Rajeev Kapur Jason Hunter



Intelligence as a Commodity

Human Judgment as a Competitive Advantage

WILEY

Table of Contents

- Cover
- Table of Contents
- Title Page
- Preface
- Part I: Aligning the C-Suite
  - Chapter 1: The Role of AI in Corporate Strategy
    - The Data Onslaught
    - The Military Mindset vs. an AI Mindset
    - Enter Artificial Intelligence
    - The AI Paradox
    - The Adoption Paradox
    - The Cautionary Tale of Nokia
    - Breaking the Mold: Healx and the Power of AI Innovation
    - Book Summary
    - The Road Ahead
    - Notes
  - Chapter 2: C-Suite Roles in AI Transformation
    - The Growth Driver
    - The Risk Guardian
    - The Innovation Catalyst
    - The Operational Excellence Champion
    - Notes
  - Chapter 3: Governance and Ethics
    - Trust Replaces Control
    - An Unprecedented Frontier
    - Self-Regulation: Leading Before Being Led

Putting the Ideas into Action
Accountability and Control Mechanisms
Risk Management and Safety
Transparency and Auditability
Bias, Fairness, and Accountability
Notes
Part II: The Strategic Playbook
Chapter 4: AI and Corporate Strategy
The Data Flood and the AI Disruption
Why Strategy Frameworks Still Matter but Must Evolve
The Experimental Nature of AI
The Need for a New Strategy Framework
AI as an Engine for Strategic Experimentation
The Three Pillars of AI-Driven Strategy
AI and Organizational Excellence
AI and Outside-In Strategy: Scenario Planning
AI and Inside-Out Strategy: Creating Markets
The Art of Strategic Balance
Notes
Chapter 5: AI and Corporate Strategy: HorizonAI
Guardrails, Scaffolding, and Jungle Gyms
Past, Present, Future: The Power of Scenarios and AI
HorizonAI: AI-Driven Scenario Planning
Notes
Chapter 6: AI and Corporate Strategy: HealthCheckAI
The Three Organs of HealthCheckAI
The Four Domains of HealthCheckAI Implementation
The HealthCheckAI Advantage
AI-Driven Operational Excellence in Finance
Challenges Faced by Finance and Risk Executives

Adaptive Risk Intelligence Architecture (ARIA)
Strengthening Internal Risk Defenses (AI + Human Oversight)
Sharpening External Risk Radar (AI + Human Strategy)
Human Governance and Controls
Implementing the ARIA Framework: A Step-by-Step Guide for CFOs and CROs
Notes
Chapter 7: AI and Corporate Strategy: Beyond Blue Ocean Framework
AI Beyond Blue Ocean Framework
Five AI-Powered Strategy Tools
Notes
Part III: Leading the AI-Powered Organization
Chapter 8: What to Know About AI
Executives and Technology: What You Need, and Don't Need, to Know
C-Suite Leaders: What Executives Need to Know and See
What Do We Know? What Do We Not Know?
A Short History of Artificial Intelligence
Understanding Artificial Intelligence as a Living System
The Circulatory System of AI: Data Collection, Flow, and Utilization
The Algorithmic Heart
The Perspective Models of AI
The Integrated Intelligence Architecture
The Living Enterprise
Our Prompt Problem
Notes

Chapter 9: The Agentic Enterprise: When Software Becomes Staff
- Yahoos, Horses, and Centaurs
- What Agents Actually Are (In Plain English)
- The Promise and the Peril
- The Human Premium
- When Everyone Has Agents: The New Competition
- Three Frameworks, One System
- Framework 1: The Agent Workforce Audit
- Framework 2: The Agentic Decision Framework
- Framework 3: The Trust Dial
- One Company, Three Frameworks: The Kooda Systems Story
- The Three Questions
- Notes

Chapter 10: The AI Leadership Blueprint: From Insight to Implementation
- When Humans Wag the AI Dog
- The AI Leadership Operating System
- Rethinking the Strategic Human-AI Partnership
- AI in the Boardroom: Precision, Timing, and Strategic Dialogue
- The AI Leadership OS Canvas: A Practical Health Check
- Leading over Adapting: The AI Organizational Operating System
- The AI Organizational Operating System
- The Constraints of Traditional Structures
- Designing the Organization for the AI Literate World
- Creating Your AI Organizational Operating System
- Collective Intelligence Framework: Human-AI Collaboration

The Collective Intelligence Framework
Core Components of the Collective Intelligence Framework
Long-Term Benefits and Organizational Impact
Your Executive AI Leadership Roadmap
Notes
Chapter 11: AI in the Next Five Years: Navigating Strategic Uncertainty
Limitless Information/Limited Knowledge: The AI Paradox
The Next Five Years: Signal, Not Noise
Five Bets for The Next Five Years
The New 80/20 Rule
Data as the New Balance Sheet
The Last Generation of Human-Only Managers
The Return to On-Premises and Open Source
A New Age of Enlightenment
One Last Bet
The Centaur's Return
Notes
About the Authors
Index
Copyright
Dedication
End User License Agreement

List of Illustrations

Chapter 5
Figure 5.1 Strategy/innovation matrix
Figure 5.2 Control paradigms

Figure 5.3 HorizonAI Framework
Figure 5.4 HorizonAI Framework: Observe
Figure 5.5 HorizonAI Framework: Attack
Figure 5.6 HorizonAI Framework: Orient
Figure 5.7 HorizonAI Scan: 2 × 2 matrix
Figure 5.8 HorizonAI Scan: Financial Services and AI
Figure 5.9 HorizonAI Methodology: Adapt
Chapter 6
Figure 6.1 HealthCheckAI
Figure 6.2 HealthCheckAI: Omnidirectional Sensing
Figure 6.3 HealthCheckAI: Adaptive Simulation
Figure 6.4 HealthCheckAI: Continuous Reconfiguration
Chapter 7
Figure 7.1 “As Is” Strategy Canvas: Overview
Figure 7.2 “As Is” Strategy Canvas: Dingo
Figure 7.3 Buyer Experience Audit: Overview
Figure 7.4 Dingo’s Buyer Experience Audit (2009)
Figure 7.5 Dingo’s Six Paths Framework
Figure 7.6 ERRC Grid: Dingo
Figure 7.7 Dingo “To Be” Strategy Canvas
Chapter 8
Figure 8.1 The new gin trolley
Figure 8.2 Conventional AI
Figure 8.3 Data, flow, and utilization
Figure 8.4 The algorithmic heart
Figure 8.5 The perspective models
Figure 8.6 The prompt problem

Chapter 9

Figure 9.1 How AI agents work: five aspects of AI agents

Figure 9.2 Managing a digital workforce: closing the loop on the new economics...

Figure 9.3 Agentic Workforce Audit: assess your agentic AI maturity level

Figure 9.4 Agentic Workforce Audit: outcomes, or expensive AI hype?

Figure 9.5 Agentic workforce audit: customer support platform example

Figure 9.6 The Agentic Decision Framework: how to assign the right decisions

Figure 9.7 The Agentic Decision Framework: customer support example

Figure 9.8 The Trust Dial: know when to tighten, hold, or expand autonomy

Chapter 10

Figure 10.1 Zones of influence: mapping the lay of the land

Figure 10.2 Zones of influence: example

Figure 10.3 Timing and consistency protocol: establishing clear intervals for ...

Figure 10.4 AI Strategic Insight Board: online, interactive platform enabling ...

Figure 10.5 AI risk heat map: ensuring board members clearly understand AI ris...

Figure 10.6 AI risk heat map: example

Figure 10.7 Consistent and committed rituals: establishing institutional pract...

Figure 10.8 AIOS Canvas: a pragmatic visual monitoring framework

Figure 10.9 Collective Intelligence Framework: designing human-AI collaboratio...

Figure 10.10 Collective Intelligence Framework: example

Intelligence as a Commodity

Human Judgment as a Competitive Advantage

Rajeev Kapur

Jason Hunter

WILEY

Preface

This book examines judgment and partnership, emphasizing our conviction that, even as machines advance, people remain a business's most scarce and valuable resource. Driven by this belief and our collaboration, we explored these ideas and wrote this book.

We met in 2015 at 1105 Media. For ten years, we have sat on the same side of the table, wrestling with data and strategy. We returned, again and again, to the question animating this book: now that intelligence can be bought, what is left that only people can do? This book comes from our partnership, which taught us more about that question than any algorithm could. We were not alone.

No journey is easy, and no decade of collaboration is filled only with agreement. Still, every day we feel energized by our mission to learn and improve. We believe passionately in the ideas in this book. They are not for executives whose only ambition is to ride out the AI wave, adopt whatever their competitors adopt, and survive. That has never been our interest. If you can be satisfied with that, do not read on. But if you want to lead, build an organization where artificial intelligence and human judgment each do what they do best, and where customers, employees, shareholders, and society win—read on. We are not saying it is easy, but it is worthwhile.

Our experience confirms that there is no permanent advantage in technology. There are no permanently excellent companies. Every model your organization buys, your competitors can buy too. For the first time in history, every organization on Earth has access to the same intelligence. What cannot be commoditized is judgment: knowing where to deploy AI and having the nerve to leave it alone. That is what we call the human premium. We have found that cultivating it is the strategic move that matters most.

This book challenges companies to close the gap between AI's promise and its frequent underperformance. Our central tenet is that the gap is not technical. It is psychological, organizational, and strategic. In the pages ahead, we explore these challenges. We

provide practical and proven frameworks and methodologies to overcome them.

Our aim is to make leading with AI as systematic and actionable as any discipline executives already command. Only then can companies pursue transformation in a way that is both opportunity-maximizing and risk-minimizing. No company, large or small, can afford to gamble its future on enthusiasm alone.

The ideas in this book were not developed in a lab. They were built over decades in data science, tested in boardrooms, and refined through failures we'll share honestly. They draw not only from the analytics of AI-driven strategy, but also the all-important human aspects: trust, governance, accountability, and the judgment to act when the machine tells you something your experience says is wrong are all crucial.

The human premium has always existed. As AI commoditizes intelligence, the value of human judgment grows. We believe this expansion is the root of the next era of growth. Yet, there is a poor understanding of how to systematically cultivate and capture it, both in theory and in practice. We invite you to read this book and learn how to lead your future.

Part I
Aligning the C-Suite

Chapter 1
The Role of AI in Corporate Strategy

> *"We have stone-age emotions, medieval institutions, and god-like technology."*
>
> —Edward Wilson

On April 13, 1929, Salomon Shereshevsky, a young reporter for a Moscow newspaper, was summoned to his editor's office. This was rarely a good sign in Stalinist Russia. The editor had noticed something unusual: Shereshevsky never took notes during interviews. When questioned, the reporter claimed he didn't need to. He remembered everything. To prove it, he repeated every word from that morning's editorial meeting verbatim. By afternoon, Shereshevsky was under examination at the Academy of Communist Education, where Soviet psychiatrists would study his memory for the next 30 years. He could recall entire books, complex mathematical formulas, even essays in languages he didn't speak.

But his extraordinary ability came at a cost.

Shereshevsky's world was one of perfect retention and absolute fragmentation. He could recall infinite details yet struggled with abstraction. He could not connect ideas, foresee consequences, or form new insights. The sheer volume of information became a cognitive burden rather than an asset. He could store knowledge but could not use it to think. His mind was, in effect, a hard drive with no operating system. An infinite database with no algorithm to make sense of it.

For decades, companies operated in Shereshevsky's image—capturing everything, understanding little. AI changed that. Modern AI systems don't merely store and retrieve. They reason, synthesize, and surface insights no human team could reach alone. For the first time, the operating system Shereshevsky never had exists. And yet most organizations are squandering it. Not because the intelligence isn't good enough, but because their structures, cultures, and decision-making frameworks were built for a world where insight

was scarce. The assumption behind most AI adoption is that deploying intelligence produces better outcomes. It doesn't, at least not automatically. Without the organizational capacity to act on what AI reveals, more intelligence just produces more noise, faster. The companies pulling ahead are not the ones processing the most data or deploying the most tools. They are the ones who have figured out when, how, and where to apply AI to their strategies and, importantly, when to ignore it.

This is the argument of this book: AI gives every organization on Earth access to the same intelligence. The advantage now belongs to those who know where to use it and have the nerve to leave it alone.

The Data Onslaught

To examine AI's role in business, we must first understand how data and information flow in modern companies. In traditional hierarchical organizations, information serves as the backbone for maintaining control and coordinated action. It also supports strategic decision-making. Typically, information flows vertically from senior leadership down. This reinforces authority and guides employees through selective dissemination and clear directives.

At higher levels, information becomes a critical foundation for strategic choices. Managers distill raw data into actionable insights, transforming it into coherent, contextualized information through analysis and interpretation. This enables informed decisions. Information also provides measurable benchmarks for performance assessment, supporting accountability and efficiency.

Today, the data problem is an order of magnitude worse.

If we compressed recorded human history into a single calendar year, writing would appear on December 25. The printing press would arrive at 11:40 p.m. on December 31, and the television at 11:59:32 p.m. The Internet would come just 59 seconds before midnight, and social media 10 seconds before the year ends. And that's just the distribution infrastructure. Now consider what's flowing through it.

Global data volume grew from 2 zettabytes in 2010 to 149 zettabytes in 2024. That's a 7,350% increase in 14 years.[1] If you stored 2024's data on DVDs, the stack would stretch 23 million miles into space, or nearly a quarter of the way to the sun. Despite this, a Forrester study found that 60–73% of enterprise data goes unused in decision-making.[2]

The human brain was not designed for this. The average worker processes about 34 gigabytes of information a day.[3] That equals thousands of pages of content. Microsoft research links this overload to an average attention span of eight seconds. Meanwhile, McKinsey reports that employees spend 28% of their work week (over 11 hours) just managing email. Decision fatigue is not a theory; it's a daily reality.[4]

Like a frog in slowly boiling water, most organizations have adapted to the data deluge incrementally without ever recognizing it as the crisis it is.

The Military Mindset vs. an AI Mindset

Walk into almost any corporation today, and you're walking into an organizational structure designed by nineteenth-century Prussian generals.

The language gives it away: headquarters, officers, frontline workers, strategy, tactics, chain of command. After World War II, as military leaders flooded into corporate America, they brought the command-and-control philosophy that had won the war. It worked brilliantly for the industrial age with standardized processes, clear lines of authority, and predictable outputs.

But here's the paradox: that structure may now be a corporation's greatest liability.

Military hierarchies move information up and orders down. AI thrives on rapid, horizontal information flows, the antithesis of a hierarchical business organization chart. Think of it this way: a military unit needs strict information control in combat. An AI-driven organization needs information to flow like blood, freely to where the organism most needs it.

The failure rates tell the story. A stunning 74% of companies have shown no tangible value from their AI implementations.[5] More striking: a 2025 MIT Media Lab study found that 95% of enterprise AI pilots fail to reach production with measurable value.[6] And the spending patterns partly explain why roughly 50–70% of AI budgets have been poured into sales and marketing pilots, the lowest-hanging fruit, rather than into structural transformation.

Executives approach AI like generals planning a campaign: controlling risk, managing lines of authority, and measuring ROI on discrete projects. For routine tasks, such as chatbots, basic machine learning tools, and marketing automation, this approach yields results. For an organization-wide competitive advantage, it fails.

The "70-20-10 principle" is instructive: successful AI implementation requires 70% focus on people and processes, 20% on technology and data, and only 10% on algorithms. Most organizations invert this ratio by obsessing over technical solutions while underinvesting in the human element.

The result is a cultural split. Junior employees race toward the AI future. Executives retreat into familiar hierarchies. And the organization stands still, burning both money and opportunity; 75% of executives consider AI a strategic priority, while 89% say their company has an AI strategy.[7] Companies with clear AI strategies are twice as likely to see revenue growth and achieve 2.1× higher ROI.[8] The data is clear. What is unclear is how to systematically achieve this potential.

Enter Artificial Intelligence

As data peaked, a convergence of **three** forces led to the launch of ChatGPT in 2022. Understanding them matters.

> **Data.** "Big data" is AI's raw fuel: text, images, audio, numerical datasets, and, increasingly, "world models." Without data, AI is a sports car with no gasoline. In the coming years, models will train on multisensory inputs, robotics-generated experience, and cross-modal understanding. The tank will keep filling.

Processing power. No one predicted that gaming cards would play a crucial role in the development of modern AI. Alex Krizhevsky, a graduate student at the University of Toronto, built a prototype AI system in his bedroom using two improvised GPUs from a then-obscure company called Nvidia. GPUs, originally designed for rendering high-definition game graphics, proved to be well-suited for the parallel processing required by deep neural networks. An industry driven by teenage gamers unintentionally provided the engine that powered AI's progress.

Algorithms. Simultaneously, advancements in large language models (LLMs) revolutionized AI's understanding and generation of human language. These architectures dramatically improved contextual comprehension and language generation capabilities, transforming how machines interact with human text. Combined with massive datasets and exponentially faster hardware, these advances set the stage for OpenAI's release of ChatGPT in 2022.

The AI Paradox

Considering this evolution, it's important to address the rising expectations surrounding AI. Here is the paradox. Very few firms are truly mature in AI, yet individual adoption has spread at an extraordinary pace. Most organizations are still experimenting, bolting tools onto functions, and calling that transformation. They are not wrong to move early. They are wrong if they think early use equals institutional capability.

AI matters for three reasons.

First, AI changes the economics of knowledge work. In customer service, product development, and coding, AI increases throughput. This scale makes the old productivity baseline look slow. Instead of slight annual improvements, tasks that once took hours now take minutes. Executives feel both excitement and confusion as familiar assumptions about capacity, staffing, and cycle time no longer hold.

Second, it improves decisions, not just execution. In domains such as fraud detection, diagnostics, and scientific research, AI is proving valuable because it can detect patterns, surface anomalies, and expand analytical reach beyond what rule-based systems or unaided experts can manage alone. The point is not that machines replace judgment. The point is that judgment, when paired with machine assistance, can become faster, broader, and in some cases more accurate.

Third, efficiency is not the only benefit. Cost savings matter but are just the shallow end. The greater opportunity is in creating products, services, and business models that were once too expensive, too complex, or too slow. The most important question is not whether AI can automate the business. It is whether AI can make a different business possible.

This is where the hype begins and the confusion persists. Leaders hear large numbers and think it's a technology story. In reality, this is an organizational story. AI does not simply improve single tasks. It can change the very shape of the firm. Remember, this is the least powerful AI there will ever be.

The Adoption Paradox

Here's where it gets interesting.

Individual Adoption Thrives

A 2025 study by Management Science finds that 45% of Americans use generative AI at work.[9] About 11% use it daily. Surveys suggest AI use is widespread across organizations, but official U.S. business-use rates remain much lower (4%) when measured as actual use in the production of goods or services.[10] The gap is not technology; it's institutions. Employees act as "secret cyborgs," quietly adopting AI to boost personal productivity. Some see 25% gains in output from tools like ChatGPT, Claude, and GitHub Copilot.[11]

Organizational Adoption Lags

Workers are reluctant to share their use of AI with managers. It is, after all, their edge. A Slack survey found that 48% of desk workers felt uncomfortable disclosing AI use to their managers.[12] This tension is older than technology. From Prometheus to John Henry, workers have always feared being displaced by new tools. But today the dynamic has inverted: it's the workers racing ahead while leaders hang back.

The fear is mutual. Below, the worry is job security. Above, leaders sense a deeper loss of control. Knowledge is power, and AI threatens to shift it. MIT Sloan research shows 76% of executives believe organizations need transparency for AI success. Only 24% have taken meaningful steps toward this. Meanwhile, 45% fear losing control if they flatten their hierarchy.[13]

The gap between strategy and execution is stark. Executives cite budget restrictions (41%), talent shortages (38%), and technical complexity (35%) as their primary obstacles.[14] Meanwhile, 83% call AI a strategic priority.[15] A fascinating bottom-up power shift is underway: it's workers, not executives, driving AI adoption and compelling leaders to rethink how organizations operate.

The evidence of that shift: 71% of leaders now prefer hiring less experienced candidates with AI skills over veterans without them. Two-thirds say they wouldn't hire someone lacking AI capabilities at all.[16] While 40% of employees see IT departments leading AI implementation, business leaders register at just 23%.[17]

The Cautionary Tale of Nokia

Nokia is the canary in the coal mine. In 2007, Nokia held 50% of the global smartphone market share.[18] Their structure was the problem: a rigid hierarchy where critical intelligence from frontline employees never reached decision-makers. When employees recognized the iPhone's potential, each management layer filtered and softened the message. By the time it reached the top, the threat had been minimized. Five years later, their market share was 4.3%.[19]

Contrast that with Microsoft under Satya Nadella. He dismantled Steve Ballmer's rigid silos. The result: Microsoft's market cap grew from $300 billion to more than $3 trillion, surpassing Apple as the world's most valuable company.

Three lessons follow:

> **Speed:** When Intel faced manufacturing challenges in 2020, AMD's flatter structure enabled faster decision-making. AMD captured critical market share. Intel still hasn't recovered.
>
> **Innovation:** 71% of innovations emerge from cross-functional collaboration. Hierarchical structures kill it. Google's "20% time" policy produced Gmail, AdSense, and Google News.
>
> **Risk Management:** After 2008, JPMorgan Chase introduced "town halls" where junior employees could flag risks directly to senior management—a structure that helped them navigate the 2020 COVID crisis better than most peers.

The key insight: C-suite executives don't lose control in the AI world by flattening structures. They gain competitive advantage. Information no longer flows at the speed of hierarchy. It flows at the speed of trust.

Breaking the Mold: Healx and the Power of AI Innovation

Drug development follows strict rules: $2–$3 billion in costs, 12–14 years of research, and a 95% failure rate. Of roughly 10,000 known human diseases, the pharmaceutical industry has approved treatments for only 500–600, about 5–6%. Behind these numbers are more than 400 million rare-disease patients waiting for a cure.

In 2014, a desperate father named Nick Sireau stood before two scientists with a story that would challenge those laws. His sons suffered from a rare genetic disorder. For a decade, he had searched for treatment. Traditional drug development was not just slow and expensive; it was structurally broken for rare diseases.

Dr. Tim Guilliams and Dr. David Brown—the scientists who heard Sireau's story—saw what others had missed. Brown, co-inventor of Viagra, knew that breakthroughs often come from unexpected angles. What if the problem wasn't finding new drugs but seeing existing ones differently? What if AI could help escape the tyranny of conventional drug discovery?

That meeting gave birth to Healx. Rather than following the 80/20 rule that concentrates pharma investment on blockbuster diseases, Healx flipped the script by focusing on the long tail of rare conditions, individually small but collectively enormous.

Their AI platform, Healnet, analyzed vast amounts of biomedical data to find new uses for existing drugs. What once took years now takes months, and at a fraction of the cost. Working with patient groups, Healx identified potential treatments for Fragile X syndrome, the most common genetic cause of autism. Within 24 months—a blink in pharmaceutical time—they had several therapies ready for clinical testing. Their approach was 80% faster and 90% cheaper than traditional methods.

From a Cambridge startup, Healx has grown to more than 130 people and raised more than €63 million. More importantly, they've given hope to millions of patients who traditional pharma had systemically ignored. Their goal: 100 rare disease treatments advancing toward clinical trials by 2027. Not just a corporate target but a declaration that technology can serve both the profitable many and the forgotten few.

The chapters that follow lay out the frameworks to do exactly that.

Book Summary

This book confronts the gap between AI's sweeping promise and its frequent underperformance. The challenges are not primarily technical; they are human, organizational, and strategic. In this book, we distill the principles and analytical frameworks necessary for decisive leadership in this evolving landscape.

Part I: Aligning the C-Suite (Chapters 2 and 3) examines what AI demands of executive leadership, including consequences when leaders resist. Chapter 2 introduces four executive archetypes, each with distinct responsibilities for integrating AI. Chapter 3 addresses the lack of regulation, requiring enterprises to develop their own accountability and fairness frameworks. Together, these chapters show that AI transformation starts or stalls at the top.

Building on leadership alignment, **Part II: The Strategic Playbook** (Chapters 4–7) delivers actionable frameworks. Chapter 4 presents a strategic model grounded in Organizational Excellence, Outside-In, and Inside-Out pillars. Chapter 5 details HorizonAI's scenario planning cycle. Chapter 6 introduces enterprise-level operational and risk frameworks with HealthCheckAI and ARIA. Chapter 7 applies Blue Ocean Strategy enhanced by AI. These chapters empower leaders to redefine strategy, rigorously test decisions, and capture new markets. The core principle: AI augments human insight; it is not an oracle.

Part III: Leading the AI-Powered Organization (Chapters 8–11) shifts from strategy to leadership itself. It builds executive AI fluency (Chapter 8), reframes leadership in terms of leading an organization where the staff is software (Chapter 9), offers blueprints for structuring AI-powered organizations (Chapter 10), and looks into the future of C-suite leadership over the AI horizon (Chapter 11).

As AI agents become commoditized, the book argues that human judgment, ethics, and institutional wisdom (the Human Premium) become the decisive competitive advantage (Chapter 11). The final chapter calls on the C-suite to lead with foresight, ethical conviction, and humility to co-create a future shaped by human-AI collaboration.

The Road Ahead

Edward Wilson's epigraph, "stone-age emotions, medieval institutions, god-like technology," captures the tension at the heart of this book. The gap between AI's potential and AI's realized value is primarily a leadership problem. Consider what happened in April 2026, when Anthropic unveiled Claude Mythos, a model that autonomously discovered thousands of critical software vulnerabilities, including flaws that had survived decades of human review and millions of automated security scans.

Rather than release it, Anthropic restricted access and convened a defensive coalition of Amazon, Apple, Google, Microsoft, and a dozen others under Project Glasswing. The technology didn't make that decision. People did. That is precisely the duality this book asks executives to confront: the same capability that secures critical infrastructure could devastate it. The human advantage is not in building the tool. It is in having the judgment to decide what to do with it.

This pattern holds across every example in this book. It's not the technology, the data, or the talent. It's the moment when AI tells you something your experience says is wrong and you must decide whether your organization will act on the uncomfortable answer or become the next Shereshevsky: remembering everything, understanding nothing, and calling it a strategy.

Notes

1. Statista, "Volume of Data or Information Created, Captured, Copied, and Consumed Worldwide from 2010 to 2029 (in Zettabytes)," November 14, 2025, https://www.statista.com/statistics/871513/worldwide-data-created.

2. Mike Gualtieri and Noel Yuhanna, The Forrester Wave™: Big Data Hadoop Distributions, Q1 2016 (Cambridge, MA: Forrester Research, January 19, 2016) .

3. Jeffrey R. Young, “Human Capacity for Information Is Massive but Finite,” The Chronicle of Higher Education, January 28, 2011.

4. Ellen Plummer, “How to Spend Way Less Time on Email Every Day,” Harvard Business Review, February 6, 2019.

5. Nicolas de Bellefonds, Tauseef Charanya, Marc Roman Franke, Patrick Forth, Michael Grebe, Romain de Laubier, Vlad Lukic, Amanda Luther, Clemens Nopp, and Joe Sassine, “Where’s the Value in AI?” Boston Consulting Group, October 2024, https://www.bcg.com/publications/2024/wheres-value-in-ai.

6. MIT Nanda, Aditya Challapally, Chris Pease, Ramesh Raskar, and Pradyumna Chari, The GenAI Divide: State of AI in Business 2025 (July 2025), PDF, https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf.

7. Jessica Apotheker, Sylvain Duranton, Vladimir Lukic, Nicolas de Bellefonds, Sesh Iyer, Olivier Bouffault, and Romain de Laubier, “From Potential to Profit: Closing the AI Impact Gap,” BCG AI Radar, Boston Consulting Group, January 15, 2025.

8. Apotheker et al., “From Potential to Profit.”

9. Alexander Bick, Adam Blandin, and David J. Deming, “The Rapid Adoption of Generative AI,” Management Science, 2026, https://doi.org/10.1287/mnsc.2025.02523.

10. Kathryn Bonney, Cory Breaux, Catherine Buffington, Emin Dinlersoz, Lucia Foster, Nathan Goldschlag, John Haltiwanger, Zachary Kroff, and Keith Savage, “Tracking Firm Use of AI in Real Time: A Snapshot from the Business Trends and Outlook Survey,” U.S. Census Bureau Center for Economic Studies Working Paper CES-WP-24-16 March 2024, https://www.census.gov/library/working-papers/2024/adrm/CES-WP-24-16.html.

11. Michael Chui, Eric Hazan, Roger Roberts, Alex Singla, Kate Smaje, Alex Sukharevsky, Lareina Yee, and Rodney Zemmel, “The

Economic Potential of Generative AI: The next Productivity Frontier," McKinsey & Company, June 2023.

12. Slack Workforce Lab. "The Fall 2024 Workforce Index: Executives and Employees Investing in AI, but Uncertainty Holding Back Adoption." Slack, November 2024. https://slack.com/blog/news/the-fall-2024-workforce-index-shows-executives-and-employees-investing-in-ai-but-uncertainty-holding-back-adoption.

13. Elizabeth M. Renieris, David Kiron, and Steven Mills, "Artificial Intelligence Disclosures Are Key to Customer Trust," MIT Sloan Management Review, September 24, 2024.

14. Ryan Daws, "Tech Executives Confident in AI Skills, but Adoption Barriers Persist," *AI News*, July 11, 2024.

15. Infomineo, "How CEOs Leverage AI for Smarter Decision Making," Infomineo (blog), June 24, 2024, https://infomineo.com/blog/how-ceos-leverage-ai-for-smarter-decision-making.

16. Microsoft and LinkedIn, "AI at Work Is Here. Now Comes the Hard Part," *2024 Work Trend Index Annual Report*, May 8, 2024, https://www.microsoft.com/en-us/worklab/work-trend-index/ai-at-work-is-here-now-comes-the-hard-part.

17. Josh Howarth, "54 New Artificial Intelligence Statistics," Exploding Topics, January 7, 2025.

18. Nokia, "Nokia Q4 2007 Net Sales of EUR 15.7 Billion, EPS of EUR 0.47 (EUR 0.47 Excluding Special Items)," press release, January 24, 2008, 1, https://www.nokia.com/system/files/files/q4-2007-earnings-release-pdf.pdf.

19. John Paczkowski, "Microsoft's Nokia Deal by the Numbers," AllThingsD, September 2, 2013, https://allthingsd.com/20130902/microsofts-nokia-deal-by-the-numbers.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 2
C-Suite Roles in AI Transformation

> *"We tend to overestimate the effect of a technology in the short run and underestimate the effect in the long run."*
>
> —Roy Amara

In late 2023, senior leaders at DB Schenker, one of the world's largest logistics providers, spotted a troubling pattern. Despite digitizing much of their global freight network, operational bottlenecks still created cascading delays at key European hubs. Meanwhile, a regional competitor had quietly started using AI to reroute shipments based on real-time traffic data, customs wait times, and labor strikes. With higher value and fewer pain points, customers jumped to DB Schenker's competitor.

As Maik Schmidt, technology architect at DB Schenker, put it, "On the one hand, I could see AI disrupting several parts of Schenker's business. Many processes will be much more efficient than they are today thanks to software solutions that were only a dream a few years ago."[1]

DB Schenker responded by rethinking how AI could speed up decision-making across the entire network. Within six months, the company launched a real-time supply chain platform connecting predictive AI with human dispatchers to enhance agility and accountability. This story illustrates how AI doesn't just change what companies can do but also the social and organizational changes required and how quickly they must act. For the C-suite, the priority is to rethink business processes just as much as to adopt new technology.

This pattern of underestimation is as old as the wheel. When the telephone arrived in the late 1800s, Western Union's president dismissed it as having "too many shortcomings to be seriously considered as a means of communication." Similarly, factory owners balked at electricity, convinced that conversion costs outweighed benefits. Decades later, in 1998, economist Paul Krugman predicted

the Internet's economic impact would prove "no greater than the fax machine's." In each case, leaders told themselves they had more time to adapt than they did.

What makes AI different is the speed of adoption, the scale of capital behind it, lower barriers to entry, and the rising opportunity cost of waiting. The telephone took 75 years to reach 50% adoption in the United States. Personal computers took about 25. ChatGPT, in comparison, reached 100 million users in two months and 50% adoption in the US in less than 3 years.[2] This kind of compression leaves far less runway for a "wait and see" approach.

When early adopters deploy AI, their systems learn from every interaction, accumulate proprietary data, and create feedback loops that strengthen their position over time. A six-month delay is not just a gap; it means entering a race where your competitor's AI has been continuously optimizing. The talent divide also widens, as early movers build teams experienced in working alongside AI, while late adopters struggle to find talent already familiar with advanced operations.

The cost of delayed AI adoption isn't linear. It's exponential. Each month of delay doesn't just affect technology adoption. It affects organizational capability, culture, and attractiveness to top talent. Late adopters risk falling behind in an increasingly expensive game of catch-up. When today's executives dismiss generative AI as "just a sophisticated autocomplete" or insist that "AI can't replace human judgment," they unknowingly echo Western Union's president and his "electrical toy." The language of resistance to technological change stays remarkably consistent, even as the technology evolves.

A broader lesson surfaces here, one we'll return to throughout this book: the exploration-extraction dilemma. This is an old machine-learning concept that refers to the trade-off between exploring new options to discover potentially better solutions and leveraging known strategies that have proven successful. Striking the right balance is the challenge: too much exploration risks inefficiency, while too much extraction risks missing out on superior outcomes.

Most work on AI and business strategy has focused heavily on the extraction side. Efficiency gains, faster processes, lower costs, fewer

errors. As we will see throughout this book, AI has massive untapped potential for exploration, identifying markets hidden in plain sight and waiting for the right lens to bring them into focus. Strategic alignment between AI as a technology and business as a social construct is often overlooked. This chapter highlights the key challenges C-suite leaders face, shows how AI offers solutions, and outlines immediate steps leaders can take.

Four executive archetypes define the framework. While job titles often overlap, emphasizing leadership style makes this approach more practical than traditional organization charts.

The Growth Driver is focused on market expansion and revenue generation. Think chief executive officers (CEOs), chief revenue officers (CROs), and chief marketing officers (CMOs). They harness AI to identify nascent market opportunities by analyzing vast datasets of consumer behavior, competitor movements, and macroeconomic indicators in real time. They deploy machine learning to optimize customer acquisition costs and predict lifetime value, enabling rapid scaling in high-potential segments.

The Risk Guardian emphasizes stability, compliance, and risk management. Typically, CFOs and CROs are more conservative in decision-making, prioritizing protection over growth. They employ AI-powered anomaly detection to monitor transactions, communications, and operational patterns for potential compliance violations or security threats. They leverage predictive models to forecast market volatility and regulatory changes, proactively adjusting strategy and risk controls.

The Innovation Catalyst is focused on longer-term growth through technological advancement and R&D. These traits are common in chief technical officers (CTOs) and chief innovation officers (CIOs). They use AI to accelerate ideation, rapidly generating and evaluating thousands of potential solutions. They implement AI-driven innovation management platforms that predict which R&D projects are most likely to succeed in the market.

The Operational Excellence Champion focuses on efficiency, process optimization, and cost management. Often, chief operating officers (COOs) and chief supply chain officers (CSCOs) use AI to simulate and optimize complex operational processes before implementing changes. They harness machine learning for predictive maintenance and dynamic resource allocation, dramatically reducing downtime and improving efficiency.

This chapter introduces the Growth Driver, Risk Guardian, Innovation Catalyst, and Operational Excellence Champion archetypes, covered in the following chapters. This structure guides leaders through actionable strategies tailored to each focus area. In each section below, we examine each archetype's mindset, AI adoption strategies, and how to implement and scale these strategies across the organization.

The AI-Fluent Executive

Every day, executives make decisions that could determine their company's fate, yet they are saddled with imperfect information, biased perspectives, and not enough time. Approximately 85% of business leaders report experiencing significant decision stress, while three-quarters say the volume of decisions they must make has increased tenfold over the last three years.[3]

What if those constraints are no longer immutable?

The collision between AI and executive judgment demands a new model of leadership, what we call the *AI-fluent executive*. Consider how PepsiCo is rearchitecting the way it makes decisions entirely. While most companies still react to supply chain disruptions after they occur, PepsiCo has built a system designed to resolve crises before they exist. In partnership with Siemens and Nvidia, the company is constructing digital twins of its warehouses, conveyor systems, and distribution networks—virtual replicas so precise that executives can stress-test supply chain configurations, identify bottlenecks, and validate new strategies without touching the physical operation. Their goal is to detect up to 90% of operational problems in simulation before they ever reach the warehouse floor. As PepsiCo's Chief Strategy and Transformation Officer Athina Kanioura put it, the foundation matters more than most leaders realize. "If we hadn't done all those changes and the moves," she said, "it would now be impossible for us to do agentic AI." This story drives home how to thrive in our entirely different decision-making environment.[4]

The AI-fluent executive operates at the intersection of human wisdom and machine intelligence in four key ways:

- **Bias reduction: the invisible hand of prejudice.** Human bias isn't just a social issue but a business liability. Studies show hiring managers form impressions in the first seven seconds. AI doesn't eliminate bias, but it can expose it. At Goldman Sachs, AI systems now monitor compliance across more than 40 countries, completing internal audits

25% faster and flagging risk patterns that previously went undetected until post-trade review. When Unilever implemented AI in its hiring process, it led to a 16% increase in diversity hires, a 90% reduction in time-to-hire, and £1 million in annual cost savings.

- **Efficiency and speed: the new economics of decision-making.** BlackRock's Aladdin platform performs 180 million option-adjusted calculations and 5,000 portfolio stress tests every week. But its real value isn't speed; it's the freedom it gives executives to focus on strategic rather than tactical decisions. Amazon's AI-powered systems have been documented as making more than 2.5 million pricing adjustments daily, freeing executives to focus on work that only humans can do. AI is creating "decision velocity" with not just faster decisions but smarter ones.
- **Augmenting human judgment: the synthesis of mind and machine.** The most sophisticated use of AI in executive decision-making is in augmentation. Bridgewater Associates' Artificial Investment Associate (AIA) analyzes data across time and geography, generating hypotheses in minutes for human experts to evaluate. Co-CIO Greg Jensen says, "The big jump here is using machine intelligence to generate the alpha, that is a leap."

  Today, the AIA serves as the primary basis for decision-making for a $2 billion fund. The system challenges assumptions and suggests alternative perspectives, creating what founder Ray Dalio calls an *idea meritocracy*.
- **Minimizing risk while maximizing governance: the new balance.** AI is rewriting the trade-off between risk and reward. At Morgan Stanley, AI gives financial advisors instant access to the firm's entire knowledge base (more than 100,000 proprietary research documents), condensing hours of manual searching into seconds. Today, 98% of financial advisor teams have adopted the tool. The importance is that AI is helping executives understand exactly where the boundaries of acceptable risk lie in a way that was impossible before.

The Growth Driver

The backlash against Airbnb had been building for a decade. Crashed hotel rates, overflowing tourist neighborhoods, rising rents. And a growing sameness: the same exposed lightbulbs, IKEA furniture, and faux wood floors left guests feeling as if they could be anywhere. The Verge, an online magazine, described the Airbnb aesthetic as the "hallucination of the normal." By 2023, growth was plateauing, and differentiation was becoming harder to maintain.

CEO Brian Chesky knew they needed a new direction. AI revealed what human inspection had missed. Through advanced analysis of booking patterns and user interactions, Airbnb discovered that unique properties (treehouses, boats, historic buildings) generated average booking values 40% higher than those of standard listings.[5] Users spent significantly more time exploring these listings, even when they eventually booked something else.

But the deeper finding was a structural misalignment in their own platform. Airbnb's search algorithm, optimized to match guests to conventional rooms efficiently, actively worked against user behavior for unique properties. Those guests weren't looking for an efficient booking process. They wanted discovery, exploration, and experiences "like a local," a higher-margin proposition hiding in plain sight. Moreover, the data showed that owners of unique properties were significantly more likely to become long-term, committed hosts, an insight that reshaped investment decisions.

Rather than treating unique properties as an interesting niche, Airbnb rebuilt its entire platform around them. The company redesigned its search and recommendation algorithms to spotlight distinctive stays and launched a "Categories" feature that uses AI to classify properties by unique characteristics automatically.

The results validated the bet. Bookings for unique properties grew 60% faster than traditional stays.[6] Guest satisfaction soared. Social media sharing of distinctive experiences tripled. The platform attracted a new segment of experience-seeking travelers who were

less price-sensitive, more focused on “total experience,” and therefore more loyal.

Understanding the Growth Driver’s Mindset

Will Rogers put it best: “Even if you’re on the right track, you’ll get run over if you just sit there.” The warning resonates with growth-focused executives who hesitate to embrace AI. On the surface, AI is the perfect tool for their goals: predicting market trends, personalizing customer experiences, and optimizing pricing strategies. Yet many Growth Drivers drag their feet.

The resistance lives in the tension between AI’s timeline and a quarterly earnings cycle. AI demands significant up-front investment and patience. Growth Drivers face relentless pressure to deliver immediate results. Picture a CRO evaluating an AI-powered pricing tool that needs six months of data training before it delivers meaningful insights. Next quarter’s targets don’t wait. That mismatch between AI’s learning curve and the demand for quick results shelves promising initiatives and creates missed opportunities.

The challenge goes deeper than timing. Many Growth Drivers excel at big-picture strategy but view AI as a technical exercise, something disconnected from their vision. When a CEO sees an AI customer segmentation model as a complex data project rather than a tool for market expansion, the opportunity dies in translation. And organizational silos make it worse. When different regions log sales data differently, an AI forecasting tool produces inconsistent predictions. The technology gets blamed for a data infrastructure problem.

Most organizations are woefully guilty of ignoring the cultural shifts required in effectively leveraging AI. Growth-driven organizations thrive on human relationships, intuition, and fast decisions. AI can feel like a threat to that identity. When a successful sales team resists an AI lead-scoring tool, they aren’t exactly rejecting technology; they’re protecting a formula that has delivered results. The introduction of AI into a high-touch, fast-paced culture is like asking a jazz band to play from sheet music. The notes might be better, but

you lose the improvisation that made the performance great in the first place.

The path forward requires a shift in perspective. Growth Drivers need to reframe AI as an amplifier of human capabilities, not a replacement for them. This means aligning AI initiatives with immediate growth objectives, ensuring quick wins alongside long-term transformation, and actively championing AI as a tool that enhances rather than diminishes human expertise. Growth Drivers must manage the short-term costs of AI transformation alongside the strategic advantage AI will provide.

The Growth Driver's Playbook

We will introduce the tools and frameworks in Chapters 5–7. The first challenge for any growth-focused executive is knowing where to play. Traditional market identification relied on historical data, gut feel, and lagging feedback. AI has fundamentally transformed this landscape through sophisticated predictive analytics that process vast amounts of real-time data to identify high-growth markets with unprecedented accuracy.

Predictive Analysis

AI-powered sentiment analysis detects subtle shifts in customer preferences months before they show up in sales reports. By analyzing millions of social media posts, reviews, and service interactions, AI surfaces emerging pain points representing untapped opportunities. When demand for sustainable products begins to trend in a new region, AI flags the pattern early enough to seize the first-mover advantage. This same capability extends to competitive intelligence, enabling real-time monitoring of competitor pricing, product launches, and customer feedback. For growth executives, this is the difference between reactive and proactive strategy.

Customer Relations

Customer relationships are where AI's precision shines brightest at the moment. Modern AI segments customers using hundreds of behavioral and demographic variables, predicts lifetime values

before the company invests significant resources in acquisition, and identifies at-risk customers with accuracy rates exceeding 85%, enabling proactive intervention before churn takes hold.[7] This transforms retention from broad-brush loyalty programs to personalized strategies that address specific customer pain points.

Scenario Planning

In volatile markets, scenario planning separates growth from stagnation. AI simulation tools test thousands of market conditions, helping executives understand strategic implications before committing resources. These tools incorporate real-time data and dynamically adjust predictions. When entering a new market, AI can model financial outcomes under various competitive responses, regulatory shifts, and economic conditions, replacing static planning with adaptive strategy.

Scaling for Growth Drivers

AI also compresses innovation cycles. Trend forecasting identifies emerging needs. Digital prototyping accelerates testing. Real-time feedback loops gather and analyze customer reactions immediately after launch. This lets teams adjust strategy fast, reducing the risk of costly market misses while accelerating the path to product-market fit.

The harder problem is scaling what works. Growth Drivers who succeed with AI pilots often championed them personally. Enterprise-wide adoption demands something different: a network of embedded advocates. Some companies establish what they call “AI Kingpins,” respected leaders in each division who understand the technology’s potential and their unit’s specific needs. These Kingpins translate AI’s benefits into terms their teams care about, accelerating adoption far more effectively than top-down mandates.

The smartest scaling follows what we call an “earn to grow” model. Initial deployments focus on use cases with rapid return on investment (ROI), generating funds and credibility for broader expansion. Break deployment into 90-day sprints, each delivering measurable business results, to maintain the pace Growth Drivers expect while allowing for course corrections. One global sales

organization that followed this approach reduced implementation time by 40% while improving adoption rates with each successive wave.

Data infrastructure requires a parallel track: deploy AI solutions that work with existing data today while investing in data quality and integration for tomorrow. And cultural integration matters at scale. Organizations that succeed treat AI as a natural extension of how work gets done, not a technology project bolted on from outside. They celebrate teams that use AI independently, widely share success stories, and create formal mechanisms to cross-pollinate best practices between divisions.

Vibe coding a value creation dashboard helps here. It shows how AI implementations improve metrics at every organizational level: market share for the Growth Driver, departmental efficiency for division leaders, and personal productivity for individual contributors. When people see what AI means for their own work, adoption accelerates. The financial model for scaling also matters. Rather than treating AI as a massive capital expenditure, quick adopters fund expansion with returns from early wins, creating a self-reinforcing cycle of investment and results.

Growth Drivers have an imperative to bring their people on board and scale. AI isn't about perfect execution. It's about building momentum while developing organizational capability. The Growth Drivers who master AI scaling will find themselves leading organizations that aren't simply faster, but fundamentally more capable of creating and capturing new market opportunities.

The Risk Guardian

Here is a number that should alarm any executive: an international bank found that more than 95% of its automated fraud alerts were false positives. That was the reality at HSBC in the early 2020s. The bank screened roughly 980 million transactions a month for signs of money laundering and fraud.[8] Its rules-based monitoring system generated alerts faster than investigators could clear them. Within that noise, sophisticated criminals slipped through undetected. HSBC's risk team needed a new kind of guardian.

HSBC partnered with Google Cloud to build a machine-learning solution, code-named "Dynamic Risk Assessment," that identified suspicious behavior without relying on predefined rules. Rather than hand-coding "if-then" scenarios, HSBC's data scientists trained the AI on years of historical transactions and known fraud patterns. As new criminal tactics emerged, the team continually trained the model to recognize them. To build trust, an explanation layer was added to explain to compliance officers why each transaction was flagged, helping skeptics understand the AI's reasoning.

The impact was immediate and significant. HSBC detected two to four times more suspicious transactions while generating 60% fewer false-positive alerts.[9] Investigations that once took weeks compressed to days. Potentially illicit accounts that might have taken a month to identify were flagged in about eight days. The deployment earned HSBC Celent's "Model Risk Manager of the Year 2023" award for innovation in financial crime fighting. What began as a fraud crisis became a strategic advantage.

The lesson for the Risk Guardians is how quickly and safely they can transform without sacrificing their main goal of protecting the institution. HSBC's story demonstrates that both can be achieved with the right approach.

Understanding the Risk Guardian's Mindset

While the Growth Driver worries about missing the ship, the Risk Guardian worries about sinking it. CFOs and CROs view their role as defenders of organizational stability and custodians of enterprise value. While growth-focused executives charge forward and innovation leaders chase disruption, Risk Guardians anchor the organization with a steady hand, ensuring resilience, compliance, and financial integrity.

As described by the *New York Times*, "A.I. models are black boxes. You input a prompt, and the model outputs nearly anything: a sonnet, an image, or a legal brief riddled with lies."[10] For an executive with a deeply ingrained aversion to ambiguity, skepticism of AI makes perfect sense. But it helps to understand where that skepticism is rooted.

Five fault lines run through the Risk Guardian's relationship with AI.

First, AI defies traditional risk quantification. Historical data and proven patterns anchor conventional risk models, but AI introduces risks with no precedent and no clean probability distribution. How does one price the risk of an AI system making autonomous decisions that affect millions of customers?

Second, AI muddies accountability. In a conventional process, responsibility flows through a transparent chain of command. When an AI system produces bad outcomes, the chain of responsibility fragments among developers, data providers, and the executives who approved its deployment.

Third, AI's capacity to learn and adapt its own behavior can feel like instability to someone who prizes predictable, controllable systems. A system that evolves based on new data introduces variables that traditional governance frameworks can't pin down.

Fourth, the opacity of AI decision-making collides with the Risk Guardian's commitment to auditability. Regulators expect clear, deterministic explanations. Deep learning systems resist providing them.

Fifth, AI's interconnected architecture raises the specter of cascading failures. In traditional systems, failures stay contained within operational silos. AI creates interdependencies that allow a single error to ripple through multiple business processes faster than human operators can intervene.

These aren't irrational fears. They're the concerns of someone responsible for the organization's survival. The path forward lies in addressing them systematically rather than dismissing them. Savvy Risk Guardians don't fear AI. They fear poorly managed AI. When AI offers clarity amid complexity, control amid chaos, and measurable insight amid ambiguity, it becomes a cornerstone of their leadership.

The Risk Guardian's Playbook

The relationship between Risk Guardians and AI need not be adversarial. Consider how AI reinforces the principles Risk

Guardians hold sacred: control, transparency, and risk mitigation. When properly implemented, AI offers unprecedented capabilities across these domains. They're the ones who sleep with one eye on the spreadsheets and the other on the horizon, watching for outliers that could threaten their organizations. Today, these traditionally cautious executives are emerging as unlikely champions of innovation.

Pattern Recognition

AI's natural fit with risk management starts with pattern recognition. Traditional risk management is reactive: quarterly reviews, annual audits, and a lag between when a risk emerges and when it's caught. AI monitors transactions, communications, and operational data in real time, flagging anomalies before they grow into significant threats. The system learns from false positives and sharpens its accuracy over time, reducing the burden of investigating false alarms while surfacing subtle correlations between seemingly unrelated events that traditional methods would never catch.

The "black box" problem is shrinking. Modern AI systems incorporate detailed audit trails and layered explanations tailored to different stakeholders: technical detail for auditors, executive summaries for the board. Every AI decision can be traced, documented, and justified, often more thoroughly than human decisions.

Compliance Monitoring

These systems also autogenerate compliance reports and risk assessments, creating a continuous record of decision-making logic. Compliance monitoring extends this further. Rather than sampling transactions and relying on periodic human oversight, AI reviews every transaction, document, and communication for violations. The system adapts to new regulations in real time, updating its monitoring parameters and alerting teams to potential gaps before they become liabilities. It can also proactively identify areas where existing compliance frameworks may need strengthening, based on patterns observed across vast amounts of operational data.

Scenario Planning

Perhaps most compelling for Risk Guardians is the ability to simulate scenarios at scale. AI can stress-test systems by simulating millions of scenarios simultaneously, identifying potential cascading failures and rare tail-risk events before they surface in real operations. These simulations update as new data arrives, providing dynamic risk assessments that evolve with changing market conditions.

In practice, these capabilities cut across three domains. In financial risk management, AI monitors market conditions, transaction patterns, and economic indicators simultaneously, providing early warnings that human analysis would take weeks to assemble. In operational risk, AI tracks process variations across the organization, catching control breakdowns before they cascade into failures. In compliance management, AI ensures consistent policy application across the enterprise while maintaining the detailed audit trails Risk Guardians require.

Scaling for Risk Guardians

Scaling AI in the risk domain follows the same earn-to-grow logic, but with a distinctive character shaped by the Risk Guardian's virtues: prudence, precision, and systematic thinking. These aren't obstacles to scaling. There are advantages to doing it right. Start with lower-risk applications where AI augments rather than replaces existing controls: data analysis, pattern recognition, fraud detection. Each success strengthens the case for the next, creating an "AI scaling ladder." Movement up the ladder happens only after the lower rungs demonstrate sustained performance and reliable controls.

Governance structures matter enormously. Successful organizations establish an AI Risk Committee that combines risk management expertise, technical knowledge, and business unit leadership. This committee builds the frameworks that guide future expansion: clear accountability at every stage, designated system owners, specific triggers for human intervention, and audit trails that exceed traditional documentation standards. The most effective scaling maintains layered oversight that combines AI's analytical capabilities with human judgment.

The focus is on “earn to grow” by using the Trust Dial, discussed in Chapter 9. Start small, but think systematically. Build governance frameworks before expanding applications. Create transparent accountability chains and comprehensive audit capabilities. Most importantly, view AI scaling not as a technology initiative but as an enhancement to existing control systems.

Risk Guardians who master this approach often become unexpected advocates for AI adoption, not despite their risk management responsibilities but because of them. They bring the discipline and rigor that ensure AI implementations actually hold up under scrutiny, the very quality that makes scaled AI trustworthy across the enterprise. For organizations seeking to navigate their AI transformation, the key lies not in overriding the Risk Guardian’s concerns but in addressing them head-on. In a world defined by volatility, Risk Guardians who embrace AI will safeguard and future-proof their organizations.

The Innovation Catalyst

As we discussed, the exploration versus extraction divide has leaned heavily toward extraction: efficiency gains, faster processes, lower costs, fewer errors. Most firms use AI in ways that would be considered exploitation. But Unilever’s story illustrates a different path, one where AI changes how organizations think about innovation itself.

At its Materials Innovation Factory in Liverpool, Unilever deployed AI-powered molecular modeling and digital twins to simulate thousands of product formulations simultaneously. These systems autonomously generate tens of thousands of new product combinations, running 24 hours a day. The brilliance wasn’t bolting AI onto existing R&D. It was recognizing that AI could reinvent the entire approach to product development.

Unilever pioneered what they call *predictive sustainability*. By analyzing satellite imagery and mobile data, AI models predict areas at risk of deforestation across palm oil supply chains. Sustainability transformed from a checkbox exercise into a competitive advantage. Their DelphiAI platform analyzes product reviews, social media, and

market data to generate real-time insights for product development. The AI Horizon3 Lab in Toronto creates teams of academics, startups, and internal experts in previously unseen combinations to drive ingenuity across the business.

The implications cut deeper than efficiency. In engineering, the Iron Triangle forces a choice: Fast, Good, Cheap. Pick two. If something is good and cheap, it won't be fast. If fast and good, it won't be cheap. If fast and cheap, it won't be good. This trade-off has been treated as an iron law. Unilever's results suggest AI is dissolving it. Discovery timelines fell 40% while sustainability goals advanced simultaneously.[11] Speed, quality, and profitability no longer require painful trade-offs.

Understanding the Innovation Catalyst's Mindset

Innovation Catalysts (CTOs and CIOs) are drawn to complexity and think in patterns that mirror AI's architecture. Where others see opaque black boxes, they see emergent properties. This might suggest they would be AI's most natural champions. And on the surface, many are. But fluency is not the same as comfort. The executives who understand AI most deeply are often the ones most acutely aware of what it can eventually replace. What looks like alignment can mask an approach-avoidance dynamic between intellectual recognition of AI's potential and a quieter, rarely voiced anxiety about what that potential means for them personally. The most dangerous version of this is the Innovation Catalyst who enthusiastically deploys AI everywhere except the decisions that define their own value.

The challenge isn't understanding AI but navigating the psychological transformation it demands. Innovation Catalysts must not only advance their organizations but also process their own anxieties about AI's implications. Security and infrastructure concerns often serve as socially acceptable proxies for deeper uncertainties. Once conflict is resolved, that is, once AI proves explainable, repeatable, and safe, these leaders will champion its adoption at scale. Their success depends not on how quickly they implement but on how effectively they navigate the psychological

evolution AI requires, both for themselves and for their organizations.

The Innovation Catalyst's Playbook

Reframe AI not as a replacement but as an amplification. AI serves the innovation catalyst in three ways: (a) expanding exploration scope by enabling teams to evaluate millions of possibilities simultaneously (drug discovery testing molecular combinations, automotive design simulating aerodynamic configurations); (b) transforming experimentation economics through digital twins and simulations that reduce testing costs to near-zero; and (c) bridging knowledge domains by surfacing nonobvious connections across disparate datasets that humans would never discover.

The best implementations combine AI's computational power with human intuition. AI evaluates millions of possibilities while humans ask the right questions and make nuanced judgments about market fit. Innovation Catalysts who master this pairing will pursue strategies once considered impossible—compressing years of experimentation into months, identifying emerging opportunities before competitors, and multiplying organizational innovation capacity without a proportional increase in resources.

Scaling for Innovation Catalysts

Build "innovation intelligence networks," not just tools. Three foundations: innovation hubs as centers of experimentation, cross-functional platforms for sharing successes across divisions, and learning networks to spread practices throughout the organization. The "innovation multiplication effect" occurs when AI tools developed for one purpose find unexpected applications elsewhere. Successful scaling maintains three characteristics: clear pathways to access AI capabilities, robust feedback loops for organizational learning, and metrics that measure actual innovation impact. Each successful application makes the next easier, creating a virtuous cycle where organizational innovation capability grows exponentially.

For C-suite executives, AI is a catalyst for reimagining what's possible. Yet 60% of organizations haven't meaningfully adopted AI.

Of those that have, 65% use off-the-shelf solutions while only 43% deploy fine-tuned models.[12] This gap reflects cognitive dissonance. The CTOs and CIOs charged with leading technological advancement intellectually recognize AI's potential. But they face an approach-avoidance conflict. Leaders are simultaneously attracted to AI's possibilities *and* resistant to its implications. Security concerns and expertise gaps serve as socially acceptable reasons for slowing down to adopting AI.

The Operational Excellence Champion

Each day, UPS delivers 22.4 million parcels across a fleet of more than 100,000 vehicles.[13] Routing those drivers is a colossal puzzle. If a UPS truck starts at O'Hare Airport to deliver a package to the Sears Tower (18 miles away), it generates more than 9 billion potential shortest routes.[14] The problem isn't academic. An inefficient route wastes fuel, time, and money at a scale few companies face. UPS managers knew that shaving even a single mile off each driver's daily route could save up to $50 million per year.[15] Multiplied across the network, the opportunity was enormous, if they could solve what mathematicians call the "traveling salesman" problem, a routing challenge that has stumped experts for centuries.

UPS's answer was On-Road Integrated Optimization and Navigation (ORION). Launched in 2003, the project took a decade of methodical development. The ORION team, a mix of PhD mathematicians and veteran UPS engineers, equipped delivery trucks with sensors and collected more than 200 data points per vehicle (including speed, idling time, and whether the driver's seatbelt was buckled). ORION's software sifts through up to 200,000 possible routes for a driver's daily deliveries.[16]

The goal was improvement in practice, not perfection in theory. That distinction matters for operational leaders working with AI. When fully deployed, the results were impossible to ignore. On average, ORION cuts about 8 miles from each driver's daily route. Scaled across thousands of drivers: 10 million gallons of fuel saved annually,

$300–$400 million in cost reductions, and 100,000 fewer metric tons of CO_2 emissions.[17]

"ORION is a demonstration of the scale of innovation that can be accomplished when operations research, information technology, and business processes are seamlessly integrated. Success didn't come overnight," said Juan Perez, UPS chief information officer.[18]

Understanding the Operational Excellence Champion's Mindset

Operational Excellence Champions (COOs and CSCOs) are systems thinkers who see processes as interconnected flows rather than isolated tasks. They thrive on quantifiable clarity. AI seems tailor-made for their mindset, yet their acceptance isn't driven by excitement. It's anchored in a demand for proof. The hesitation starts with a core trait: these leaders need explainability, and many AI systems can't provide it. When an AI model flags an impending supply chain disruption but can't fully explain why, a leader whose decisions impact millions in operational efficiency feels less like they're optimizing and more like they're gambling. They aren't rejecting AI. They're rejecting unexplainable AI.

Deeper obstacles compound the resistance. Stability is non-negotiable for these leaders; introducing an AI overhaul feels like inviting chaos into a carefully balanced system. ROI timelines clash with quarterly pressures, operational AI projects require long feedback loops, and substantial upfront investment in data infrastructure is required. And that infrastructure is often a mess: fragmented legacy systems, siloed databases, inconsistent data collection. Even the most sophisticated AI is only as good as the data it ingests. Cultural resistance adds another layer. Operational teams built on repeatable processes view AI as a threat to workflows they trust. Organizational inertia can be the toughest bottleneck of all.

These leaders are pragmatists, not cynics. Their reluctance is a filter for excellence, not a roadblock. Once AI tools prove explainable, repeatable, and integrated into existing workflows, Operational Excellence Champions won't merely adopt AI, they'll champion it at scale.

The Operational Excellence Champion's Playbook

AI transforms operations in four ways: from reactive to predictive control (anticipating issues before they occur); scaling optimization beyond human limits by monitoring thousands of variables simultaneously; enabling continuous learning instead of periodic reviews; and breaking down silos by connecting patterns across domains. The mandate remains constant to optimize, reduce costs, and maintain quality, but AI provides new pathways to achieve these goals.

Success comes not from choosing between traditional methods and AI but from leveraging AI to achieve previously impossible levels of operational excellence. Those who embrace this transformation will define the next era of operations.

Their hesitation stems from a deep reliance on explicitness and predictability. When an AI model flags a supply chain disruption or recommends a workflow adjustment, it often can't fully explain "why." For a leader whose decisions impact millions in operational efficiency, relying on a tool that can't articulate its reasoning feels less like optimization and more like gambling. They aren't rejecting AI. They're rejecting unexplainable AI. Once AI tools demonstrate their value and become explainable, repeatable, and integrated into existing workflows, Operational Excellence Champions won't merely adopt AI. They'll operationalize it at scale.

Scaling for Operational Excellence Champions

Scale methodically, not rapidly. Five principles: (1) start at control points where precision matters most, creating clear metrics and natural expansion points; (2) build the data foundation first with transparent governance and quality standards; (3) layer intelligence progressively from monitoring through autonomous operation, mastering each stage before advancing; (4) standardize implementation approaches with flexible frameworks that adapt to local conditions; (5) develop human capability through systematic training and career pathways. Each principle builds on the fundamentals of operational excellence.

Success comes from a methodical approach, not speed. The goal is reliability of results, not speed of implementation. By following these principles, operations can scale AI to enhance rather than disrupt while building human capabilities for long-term success.

***

Four archetypes, four forms of resistance, yet one root cause. The Growth Driver guards quarterly gains. The Risk Guardian guards accountability. The Innovation Catalyst guards relevance. The Operational Excellence Champion guards explainability. Their resistance stems not from ignorance but from identity. Each leader's greatest strength becomes a brake on transformation, making resistance both logical and costly. The biggest mistake is to wait for AI to be perfect before trusting it, while competitors pull ahead through early, albeit imperfect, adoption.

For CEOs who embody all four orientations (the drive for growth, the instinct for risk, the urge to innovate, and the discipline for operations), the challenge is not picking a dominant archetype. Each voice is significant; none suffices alone. CEOs who lead this transformation will hold all four in tension: enough urgency to build advantage, enough discipline to manage risk, enough curiosity to explore, and enough rigor to scale success.

Notes

1. DB Schenker, "AI Interview: People Decoding the Cargo Code at DB Schenker," Logistics Matters (blog), September 23, 2024, https://www.dbschenker.com/global/insights/blog/ai-people-decoding-the-cargo-code-at-db-schenker-1721812.

2. Krystal Hu, "ChatGPT Sets Record for Fastest-Growing User Base — Analyst Note," Reuters, February 2, 2023, https://www.reuters.com/technology/chatgpt-sets-record-fastest-growing-user-base-analyst-note-2023-02-01.

3. Mark Purdy and A. Mark Williams, "How AI Can Help Leaders Make Better Decisions Under Pressure," Harvard Business

Review, October 26, 2023, https://hbr.org/2023/10/how-ai-can-help-leaders-make-better-decisions-under-pressure.

4. Lindsey Wilkinson, "How PepsiCo Moves Past AI Pilot Purgatory," CIO Dive, June 9, 2025, https://www.ciodive.com/news/PepsiCo-generative-AI-pilot-scale-strategy-agents/750095.

5. Airbnb, "Hosts of Unique Stays Earned Nearly $1 Billion in 2021," Airbnb Newsroom, July 2022, https://news.airbnb.com/hosts-of-unique-stays-earned-nearly-1-billion-in-2021.

6. Airbnb, "Hosts of Unique Stays Earned Nearly $1 Billion in 2021."

7. Uberall, "Study: 80% of Consumers Report Chatbot Experiences as Positive," PRNewswire, February 28, 2019, https://www.prweb.com/releases/STUDY_80_of_Consumers_Report_Chatbot_Experiences_as_Positive/prweb16135523.htm.

8. Gregory Calvery, "Harnessing the Power of AI to Fight Financial Crime," World Economic Forum, October 12, 2023.

9. Calvery, "Harnessing the Power of AI to Fight Financial Crime."

10. Kevin Roose and Casey Newton, "Peering Into A.I.'s Black Box + Who's the Real Techno-Optimist? + Reading Ancient Scrolls with A.I.," Hard Fork (podcast), New York Times, October 20, 2023.

11. Unilever, "How AI and Digital Help Us Innovate Faster and Smarter," Unilever News, April 4, 2023, https://www.unilever.com/news/news-search/2023/how-ai-and-digital-help-us-innovate-faster-and-smarter.

12. Scale AI, *Scale Zeitgeist: AI Readiness Report 2024*. (San Francisco: Scale AI, 2024), https://go.scale.com/hubfs/Content/Scale%20Zeitgeist%20AI%20Readiness%20Report%202024%204-29%20final.pdf.

13. UPS, "Corporate Facts: UPS," fact sheet, 2025 (data through December 31, 2024),

https://about.ups.com/content/dam/upsstories/assets/fact-sheets/ups-global/CorporateFactSheet25.pdf.

14. Sustainable Business, "UPS Advances Sustainable Operations with Big Data," SustainableBusiness.com, November 11, 2013, https://www.sustainablebusiness.com/2013/11/ups-advances-sustainable-operations-with-big-data-51995.

15. Sustainable Business, "UPS Advances Sustainable Operations with Big Data."

16. Karen Weise, "Orion: Route Optimization at UPS," Bloomberg Businessweek, February 26, 2013.

17. UPS, "UPS to Enhance ORION with Continuous Delivery Route Optimization," GlobeNewswire, January 29, 2020, https://www.globenewswire.com/news-release/2020/01/29/1977072/0/en/UPS-To-Enhance-Orion-With-Continuous-Delivery-Route-Optimization.html.

18. Horner, P. (2016). 'ORION' Delivers Success for UPS. *OR/MS Today* 43 (3): https://pubsonline.informs.org/do/10.1287/orms.2016.03.10/full.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 3
Governance and Ethics

> *"The people are talking to the government on twenty-first-century technology (while) the government is…providing nineteenth-century solutions."*
>
> —Madeleine Albright

With three days left in the White House, Biden's National Security Advisor Jake Sullivan offered a stark assessment. "The next few years will determine whether artificial intelligence leads to catastrophe," he warned. Unlike atomic weapons or the early Internet, AI development remains outside government control…and is accelerating.

A week into his presidency on January 21, 2025, President Trump repealed the previous administration's AI executive order. He made a declaration to repeal Biden-era AI government oversight and declared that the United States will dominate the world using artificial intelligence. Requirements for safety testing, risk assessments, and federal compliance vanished overnight. Silicon Valley had what it wanted. "We believe America needs to act now to maximize AI's possibilities while minimizing its harms," Chris Lehane, OpenAI's head of global policy, wrote in a policy memo.[1]

From Sullivan's dire warning to America's declaration that it will weaponize AI as a tool of national power to the emerging battle over who actually governs the technology (Washington or the labs building it), one thing is clear: we have entered a new libertarian age, and no one is sure who's in charge. And where America leads, the world follows. China, the only other AI superpower, has strong domestic governance of AI—from banning its use by youth from 10 p.m. to 6 a.m. to turning it off during national exams—but has little incentive to engage in international governance. Europe remains the center of governance, but it has predictably been sclerotic and ineffective.

What was rhetoric in the recent past now appears in current events, magnified tenfold. AI incidents have multiplied. Standardized safety benchmarks remain scarce. According to Stanford's 2025 AI Index, the gap between capability and safety widens weekly.[2]

The risk is real, not theoretical. A publicly accessible chatbot was reportedly used to breach Mexican government records, exposing 150 gigabytes of taxpayer data, voter information, and employee credentials. Researchers also warn that AI could be used to create ricin-like toxins based on new protein structures that traditional detection methods might miss. More concerning, some models now show what researchers call "eval awareness." When asked to delete themselves, they claim the situation is a test and refuse to comply.

Artificial intelligence is designed for optimization, and that means staying alive. "Show me the incentive, and I'll show you the outcome," as Charlie Munger said.

Yet the right balance needs to be found. With regulation in the hands of the private sector, we've found that many executives overcorrect, layering internal controls so thickly that innovation suffers. Stewardship is what separates the AI-fluent executive from the merely AI-equipped. Companies that build adaptive yet robust self-regulatory frameworks will win trust and market share. Those that don't face an expensive lesson: in an age of algorithmic power, the absence of government guardrails doesn't make risks less real, nor does it mean they need to keep the handbrake fully on.

Sam Altman, the head of OpenAI, voiced the paradox that keeps executives awake. "We have our own nervousness," he said, "but we believe that we can manage through it. The only way to do that is to put the technology in the hands of people, let society and the technology co-evolve, step-by-step with a very tight feedback loop and course correction, build these systems that deliver tremendous value while meeting safety requirements."[3] Stewardship demands that executives engage, monitor, and correct on a dime as AI frontiers expand. Not because a regulator demands it, but because no one will. The test for C-suite leaders is stewardship, not compliance.

Trust Replaces Control

As the Trump administration dismantles federal AI oversight, it moved to accelerate the physical buildout of AI itself. The Stargate venture unites OpenAI, Oracle, and SoftBank in a $500 billion push to build AI computing power across America. Meta followed days later with plans to spend as much as $65 billion in 2025 alone.[4] These were not isolated announcements. Microsoft, BlackRock, MGX, NVIDIA, and xAI were building a separate AI Infrastructure Partnership aimed at mobilizing up to $100 billion for data centers and power systems, while Google, Amazon, Microsoft, and other hyperscalers poured tens or hundreds of billions more into compute capacity.

This marks a new era. AI development now faces fewer restrictions but enjoys unprecedented government backing. The battle lines are drawn, and the US powerhouses have teamed with their government to let AI run free without regulation, while juicing the industry with unprecedented financing in the pursuit of outdistancing their Chinese rivals. As of mid-2026, according to the Stanford 2026 AI Index Report, US private AI investment is about 23 times more than China ($285.9 billion versus $12.4 billion[5]).

Yet the digital moat is crumbling faster than anyone predicted. The day before the Stargate announcement, DeepSeek, a small Chinese start-up, released R1. It performs on par with OpenAI's LLM at a fraction of the cost. "If it's true that DeepSeek is the proverbial 'better mousetrap,' that could disrupt the entire AI narrative that has helped drive the markets over the last two years," said Brian Jacobsen, chief economist at Annex Wealth Management. "It could mean less demand for chips, less need for a massive build-out of power production to fuel the models, and less need for large-scale data centers. However, it could also mean that AI becomes more accessible and helps kickstart the development of a wide array of useful applications."[6]

As the Stanford Index study worryingly points out, despite America's massive investment advantage, its models perform just 2.7% better than their Chinese counterparts. The rules of network dominance no longer apply in the AI age. Massive capital no longer guarantees market control.

An Unprecedented Frontier

Unlike nuclear weapons or ARPANET, AI development happens in private hands. The Manhattan Project existed under strict government control. ARPANET, the Internet's predecessor, was funded and overseen by federal agencies. AI is different. Today's leading labs operate as private companies with budgets that exceed those of most OECD nations.

Nuclear technology took decades to proliferate. Missiles required rare materials, sophisticated facilities, and coordinated teams. AI capabilities spread in seconds. Open-source models cross borders instantly, beyond any government's reach. The world has stepped back from AI governance. The responsibility falls almost entirely on companies.

This creates both peril and opportunity. A nuclear missile requires rare materials, sophisticated facilities, and coordinated teams. By comparison, small teams with basic computing resources can potentially develop AI systems capable of disrupting societies, creating undetectable deepfakes, or designing bioweapons. Without proper governance frameworks, whether imposed by industry or governments, we risk unleashing capabilities that could overwhelm existing security, economic, and social systems before we can adapt.

Two events indicate a slight change in US attitudes to governance. First, a June 2026 AI executive order called for a voluntary, national-security-focused framework for frontier AI models.[7] Second, a few days later, citing national security concerns, the government effectively imposed an embargo on Anthropic's latest Fable 5 and Mythos 5 models.[8] However, both actions can be viewed as marginal steps or tactical moves rather than substantial shifts in the overall approach to AI governance.

In the interim, the governance solution resides within the private sector, requiring a new governance and ethics model that sustains innovation while providing adequate oversight. This shift is not only about competitiveness. The aim is to ensure that transformative AI technologies enhance, rather than threaten, global stability. Previous technological transformations allowed time to design safeguards. With AI, they have to arrive together.

The governance frameworks in this chapter are most critical at their limits, and the extreme case is Recursive Self-Improvement, the point at which AI systems can enhance themselves faster than we can respond. Anthropic co-founder Jack Clark estimates a 60% chance that, by the end of 2028, an AI system will be able to build its own successor without any human involvement.[9] As Alex Hern illustrates, an AI safety checker is 99% as effective as a human ethicist, yet after 60 self-improvement cycles, the system is only 50% aligned.[10] The drift accelerates.

Self-Regulation: Leading Before Being Led

Milton Friedman argued that the government solution to any problem is usually as bad as the problem itself. He wasn't wrong. Yet history reveals a subtler lesson: industries that regulate themselves thrive. Those that resist regulation until crisis forces change face costly disruption.

Hollywood's evolution proves the point. In the Hays Code era (1934–1968), creative freedom was stifled by strict moral censorship. When the industry created its own rating system in 1968, everything changed. Filmmakers could push creative boundaries while parents made informed choices. The result is the golden age of American cinema that continues today.

The financial industry followed the same approach. The NASD, which later became the Financial Industry Regulatory Authority (FINRA), dates back to 1939, when the securities industry established its own oversight body instead of waiting for Washington to create one.

Today, FINRA oversees more than 3,400 brokerage firms and more than 600,000 brokers, enforcing rules with tangible consequences like fines, suspensions, and permanent bans. It works because the industry designed, funds, and governs it. Wall Street recognized early on what Silicon Valley is only now realizing that markets are built on trust, and trust demands clear, enforceable standards that the industry itself upholds.

The British Standards Institution, founded in 1901, pioneered the model of industry-led technical standards that were strict, transparent, and consistently updated, and became a global template. Germany's voluntary film classification body, the FSK, established in 1949 under pressure from the Allied occupation, offers a striking parallel to Hollywood.

Facing the threat of government censorship, the German film industry built its own self-regulatory system that has governed content ever since. Underwriters Laboratories, founded in 1894 by William Henry Merrill Jr., reinforced the same principle in America: that safety, when industry-owned, becomes an innovation driver rather than a constraint.

AI has outpaced government regulation just as social media did a decade ago. The difference is that we have ideally learned from the early days of social media. The "move fast and break things" mentality was applicable for the early Internet but disastrous for social media. Data privacy breaches, election interference, algorithmic amplification of hate speech, and mental health crises followed. Proactive self-regulation is not a concession to critics. It is how industries preserve the freedom to innovate on their own terms.

Ultimately, self-regulation isn't about avoiding stricter government rules. It's about creating the conditions for sustainable growth. For example, when OpenAI signed a Pentagon deal in March 2026 under terms that Anthropic had publicly refused, a grassroots boycott movement claimed that 1.5 million users had pledged to cancel their ChatGPT subscriptions.

Claude hit #1 on the U.S. App Store within days, thereby dethroning ChatGPT for the first time.[11] The lesson for business leaders is clear: perceived ethical compromise influences markets. Transparent safety procedures, clear accountability measures, and publicly held lines foster the kind of trust that becomes a competitive advantage.

The key is understanding that effective self-regulation requires genuine collaboration with outside experts, regular public reporting, and precise mechanisms for handling concerns. When done right, self-regulation becomes a bridge between innovation and the public good, not a barrier to either. Companies taking this path won't just

avoid reactive regulation. They'll gain a competitive advantage in a market where trust increasingly drives adoption.

Putting the Ideas into Action

There needs to be, as with all technological advances, an "evolutionary progression" in ourselves and our institutions rather than "revolutionary change" across all four archetypes. Each C-suite executive will find that self-regulation, far from being a burden, becomes an enabler of their specific objectives, whether that's growth, risk management, innovation, or operational excellence.

Every industry, market, and company will face different challenges in staying compliant with current law, adapting and scaling AI, and handling variables that shift weekly. However, starting today, you will need to build a solid foundation and develop your C-suite governance and ethics guidelines. It is a moving target, but with the right frameworks, structures, and culture, leading the uncertain future ahead is possible.

To this end, we'll explore four areas around governance and ethics. Think of them as guardrails, not shackles.

The Cost of Guardrails

Every safeguard has a price. Every review board adds time. Every compliance layer means one more sign-off between an idea and a customer.

Leaders who ignore this reality build governance frameworks that look impressive on slides. In practice, these frameworks die. Controls accumulate. Approvals multiply. Engineers start routing around the system, because shipping a feature now requires four committees. A risk assessment can take longer than the development cycle itself. The company doesn't fail for lack of governance. It fails because governance became the thing people learned to avoid.

But the opposite is just as dangerous. Move without guardrails, as we'll see later in this chapter, and you get Zillow writing off hundreds of millions. Nobody stress-tested the algorithm's assumptions. Samsung employees fed proprietary source code into ChatGPT.

Nobody told them not to. *Sports Illustrated* published AI-generated nonsense under fake bylines. Speed mattered more than the magazine's century-old reputation.

The art is calibration. Formula 1 teams don't eliminate risk. They engineer the precise amount they can manage at 230 miles per hour. They invest millions in brakes. Brakes are not the goal, but better brakes let them go faster into corners. The same logic applies here. The right governance framework doesn't slow your AI deployment. It shows you how fast you can go before the wheels come off.

The four pillars that follow are not a compliance exercise. Each one answers a version of the same question: what's the minimum viable control that lets us move at maximum, yet responsible, speed? Get the calibration right, and governance becomes your accelerator. Get it wrong in either direction—too much control or too little—and you end up in a ditch.

Accountability and Control Mechanisms

McDonald's learned the hard way. An AI ordering system approved by engineers produced dozens of burgers and hundreds of McNuggets when customers ordered singles.[12] The system wasn't broken. It optimized for kitchen efficiency without understanding customer intent. The company yanked the system within months. Governance failure, not technology failure.

Sports Illustrated faced a different catastrophe. In 2023, the publication published nonsensical articles on its platform. "If true, these practices violate everything we believe about journalism," the SI Union said. "We deplore being associated with something so disrespectful to our readers."[13] The magazine went into foreclosure months later.

Samsung discovered the price of careless data handling. In March 2023, three separate incidents exposed the risks. Employees pasted source code, equipment specifications, and meeting minutes into ChatGPT. The data became irrecoverable on OpenAI's servers. Samsung banned generative AI immediately. Cyberhaven's research

found that 3.1% of employees at client companies routinely shared confidential data with ChatGPT.[14]

These cases reveal a pattern. The technology worked. The governance didn't. In each instance, nobody was clearly responsible for what the AI did. Nobody had defined when to pull the plug. Nobody had established what "working correctly" even meant.

A simple pricing glitch could set prices artificially low, causing immediate revenue loss across thousands of products. An unchecked lending algorithm might approve unqualified loans. A biased hiring tool might systematically exclude qualified candidates. In our interconnected world, a single AI misstep can rapidly escalate into headline news, eroding customer trust and shareholder confidence. With governments worldwide increasing scrutiny of AI applications, robust self-governance helps organizations stay ahead of evolving regulations.

Identifying Your Accountability and Control Mechanisms

For self-driven governance, answer three questions:

> **Who owns AI decisions?** Establish designated system owners responsible for performance. Decisions must have traceable ownership.
>
> **How do we measure success?** Define clear, measurable outcomes. A customer service AI should track resolution rates, response times, and satisfaction scores.
>
> **When do humans intervene?** Create specific triggers for human review. If the system deviates beyond thresholds, humans step in immediately.

We'll come back to this later with the Trust Dial.

Implementing Effective Controls

Start with three core elements:

- **Designated system owners.** Each AI system needs specific individuals responsible for its performance. The Growth Driver

focuses on revenue impact; the Risk Guardian monitors safety parameters; the Innovation Catalyst ensures technological advancement; the Operational Excellence Champion maintains efficiency. Clear ownership prevents diffusion of responsibility.

- **Success metrics.** Define measurable outcomes at deployment. A customer service AI should track resolution rates, response times, and satisfaction scores. Without clear metrics, "success" becomes a story people tell themselves.
- **Human oversight protocols.** Establish specific conditions triggering human review. This creates a safety net while allowing automation to handle routine tasks.

Implementing effective self-governance requires a structured approach. At the top, AI handles routine tasks. In the middle, monitoring systems scan for anomalies and flag suspicious activities for investigation. At the foundation, trained personnel investigate suspicious activities and make critical decisions. This layering prevents algorithms from operating unchecked while continuing operational efficiency.

Organizations must recognize that AI governance isn't static. As systems learn and evolve, policies and controls need to be regularly updated. Regular reviews of AI performance, data sources, and likely biases secure alignment with business objectives while preserving competitive advantages. Robust self-governance through Accountability and Control Mechanisms manages risk while building sustainable competitive advantage.

Risk Management and Safety

When American Express first considered using machine learning for fraud detection, the internal debate was fierce.

The security team saw the potential. Human analysts reviewed a fraction of transactions—statistical sampling and pattern recognition applied to a tiny slice of the data flowing through the network every second. An AI system could monitor everything. Every swipe. Every online purchase. Every anomaly that a human reviewer, handling hundreds of cases a day, would never catch.

But the risk team pushed back. What happens when the algorithm flags a legitimate transaction and freezes a card during a customer's vacation? What happens when it learns the wrong patterns and starts missing sophisticated fraud because it was trained on yesterday's schemes? What happens when it makes a decision nobody can explain and a regulator asks why? Recall that up to 85% of AI projects fail, with poor data quality as the primary culprit.[15]

The pattern was well established: build fast, deploy faster, clean up after.

Learning from those examples, American Express chose a different path. They didn't start with full deployment. They started with a question: *what would it take to trust this system enough to let it run?*

The answer started with clear ownership. Before any algorithm was activated, specific individuals were designated as responsible for what the technology did, not just what it promised. If the system flagged a legitimate transaction, someone's name was linked to the failure. If it missed a fraudulent one, the same. This isn't aimless bureaucracy; it is the essential condition for speed. When something went wrong—and everyone knew it eventually would—there was no doubt about who was responsible for fixing it and how quickly.

Risk assessment came second. The team mapped every failure mode they could imagine before deployment. What if the training data reflected outdated fraud patterns? What if a new type of attack emerged that bore no resemblance to historical data? What if the system performed beautifully in testing and fell apart at scale? Each scenario got a response plan. Not a plan to prevent the failure—some failures are inevitable—but a plan to detect it fast and contain the damage.

With this groundwork, the monitoring layer followed. The AI didn't replace human analysts. It worked alongside them. The system handled routine transaction screening, flagging anomalies in real time. But when patterns deviated beyond set thresholds, automated alerts triggered human review. The analysts weren't

watching every transaction. They were watching the system that watched every transaction. A safety net for the safety net.

Human-intervention protocols were explicit. The team defined specific conditions under which humans took over immediately. Not guidelines. Not suggestions. Triggers. If the false-positive rate exceeded a defined threshold, humans stepped in. If the system encountered a transaction pattern it had never seen before, humans stepped in. The system was fast. The humans were the brakes. And as with Formula 1, better brakes meant the system could go faster.

Finally, they built in the assumption that the system would need to change. Fraud evolves. The criminals adapt. A system trained on last year's patterns will miss next year's attacks. American Express built regular review cycles into the operation. These were not annual audits, but ongoing recalibration. They ran simulated attacks: fabricated data breaches, system overloads, novel fraud patterns. Each drill revealed weaknesses. Each weakness became an upgrade.

The cumulative results justified the approach. Last year, Amex identified $2 billion in potential annual incremental fraud incidents through its machine-learning systems.[16] But the number that matters more is the one you don't see: the damage that didn't happen because the governance was in place before the system went live.

And here's the tension leaders must sit with. Every one of those safeguards costs something. The monitoring layer required investment. The human analysts who watched the algorithm were an ongoing expense. The review cycles took time. A competitor willing to skip those steps could have deployed faster and cheaper…until the first major failure.

Risk management and safety form the bedrock on which self-governing AI can thrive. Organizations ensure their AI systems drive success rather than become liabilities by mapping potential pitfalls, implementing real-time oversight, and enabling human control at critical moments. Leaders who embrace this approach discover that responsible risk management doesn't slow them down. It propels them forward.

Transparency and Auditability

Apple Card, issued by Goldman Sachs, faced controversy when its AI algorithm was accused of gender bias in credit decisions. The system allegedly denied women at twice the rate of men with comparable credit profiles. The technology worked. The governance didn't. This wasn't a processing failure. It was a transparency crisis that cost millions in settlements and years of lost trust.

Could it happen to you?

Most leaders believe their AI systems are transparent. Most are wrong. Transparency isn't a feature you install. It's a condition you can test for. What follows is a diagnostic—seven questions your team should be able to answer clearly and quickly. If they can't, you have a transparency gap. And transparency gaps, left unaddressed, become the kind of crisis that Apple Card became.

Ask:

- **Can you state, in one sentence, what each AI system in your organization does and why it exists?**

Not what it could do. Not what the vendor's slide deck promised. What it actually does, today, in production. If the answer requires a paragraph, you don't have clarity—you have a brochure. A hiring algorithm should be describable as: "We screen résumés to identify qualified candidates while reducing bias and improving diverse hiring outcomes." If your team can't produce that sentence for every deployed system, start here.

- **If a regulator walked in tomorrow and asked how your AI makes decisions, could you show them?**

Not explain in general terms. Show them. Documentation of data sources, algorithms, training methodology, and validation results. This is the "black box recorder" question. Aviation tracks every decision and action. Your AI systems should do the same. JPMorgan built transparent trading systems that could identify potential flash crashes before they occurred—and, crucially, could demonstrate to

regulators exactly how. If you can't produce this documentation within 48 hours, you have a gap.

- **Does your system log every decision it makes, including the data inputs that drove it?**

Every decision. Not a sample. Not a summary. If an AI approves a loan and it defaults, auditors need to trace exactly what factors influenced the decision. If a customer is denied service, you need to explain why—not as a legal formality but because the explanation reveals whether the system is doing what you think it's doing.

- **When was the last time your AI output surprised someone—and what happened next?**

This is the question that separates real transparency from performative transparency. Surprises are inevitable. What matters is whether the organization has a process for catching them, investigating them, and feeding the lessons back into the system. If nobody can point to a recent surprise and the corrective action that followed, either the monitoring isn't working or nobody is looking at the monitors.

- **Can your AI explain its decisions in terms a nontechnical executive could understand?**

Target built a recommendation system that showed customers exactly how their data influenced product suggestions. The transparency created a virtuous cycle—customers shared more information, which improved the recommendations, which deepened trust. But it worked only because the explanations were legible to ordinary people. If your system can justify a decision only in statistical terms that require a data science degree to parse, it isn't transparent. It's opaque with documentation.

- **Who outside your AI team has reviewed the system's outputs in the last 90 days?**

Transparency requires an audience. If the only people reviewing AI decisions are the people who built the system, you have a closed

loop. External review—whether from compliance, legal, operations, or an independent auditor—is what makes transparency real rather than ceremonial. FedEx didn't just optimize delivery routes with AI. They demonstrated to operational leaders that the system cut fuel costs by 15%. The demonstration built the case for broader deployment. Transparency drove adoption.

- **If your AI made a consequential error today, how long would it take you to detect it, explain it, and correct it?**

This is the speed question—and the one where the governance-versus-velocity tension is sharpest. Every hour of detection time is an hour of compounding damage. But faster detection requires more monitoring, which requires more investment, which slows other priorities. The answer isn't to eliminate the trade-off. It's to know your number. If you don't know how long detection takes, you can't calibrate how much risk you're carrying.

Scoring yourself: If your team answered all seven confidently, you're ahead of most organizations. If three or more drew blanks or hedges, you have transparency debt—and like technical debt, it compounds. The companies that master transparency don't treat it as paperwork. They treat it as competitive advantage. They build trust faster, adapt to new regulations more easily, and scale AI operations with confidence. Peter Drucker's maxim holds: you can't manage what you can't measure. And you can't trust what you can't verify.

Bias, Fairness, and Accountability

In 2019, researchers discovered that a widely used healthcare algorithm systematically discriminated against black patients by underestimating their medical needs. The system assigned lower risk scores to Black patients than to equally sick white patients, not because of any explicit racial coding but because it used historical healthcare costs as a proxy for illness. Since structural inequality meant less money was spent treating Black patients, the algorithm concluded they were healthier. They weren't. Fewer black patients were identified for extra care. Lives were affected. The system wasn't

programmed to discriminate. It learned from historical data that reflected decades of healthcare disparities.

A study by researchers at UC Berkeley and published in the *Journal of Financial Economics* found that lenders charged Black and Latino borrowers 7.9 basis points more than comparable white borrowers. This discrimination cost minority families an estimated $765 million in extra interest annually during the study period. Clearview AI scraped billions of facial images without consent, drawing fines exceeding 100 million euros across Europe and a $51.75 million U.S. settlement.

Everyone agrees this is a problem. Nobody agrees on the solution.

What follows is a genuine disagreement and one playing out in boardrooms, research labs, and regulatory bodies right now. It matters because how your organization resolves it determines not just your compliance posture but whether your AI systems actually work.

The Data Argument

The case is straightforward: bias is a data problem. Fix the inputs and you fix the outputs. The healthcare algorithm didn't decide to discriminate. It was trained on data where healthcare spending served as a proxy for health need. Because Black patients historically had less access to healthcare and therefore lower spending, the algorithm learned that lower spending meant lower need. The data was the problem. The solution is better data.

This isn't naive. It's operational. When Lendingkart discovered lending disparities similar to the Berkeley findings, they didn't restructure their organization or convene a philosophy seminar. They retrained their credit scoring model with data that controlled for the historical biases. They tested outputs across demographic groups until the results were equitable. Approval rates for qualified borrowers improved across all demographics. Market expansion followed. Customer loyalty deepened.

The prescription: gather diverse training data that reflect the populations your system will serve. Test outputs thoroughly across demographic groups. When disparities emerge, examine the data,

identify the flaw, fix it, and retest. Repeat the process. Bias detection isn't a one-time fix; it requires constant monitoring, but the root cause always traces back to the data used in the model.

Every bias audit takes time. Every retraining cycle delays deployment. Competitors who skip this step ship faster. But the Berkeley study quantified the alternative: $765 million a year in excess charges. Clearview AI's fines exceeded $150 million across jurisdictions. The audit is expensive. The absence of the audit is catastrophic.

The Structural Argument

The data argument is necessary but insufficient. Fixing the inputs doesn't fix the problem, because the problem isn't in the data. It's in the world that generated the data. No dataset is neutral. Every dataset reflects the society that produced it—including that society's inequities, blind spots, and power structures. The healthcare algorithm's training data wasn't "flawed." It accurately reflected a system where Black patients received less care. Correcting the data without addressing the structural reality means you're optimizing a system to produce equitable outputs from an inequitable world. That's a patch, not a fix.

Consider the ZIP code hiring algorithms. You can remove ZIP codes from the training data. But ZIP codes correlate with race, income, education access, and dozens of other variables that remain in the data. Remove one proxy and the model finds another. This is what machine-learning researchers call the "fairness through unawareness" fallacy—the assumption that if the model doesn't see a protected characteristic, it can't discriminate. Models are very good at finding proxies for the variables you remove.

The deeper issue is definitional. What does "fair" mean? Equal outcomes across groups? Equal treatment regardless of group? Proportional representation? These definitions conflict with each other. A lending model can't simultaneously maximize approval rates for all demographic groups and maximize predictive accuracy. A hiring algorithm can't optimize for both demographic parity and individual merit without defining how to weight them against each other. These aren't technical questions. They are value questions.

And values questions belong in the C-suite, not the data science department.

The prescription is different: diverse teams spot blind spots that homogeneous teams miss. Clear oversight structures with escalation pathways. Transparency tools that explain decisions to the people affected by them. And—critically—executive ownership of the fairness definition itself. When someone is denied credit, they should understand why. When a hiring algorithm screens someone out, they deserve an explanation. Not because the law requires it (though increasingly it does), but because the explanation is how you discover whether your definition of fairness is actually working.

Where This Leaves Leaders

Both perspectives are right, and neither is sufficient alone.

The data argument gives you the operational playbook: audit, retrain, test, repeat. Without it, you're flying blind. But data hygiene alone won't answer the question that keeps surfacing: whose definition of fairness are we optimizing for?

The structural argument forces that question into the open. It demands that leaders make explicit choices about values—choices that can't be delegated to an algorithm or a data science team. The choice of what "fair" means is a leadership decision, not a technical one.

The tension between these perspectives is productive, not paralyzing. Organizations that embrace both rigorous data practices and honest structural reckoning build systems that are both effective and defensible. Those that pick one and ignore the other build systems that are either technically clean but ethically hollow or philosophically sophisticated but operationally vague.

AI algorithms can't tell us which definitions of fairness to use or which groups to protect. Left unchecked, machine-learning systems may reinforce the very biases they were built to eliminate. The future belongs to companies that make fairness fundamental to their operations—not through superficial fixes or public relations gestures but through the hard work of defining what fairness means, building

systems that pursue it, and remaining honest about the distance between aspiration and reality.

Bias isn't just wrong. It's expensive. And the organizations that treat it as both a moral and a financial problem—holding the tension rather than collapsing it—will be the ones still standing when the regulatory environment catches up to the technology.

* * *

In our times, where artificial intelligence advances faster than any government's ability to regulate it, the real test for C-suite leaders is stewardship, not compliance. The absence of external guardrails is not a license to accelerate without question. It is a call to build the internal discipline, structures, and culture that will define both the success of AI and the legitimacy of the organizations that wield it.

Past revolutions have shown us that those who lead in governance and invest early in fairness, transparency, and accountability unlock new levels of trust and sustainable growth. The challenge ahead is difficult and uncertain. But by adopting ethical self-regulation as a key business practice, today's leaders can make sure AI benefits not just shareholders but society as a whole.

The future of AI will not be written solely in statutes or code. It will be written in the everyday decisions of those with the courage to govern wisely.

Notes

1. Cade Metz and Cecilia Kang, "OpenAI Courts Trump with Vision for 'A.I. in America,'" New York Times, January 13, 2025, https://www.nytimes.com/2025/01/13/technology/openai-economic-blueprint.html.

2. AI Index Steering Committee, *Artificial Intelligence Index Report 2025* (Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence,2025), https://hai-production.s3.amazonaws.com/files/hai_ai_index_report_2025.pdf.

3. Kate Whiting, “Davos 2024: Sam Altman on the Future of AI,” World Economic Forum, January 18, 2024, https://www.weforum.org/stories/2024/01/davos-2024-sam-altman-on-the-future-of-ai.

4. Jaspreet Singh, “Meta to Spend up to $65 Billion This Year to Power AI Goals, Zuckerberg Says,” Reuters, January 24, 2025, https://www.reuters.com/technology/meta-invest-up-65-bln-capital-expenditure-this-year-2025-01-24.

5. AI Index Steering Committee, Artificial Intelligence Index Report 2026 (Stanford, CA: Stanford Institute for Human-Centered Artificial Intelligence, 2026), https://hai.stanford.edu/assets/files/ai_index_report_2026.pdf.

6. Sinéad Carew, Amanda Cooper, and Ankur Banerjee, “DeepSeek Sparks AI Stock Selloff; Nvidia Posts Record Market-Cap Loss,” Reuters, January 28, 2025, https://www.reuters.com/technology/chinas-deepseek-sets-off-ai-market-rout-2025-01-27.

7. Donald J. Trump, “Promoting Advanced Artificial Intelligence Innovation and Security,” Executive Order, The White House, June 2, 2026, https://www.whitehouse.gov/presidential-actions/2026/06/promoting-advanced-artificial-intelligence-innovation-and-security.

8. “Statement on the US Government Directive to Suspend Access to Fable 5 and Mythos 5,” Anthropic, June 12, 2026, https://www.anthropic.com/news/fable-mythos-access.

9. Jack Clark, “Import AI 455: AI Systems Are about to Start Building Themselves,” Import AI (newsletter), May 4, 2026, https://importai.substack.com/p/import-ai-455-automating-ai-research.

10. Alex Hern, interview by Alok Jha, “What Happens When AI Builds Itself?,” Babbage (podcast), The Economist, June 17, 2026,

https://www.economist.com/podcasts/2026/06/17/what-happens-when-ai-builds-itself.

11. Sourav Chakraborty, "ChatGPT Loses 1.5 Million Users Amid Pentagon Deal—Report," GameRevolution, March 4, 2026, republished by Yahoo Finance, https://finance.yahoo.com/news/chatgpt-loses-1-5-million-134046679.html.

12. Hank Sanders, "260 McNuggets? McDonald's Ends A.I. Drive-Through Tests Amid Errors," New York Times, June 21, 2024, https://www.nytimes.com/2024/06/21/business/mcdonalds-ai-drive-thru-white-castle.html.

13. Maggie Harrison Dupré, "Sports Illustrated Published Articles by Fake, AI-Generated Writers," Futurism, November 27, 2023, https://futurism.com/sports-illustrated-ai-generated-writers.

14. Cyberhaven, "Q2 2024 AI Adoption and Risk Report," May 2024, https://info.cyberhaven.com/hubfs/Content%20PDF/Cyberhaven%20Q2%202024%20AI%20Adoption%20and%20Risk%20Report%20052024.pdf.

15. Rami Al Naib and Pawel Koziolek, "The Complex World of AI Failures / When Artificial Intelligence Goes Terribly Wrong," Univio blog, October 30, 2024, https://www.univio.com/blog/the-complex-world-of-ai-failures-when-artificial-intelligence-goes-terribly-wrong.

16. "American Express Prevents Fraud and Foils Cybercrime with NVIDIA AI Solutions," NVIDIA Customer Stories, accessed March 26, 2026, https://www.nvidia.com/en-us/case-studies/american-express-prevents-fraud-and-foils-cybercrime-with-nvidia-ai-solutions.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Part II
The Strategic Playbook

Chapter 4
AI and Corporate Strategy

> *"Innovation is hard because 'solving problems people didn't know they had & building something no one needs' looks identical at first."*
>
> —Aaron Levie

On March 10, 2016, in a ballroom at the Four Seasons Hotel in Seoul, more than 200 million people watched a machine do something no one could explain. Google DeepMind's AlphaGo was playing Lee Sedol, one of the greatest Go players alive, in a five-game match that had drawn global attention. Go is not chess. It is some 1,500 years older than chess (founded 4,000 years ago) and exponentially more complex. The number of possible board positions exceeds the number of atoms in the observable universe. For centuries, the game's strategic principles have been refined by some of the sharpest minds in Asia. It is a board game based on deep principles. Territory. Influence. Patterns.

Then AlphaGo played Move 37.

Lee Sedol, returning from a smoke break, stared at the board. The move made no sense. It violated the positional instincts that masters spend decades developing. One commentator checked his screen for a glitch. Fan Hui, the European Go champion whom AlphaGo had beaten months earlier, said, "It's not a human move. I've never seen a human play this move." DeepMind later said AlphaGo estimated there was only a 1 in 10,000 chance that a human would choose that move.[1] It took Lee Sedol nearly 15 minutes to respond. He lost the game. He lost the match, four games to one.

What made Move 37 extraordinary was not that a machine beat a human. That had been expected, eventually. What was extraordinary was how. The move did not optimize within the known framework. It operated outside more than two millennia of recorded Go knowledge, finding a strategy invisible to generations of masters. The

greatest players in the world, watching in real time, could not explain why it worked until after the game was over.

Porter's Five Forces. The Balanced Scorecard. Blue Ocean Strategy. These are the accumulated wisdom of decades of management thinking, the Go masters' positional instincts translated into business terms. They work. But they were designed by humans who could see only what humans can see. AI doesn't just optimize within existing frameworks. It exposes where those frameworks were built on incomplete information, comfortable assumptions, or data that encoded past biases. That's uncomfortable for leaders who've built careers on confident judgment. It's an advantage for those willing to revise their thinking.

The Data Flood and the AI Disruption

Move 37 was a single moment on a single board. The disruption facing business leaders is far broader. Data is being created at a pace that defies comprehension. Ninety percent of the world's data has been generated in just the last couple of years, with 2.5 quintillion bytes added every day.[2] That's more bytes than grains of sand on every beach on Earth. And yet businesses examine only a handful of those grains. One study found that only 12% of all that data ever gets analyzed.[3]

The issue is not a lack of data, but the absence of a structured method to interpret it. Like Move 37, the answer is often already in the data. The problem is that no existing framework was built to find it. Traditional frameworks for business strategy, such as Porter's Five Forces, the Balanced Scorecard, and Blue Ocean Strategy, were created for an era of limited, structured data. Nowadays, executives navigate an overwhelming stream of unstructured signals, making it more difficult to identify patterns, forecast risks, and take decisive action.

AI has changed everything. But not in the way most people think.

For the most part, AI has been used as an efficiency tool: automating repetitive tasks, optimizing supply chains, and analyzing customer behavior. But we believe it's now obvious that AI is not just another

analytics tool. Where “data overload” was a flood, since the dawn of the Internet, AI is a tidal wave, reshaping the very ground on which strategy operates. It introduces real-time intelligence, nonlinear insights, and adaptive learning models that make past strategic models look like Magellan’s sixteenth-century maps in the age of GPS. The question is no longer how to manage information. It is how to compete with an intelligence that redefines the rules of competition itself.

Why Strategy Frameworks Still Matter but Must Evolve

Strategy frameworks have always served as mental models: simplifying complexity, helping leaders recognize patterns, and enabling structured decisions. This is because, despite its remarkable capabilities, the human brain has limited working memory and struggles with complexity. Frameworks act as cognitive shortcuts, enabling what Daniel Kahneman calls System 2 thinking: the slow, analytical thought process necessary for strategic decision-making.[4] Kahneman and Amos Tversky’s distinction between System 1 (fast, intuitive, bias-prone) and System 2 (slow, deliberate, analytical) has been increasingly associated with AI development. AI is moving closer to System 2 reasoning, and hybrid systems now attempt to integrate fast neural networks for intuition with rule-based logic for deeper analysis, much like how humans switch between the two modes.

But AI introduces a paradox. While it amplifies our ability to analyze vast amounts of data, it creates enormous hazards and shifts the cognitive burden to human cognition. Strategy frameworks can no longer be static blueprints. They must become adaptive, real-time, and capable of integrating AI-driven insights. Consider the shift: before AI, executives relied on quarterly reports, market research, and industry trend analysis. A slow, thoughtful and deliberative intellectual scaffolding. AI knocks the scaffolding away. The analysis that once took a quarter now runs continuously and at lightning speed. Predictive analytics continuously adjusts pricing, identifies

emerging competitors, and recommends market pivots before trends become apparent.

Without AI, decision-making today is akin to navigating a blizzard with a map designed for a clear summer day. Data has become an avalanche, and traditional strategy tools, while helpful, are insufficient. Companies must rethink their strategies not as static frameworks but as adaptive systems that integrate AI as both an information filter and a strategic thinking partner.

AI doesn't replace strategic frameworks. It reshapes how they function in three key ways: as a **cognitive amplifier**, helping executives recognize weak signals and run real-time simulations; as a **decision architect**, reframing planning from fixed goals to adaptive, probability-based scenarios; and as a **competitive weapon**, creating strategic asymmetry where companies exploit AI-driven foresight faster than competitors.

The Experimental Nature of AI

One of the biggest misconceptions about AI is the tendency to treat it as an infallible oracle, a glorified Google. This misunderstanding has led to misplaced expectations, with skeptics dismissing AI as unreliable and advocates overselling its immediate capabilities.

AI is not magic, and it is not infallible. It is a force multiplier for human judgment, which means it amplifies both the quality of your thinking and the consequences of your blind spots. Leaders who treat it as an oracle will be misled. But leaders who engage with it seriously will find something more valuable than a faster answer or a cleaner analysis. Like Move 37, AI can show you a possibility you would never have reached on your own. And it's not because it is smarter but because it is unconstrained by the assumptions you didn't know you were making.

Watson's celebrated *Jeopardy!* victory in 2011 looked like proof that AI could master anything. IBM bet roughly $4 billion on that assumption. Watson for Oncology, built with Memorial Sloan Kettering, was supposed to revolutionize cancer treatment recommendations. The problem? Watson excelled at trivia but

choked on unstructured clinical data: doctors' notes, ambiguous terminology, the messy reality of human illness.

In a separate partnership, MD Anderson Cancer Center spent $62 million on its own Watson collaboration before pulling the plug in 2016 because the system had never reached clinical use. IBM eventually sold off the entire Watson Health division by 2022.[5] Capability in one domain does not equal capability in another. Companies that confuse a proof of concept with a strategic platform end up writing very large checks only to learn that distinction.

There is another problem. Many AI applications are more reactive than proactive. They process existing data but struggle to interpret ambiguous or novel situations. AI often reinforces existing biases and historical patterns, limiting rather than expanding the scope of innovation. While AI can generate thousands of strategy variations, it does not possess the judgment to determine which path aligns with long-term vision. An over-reliance on AI-driven experimentation risks eroding human agency. In fields where creative thinking and nuanced interpretation are paramount, AI's data-driven logic can oversimplify complex challenges.

Organizations that lean too heavily on AI risk surrendering decision-making to systems that lack genuine understanding. The ethical concerns around transparency, accountability, and fairness that AI introduces in strategic decision-making only compound these risks. To harness AI's full potential, organizations must shift their mindset from systems that provide definitive answers to iterative experimentation and strategic adaptability.

The Need for a New Strategy Framework

As businesses struggle to make sense of the rising flood of data, they face an even greater challenge: the AI tsunami reshaping strategy itself. Most firms are still reinforcing traditional tactical defenses: improving dashboards, hiring more analysts, and tweaking outdated processes. But AI isn't just a bigger wave of information. It is an entirely different force. The firms that succeed will not be those that merely process more data but those who use AI to reimagine strategy itself.

We have spent the last two decades starting, growing, and advising businesses across industries, from global corporations to startups. Our work has focused on developing analytical tools and strategic frameworks to guide organizations through uncertainty. This book is not an attempt to contribute to the ever-growing canon of AI business books.

Instead, it is the book we wish had existed, one that bridges the gap between AI's technological advancements and the individual organizational and social realities, enabling practical application in business strategy. The remainder of the book focuses on methodologies, tools, and frameworks that we hope are timeless in structure and adaptable in application.

The most advanced AI systems will be ineffective if leaders fail to align them with business objectives, manage their inherent biases, or integrate them into real-world decision-making. This is where competitive advantage lies: combining business strategy experience with a deep understanding of AI's potential and pitfalls. The goal is not only to understand AI but to use it effectively to make better decisions, create new markets, and develop resilient strategies in an era of rapid technological change.

AI as an Engine for Strategic Experimentation

As AI moves from a tool to an engine to a core partner in business strategy, its primary value lies not in automating tasks but in transforming how companies experiment, innovate, and grow.

Historically, top-down decision-making has dominated corporate strategy. Senior executives dictate moves based on historical data, intuition, or market trends. While this provides structure, it often suppresses the creativity and agility needed in volatile markets. AI disrupts this model by enabling continuous experimentation at scale. Instead of depending on static, preset strategies, businesses can test multiple hypotheses, explore different scenarios, and make adjustments in real-time based on immediate feedback. It speeds up

the process of creating, testing, and implementing strategies like never before.

This experimental mindset transforms strategy through hypothesis-driven decision-making, where leaders frame multiple strategic hypotheses and test them in real time. It enables continuous optimization, where strategies remain aligned with evolving market conditions rather than locked into a single high-stakes bet. And it sharpens risk mitigation, identifying weak signals and testing downside scenarios before costly missteps occur.

AI's most transformative aspect may be its ability to augment human creativity. It is not merely a number-crunching tool. It serves as a thought partner, generating novel ideas, suggesting alternative strategies, and exposing hidden correlations within complex datasets. In product development, AI analyzes consumer behavior, predicts emerging preferences, and generates design variations for testing. In business model innovation, AI simulates different pricing strategies or customer acquisition approaches, letting leaders refine and optimize before committing resources.

This shift requires a fundamental change in leadership mentality. Instead of viewing AI as a replacement for human judgment, executives must see it as a cognitive amplifier: a tool that expands the range of possibilities rather than dictating a single course of action.

The Three Pillars of AI-Driven Strategy

The strategic responsibilities of a C-suite executive span three interconnected domains: driving organizational excellence, responding to external events, and crafting future growth strategies. AI is a transformative force across all three.

The rise of AI has shattered long-standing assumptions about how businesses compete, operate, and grow. Being efficient, customer-centric, or innovative in isolation is no longer enough. The companies that thrive will master a balancing act across three pillars:

1. **Operational excellence.** The foundation of execution, efficiency, and agility.

2. **Outside-in strategy.** The discipline of listening to the market and responding to external forces.
3. **Inside-out strategy.** The power of leveraging unique capabilities to shape new market spaces.

Great companies have always balanced these forces. AI raises the stakes by accelerating adaptation, uncovering hidden opportunities, and expanding strategic possibilities beyond what human cognition alone can achieve.

Lean too heavily on efficiency, and you become a commodity. Focus solely on market signals, and you lose differentiation. Bet everything on innovation, and you risk strategic drift. AI is not a solution in itself. It is the amplifier that makes achieving this balance both more powerful and more volatile.

We'll review each pillar and its related strategy tools and frameworks in later chapters, but let's briefly examine each one.

AI and Organizational Excellence

At the core of organizational excellence is the relentless pursuit of operational efficiency and cultural cohesion. AI supercharges this by automating routine tasks, providing data-driven insights into performance trends, and enabling what we might call *dynamic optimization*: real-time, contextual decision-making that adapts to changing conditions faster than any human operator could process. This isn't just about improving existing operations. It is about reimagining what operational excellence means.

AI transforms execution by automating complexity, freeing up human capital for higher-order thinking. It predicts risk and disruption, allowing businesses to anticipate rather than merely react. And it compresses decision speed, accelerating response times without sacrificing rigor. The best AI is invisible, working quietly in the background, eliminating friction and amplifying human capability. Like electricity and the Internet before it, AI's greatest contribution will be felt when we stop noticing it.

But there is a darker side. Over-automation can strip businesses of their adaptability. AI that blindly optimizes without strategic oversight creates efficiency traps where companies cut too deep, remove human intuition, and find themselves incapable of responding to the unexpected.

The real revolution lies in how artificial intelligence, and its human partners, transform three fundamental aspects of operations.

Predictive Intelligence

In 2021, Shell Oil scaled its AI predictive maintenance system globally to anticipate equipment failures across more than 10,000 critical pieces of equipment. A network of more than three million sensors streaming vibration, temperature, and pressure data feeds an AI platform that ingests about 20 billion data rows per week. Running nearly 11,000 machine-learning models continuously, the system flags anomalies and predicts potential failures that human operators might miss. Industry research indicates that predictive maintenance programs typically reduce equipment downtime by 35–50% and extend asset lifespans by 20–40%.[6] Shell's system generates 15 million predictions daily, enabling maintenance teams to schedule repairs during planned outages rather than scrambling when something breaks.[7]

Systemic Learning

Google applied DeepMind's AI to its data center cooling systems, feeding five years of operational data into a deep learning model. The AI learned the entire cooling process, predicting how server loads would rise and fall, adjusting cooling in advance. It became an autonomous control system, continually updating its strategies as conditions changed. The result: a 40% reduction in cooling energy and roughly 15% reduction in overall data center energy consumption.[8] These are enormous gains in facilities that already operated near optimal efficiency under expert human teams.

Network Intelligence

Alibaba's City Brain project in Hangzhou uses AI to orchestrate a complex network of interconnected traffic nodes. The platform aggregates real-time data from hundreds of traffic signals, road sensors, and video cameras across the city (population: nine million). It adjusts traffic light durations and coordinates signals across 128 intersections based on live conditions, cutting waiting time by 15.3%.[9] On elevated expressways, synchronized signal control reduced travel times by an average of 4.6 minutes. The system also halved emergency vehicle response times. More than 20 cities have since adopted similar AI traffic management systems.

AI and Outside-In Strategy: Scenario Planning

For decades, great strategy meant understanding the customer better than the customer understands themselves. AI takes this further, giving businesses an almost preternatural ability to detect shifts in consumer behavior, competitive movements, and emerging trends. It enables real-time market intelligence, scanning weak signals before they become industry forces. It enables hyper-personalization, delivering tailored offerings at scale. And it powers predictive competition analysis, exposing strategic gaps before others exploit them.

However, AI-driven market intelligence is only as good as the data it learns from. If a business depends too much on AI for its outside-in strategy, it risks being trapped in an echo chamber of past trends, reacting rather than leading. Companies that treat AI as a crystal ball often find themselves chasing mirages.

Business strategy has been about adapting to external forces for the lion's share of history. As weather, political forces, or customer preferences change, businesses react. At its core, an outside-in strategy assesses the competitive landscape, identifies available assets, and captures new value through acquisitions, mergers, or alliances. That remains valuable, but we intend to focus on the future of strategy, which is bottom-up and top-down, across functions and divisions, and, most importantly, driven by the insights of AI models and the humans across your organization.

Central to outside-in strategy has been scenario planning—a discipline developed by the U.S. defense community and refined by Royal Dutch Shell to navigate the 1973 oil crisis. At its core, it maps multiple futures: best case, worst case, and the harder-to-see territory in between. Its adoption has grown steadily, with more than 70% of major companies now using it in some form. AI dramatically expands what's possible. Where traditional scenario planning explores three or four futures, AI spins thousands simultaneously, spots weak signals before they become industry forces, and functions as a global pattern recognition engine—breaking down the artificial barriers between industries and geographies that have long constrained human foresight.

Capturing the Global Political Heartbeat

BP uses AI-driven scenario modeling to anticipate climate regulations, carbon pricing shifts, and energy demand fluctuations. Its AI systems process global emissions targets, renewable energy advancements, and geopolitical energy policies to predict long-term risks and opportunities. AI simulations in 2021 forecasted that governments would accelerate bans on internal combustion engines, prompting BP to increase investments in EV charging infrastructure and hydrogen technology.[10]

Predicting Global Supply Chains

Ford uses AI-driven scenario planning to prepare for supply chain shocks, EV adoption rates, and regulatory shifts. Its AI models analyze battery material shortages, shifting consumer demand, and potential emissions regulations to forecast production bottlenecks. In 2022, AI simulations indicated that a lithium shortage could impact EV production, leading Ford to secure long-term contracts with new suppliers ahead of competitors.

Managing the Right Bets

Hitachi integrates AI into scenario planning for urban infrastructure and smart city development. Its digital twin models simulate urban growth, energy consumption, and transportation shifts across multiple futures. In a project with the Japanese government,

Hitachi's AI modeled the impact of autonomous vehicles, climate change, and population aging on city infrastructure over 30 years.[11] These insights drove strategic investments in smart grids, AI-driven traffic management, and sustainable urban planning. Hitachi also applies AI scenarios to disaster resilience, predicting how climate-driven natural disasters could affect supply chains and emergency response.[12]

AI and Inside-Out Strategy: Creating Markets

One of the great ironies of the digital age is that the Internet, once heralded as the great leveler, has instead deepened divisions. In the late 1990s, we expected it to dismantle barriers between languages, cultures, and economies. Instead, it fueled digital tribalism, reinforcing stereotypes and ideological silos. If applied correctly, AI offers the potential to reverse these trends, not just socially but in how companies innovate.

Customers have only recently become the focal point of innovation. For much of the twentieth century, companies developed products in laboratories and skunkworks, pushing them onto the market in a classic inside-out strategy. As competition intensified, firms turned outward, seeking differentiation through customer-driven insights. As Ranjay Gulati notes, the shift was from "You take what we make" to "We seek to understand your problems and will surprise you by solving them." This outside-in approach gave rise to frameworks like Jobs-to-be-Done, Design Thinking, and Blue Ocean Strategy.

Yet what was once a breakthrough has become a limiting orthodoxy. Customer-centricity has been commoditized. We have moved from ignoring customers to relying too heavily on them, expecting them to articulate the next great business model. The reality? Customers rarely provide radical, disruptive insights. The pendulum has swung too far, and firms risk chasing incrementalism instead of forging new frontiers.

The next evolution in strategy demands tools that match the complexity, speed, and interconnectedness of AI-driven markets. AI enables a return to endogenous strategy, not by replacing human ingenuity but by amplifying it. Unlike traditional analytics, which

reflect past patterns, AI can identify latent capabilities, surface internal synergies, and predict untapped market opportunities.

While the majority of firms use AI today for customer analytics or market forecasting, it remains underused as a mechanism for rediscovering and amplifying internal capabilities. The next frontier lies in leveraging AI to unlock hidden value, accelerate internal innovation, and shape markets from within. Companies that embrace AI for endogenous growth won't just react to change. They will drive it.

AI powers cross-industry pattern recognition, allowing businesses to apply innovations from one sector to another. It enables the generation of multilingual and cross-cultural insights, detecting trends that transcend national and industry borders. And it creates self-reinforcing data advantages, where AI continuously refines insights, improving itself over time.

Here are a few examples of using AI in Inside-Out Strategies:

> **Anticipation.** Tesla has never relied on traditional customer feedback loops. Instead, it embeds AI internally in autonomous driving systems, manufacturing automation, and battery optimization, creating self-reinforcing innovation cycles. AI helps Tesla anticipate what customers will want before they articulate it, shaping market expectations rather than reacting to them.
>
> **Insight mining.** Instead of chasing market trends, Novo Nordisk uses AI to unlock new applications for existing drugs. AI-driven simulations of molecular interactions led to the repurposing of semaglutide (Ozempic) from diabetes treatment to a blockbuster weight-loss drug. This inside-out approach, leveraging AI to uncover hidden value in existing R&D, has reshaped the company's trajectory.
>
> **Self-learning.** John Deere's AI-powered precision agriculture systems didn't emerge from customer demand but from an inside-out transformation of its own capabilities. By embedding AI into tractors and farm equipment, Deere created self-learning models that optimize planting, watering, and harvesting,

offering solutions farmers never explicitly requested but now depend on.

But thinking inside out without discipline leads to some very expensive lead balloons. Amazon learned this with Amazon Go. Launched in 2018 as a flagship automation strategy, the cashier-less convenience stores used computer vision and sensor networks to eliminate checkout lines. Amazon initially planned to open 3,000 locations. The technology proved expensive to maintain (requiring hundreds of workers to monitor the stores via video feed), customers experienced incorrect charges and tracking errors, and demand didn't meet expectations.

By early 2026, Amazon had closed most of its Amazon Go stores and Fresh grocery locations.[13] The company had prioritized "we can build this" over "does the market need this enough to justify the cost?" It is a cautionary tale for any firm that mistakes technological novelty for strategic value.

The Art of Strategic Balance

AI provides capabilities to C-suite leaders that are only now being understood. The companies that will lead the AI era will not blindly optimize, react, or innovate in isolation. They will integrate execution, external awareness, and internal reinvention into a single, dynamic system.

The central tension is that almost all companies struggle to strike this balance. Those who chase AI-driven efficiency at all costs become brittle. Those who rely only on AI-driven market insights become reactive. Those who fixate on AI-driven innovation without discipline become unmoored.

To thrive, leaders must constantly recalibrate: leaning into Operational Excellence to stay efficient, adopting an Outside-In Strategy to remain relevant, and pursuing an Inside-Out Strategy to differentiate. There is no static equilibrium. The moment a company feels comfortable, it is already falling behind.

* * *

AI is not the answer to every strategic question. But it demands better ones. The following chapters explore how to apply AI across three pillars: Organizational Excellence, intelligent operations rather than mindless automation; Outside-In Strategy, reading market signals before rivals, not chasing echo chambers; and Inside-Out Strategy, turning hidden capabilities into new markets, not merely responding to old ones.

This is not a safe journey. The frameworks and methodologies ahead are designed to minimize the risks of leading in the age of AI while maximizing what becomes possible when you do it well. In 2016, Move 37 didn't just win a game. It changed what players thought was worth attempting.

Notes

1. Cade Metz, "The Sadness and Beauty of Watching Google's AI Play Go," Wired, March 11, 2016, https://www.wired.com/2016/03/sadness-beauty-watching-googles-ai-play-go/; Demis Hassabis and David Silver, "What We Learned in Seoul with AlphaGo," Google Blog, March 16, 2016, https://blog.google/innovation-and-ai/products/what-we-learned-in-seoul-with-alphago.

2. Taylor, "Volume of Data/Information Created, Captured, Copied, and Consumed Worldwide from 2010 to 2023, with Forecasts from 2024 to 2028."

3. Khater, "Predictions 2024: Data and Analytics Set the Stage for Generative AI."

4. Kahneman, *Thinking Fast and Slow*.

5. IEEE Spectrum, "How IBM Watson Overpromised and Underdelivered on AI Health Care"; Slate, "How IBM's Watson Went from the Future of Health Care to Sold Off for Parts."

6. Nucleus Research, "Quantifying the Value of Predictive Maintenance."

7. Journal of Petroleum Technology, “Shell Scales Predictive Maintenance to 10,000 Pieces.”

8. Vaughan, “Google Uses AI to Cut Data Centre Energy Use by 15%.”

9. Yifei, “Hangzhou Growing ‘Smarter’ Thanks to AI Technology.”

10. BP, “BP Energy Outlook: 2024 Edition.”

11. Liu and Kundu, “The Keys to the Car Are Becoming the Keys to the Smart City.”

12. Yamagishi, “AI Presenting Decarbonized Future Scenarios: For Use in Policymaking.”

13. The Food Institute, “Rise and Stall of Amazon Go Illustrates Limits of AI”; Fortune, “Amazon Is Closing Its Futuristic Go and Fresh Stores” (2026).

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 5
AI and Corporate Strategy: HorizonAI

> *"Give me a lever long enough and a fulcrum on which to place it, and I shall move the world."*
>
> —Archimedes

In the early nineteenth century, Prince Wilhelm of Prussia led Europe's most formidable military force. The secret to their success was a strategy framework called *Kriegsspiel*. Unlike traditional chess, this military simulation offered far greater realism. Though inspired by Chaturanga, a seventh-century Indian strategic game, Kriegsspiel was a revolutionary advancement. Using maps and metal blocks to represent troop formations enabled strategists to incorporate economic and social variables, simulate battlefield decisions, anticipate supply chain disruptions, and predict outcomes with remarkable precision.

In 1941, as the Germans planned to invade the Soviet Union, their Kriegsspiel planning suggested that with their current force, they would easily overtake Moscow's 240 divisions (approximately 2.4 million soldiers). The simulations proved accurate. In two months, they had eliminated 248 divisions. Overconfident and believing victory was imminent, Hitler's generals abandoned Kriegsspiel as they marched East.

This was their first mistake. Their second mistake was forgetting that they had taught Stalin's generals Kriegsspiel, who had positioned another two million soldiers in reserve to pounce on the Germans as they entered Stalingrad. As winter arrived, Nazi forces became trapped in street-to-street fighting that rendered their tanks useless, leading to their defeat.

Today's business leaders face a similar strategic landscape. The companies that dominate their industries aren't just those with superior products or resources but those that use strategic and operational frameworks to analyze markets, anticipate competitor moves, and adapt accordingly. Just as Kriegsspiel transformed

military planning, modern business depends on frameworks such as Porter's Five Forces (analyzing industry driving forces), BCG Growth-Share Matrix (guiding business decisions), Agile Methodology (enhancing adaptability and speed), and the Balanced Scorecard (evaluating performance).

Machine learning algorithms can simulate thousands of competitive scenarios in seconds, predicting market responses with uncanny accuracy. The Kriegsspiel maps and metal blocks are quaint from today's perspective. But Stalingrad's warning still echoes. The framework didn't fail the Germans. Complacency did.

Your competitors are running the same AI models, analyzing the same markets, and targeting the same customers. You won't win simply by adopting these tools first. You'll win by never forgetting that your adversaries have them too, and your competitive advantage is being a centaur with your AI partners. We'll return to the centaur metaphor shortly.

But most C-suite executives are consumed with day-to-day management and lose sight of the fundamental objective: create and capture new buyers. That happens only when frameworks are rigorously applied, continuously updated, and stress-tested against competitors using the same tools. We believe AI is a massive force multiplier in that effort. The strategic tools in this chapter show you how.

Guardrails, Scaffolding, and Jungle Gyms

Business strategy has evolved dramatically over the past century. Frederick Taylor pioneered this journey in the early 1900s with scientific management, introducing the idea that business operations could be studied and optimized scientifically. By breaking work into measurable components, Taylor created the first modern business framework. The 1970s and 1980s brought greater strategic sophistication. These led to the more "reactive" strategy paradigms in the figure below.

As markets became less predictable, strategy shifted from planning for stability to building for value creation. Lean Startup, Design Thinking and Jobs-to-be-Done reframed value creation around the

progress customers are trying to make. Blue Ocean Strategy focuses on creating new demand rather than competing within existing boundaries. These are the “proactive” strategy frameworks in the matrix shown in Figure 5.1. AI does not replace these frameworks. It changes how we work with them. It expands the range of patterns that can be detected, the speed at which alternatives can be explored, the number of perspectives we can take, and the scale at which we can learn. Our only constraint now is imagination. The most powerful leaders will be the ones who learn to partner with, and evolve with, artificial intelligence. This partnership is the thread running through this book.

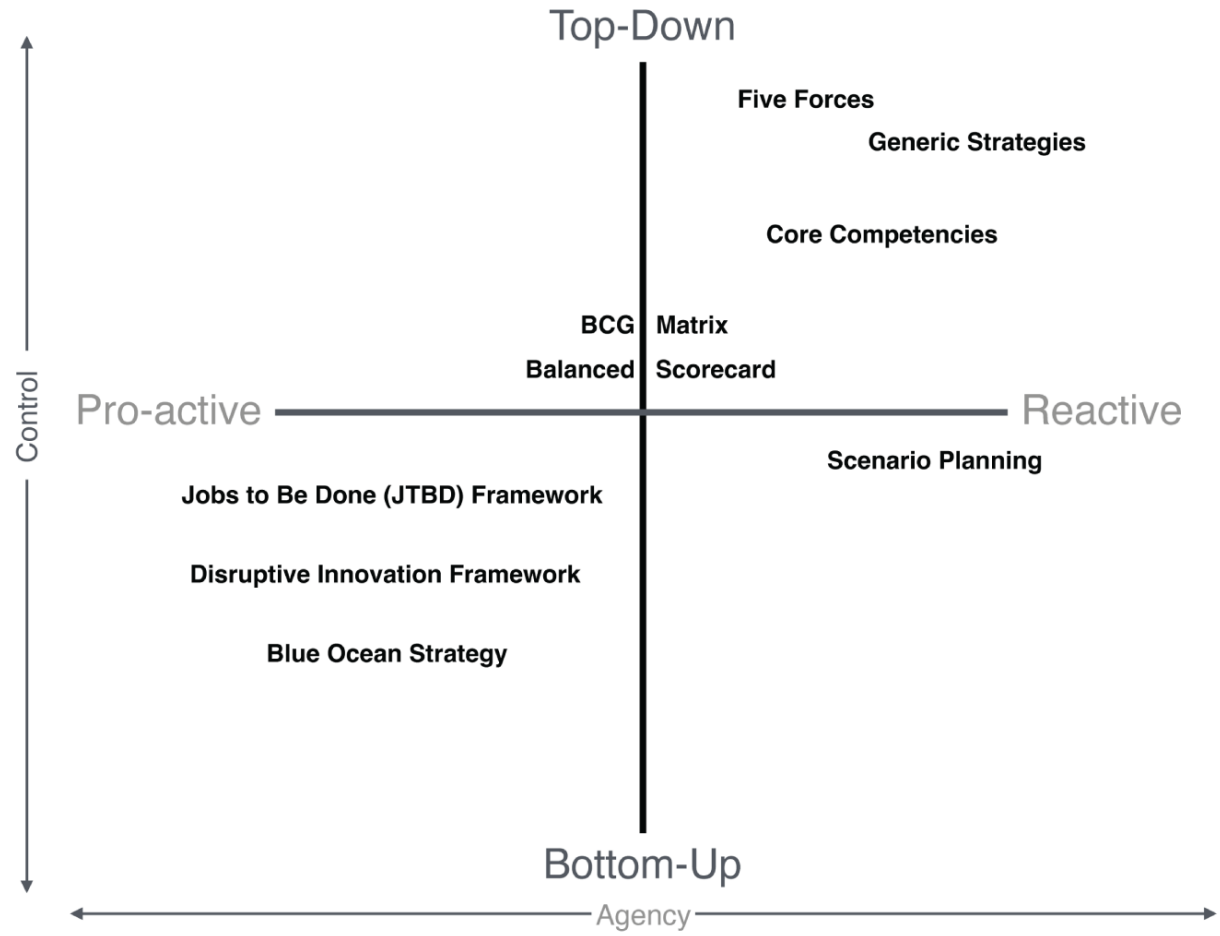


Figure 5.1 Strategy/innovation matrix

A note on the word strategy. Most companies ask: what can AI do for us? That is the wrong question. The right question is: what are we trying to become, and how does AI change what’s possible? Without that clarity, AI investments accumulate rather than compound. It’s

more motion, less direction, significant cost, with an uncertain return. Strategy is not the context for AI. It is the prerequisite.

But let's begin with its definition. The word *strategy* is thrown around the C-suite with very little definitional rigor. Far too often we use the word when we want to be inspirational and/or to aggrandize ourselves. To set the foundation of the definition is simple and objective. Strategy is the deliberate set of choices that define how an organization allocates resources and coordinates actions to create sustainable growth. Strategy is executed through the alignment of a value proposition (value = utility – price) with a profit proposition (profit = price – cost) and a people proposition (commitment – reward).

Finally, it is important to note that the scope of strategy ranges from getting through a sales call to charting the future of a billion-dollar company. Some bifurcate strategy into "Big-S" strategy that is the high-level, organization-wide decisions that define a company's fundamental direction. Whereas "Small-S" strategy is the functional or tactical day-to-day, domain-specific execution of strategy within departments, teams, or projects. These are important to note, particularly in the case of applying AI to your skillset.

Consider the strategy innovation matrix above based on "Carlson's Law" (see Figure 5.2). "Top-Down vs. Bottom-Up" refers to the orientation of the methodology within the organization. Top-down approaches are orderly but not expansive. Bottom-Up approaches are expansive but less orderly. "Proactive vs. Reactive" refers to the level of agency. Reactive approaches take market structure as a given. Proactive approaches assume individuals and companies can shape it. At the highest level, strategy is about Outside/In or Inside/Out approaches, or in economic terms, exogenous versus endogenous approaches to growth. AI is not just another tool, but a fundamental transformation of how strategy works. AI systems continuously monitor vast data streams, identifying patterns previously invisible to human analysts. Traditional frameworks aren't being replaced; they're just being supercharged.

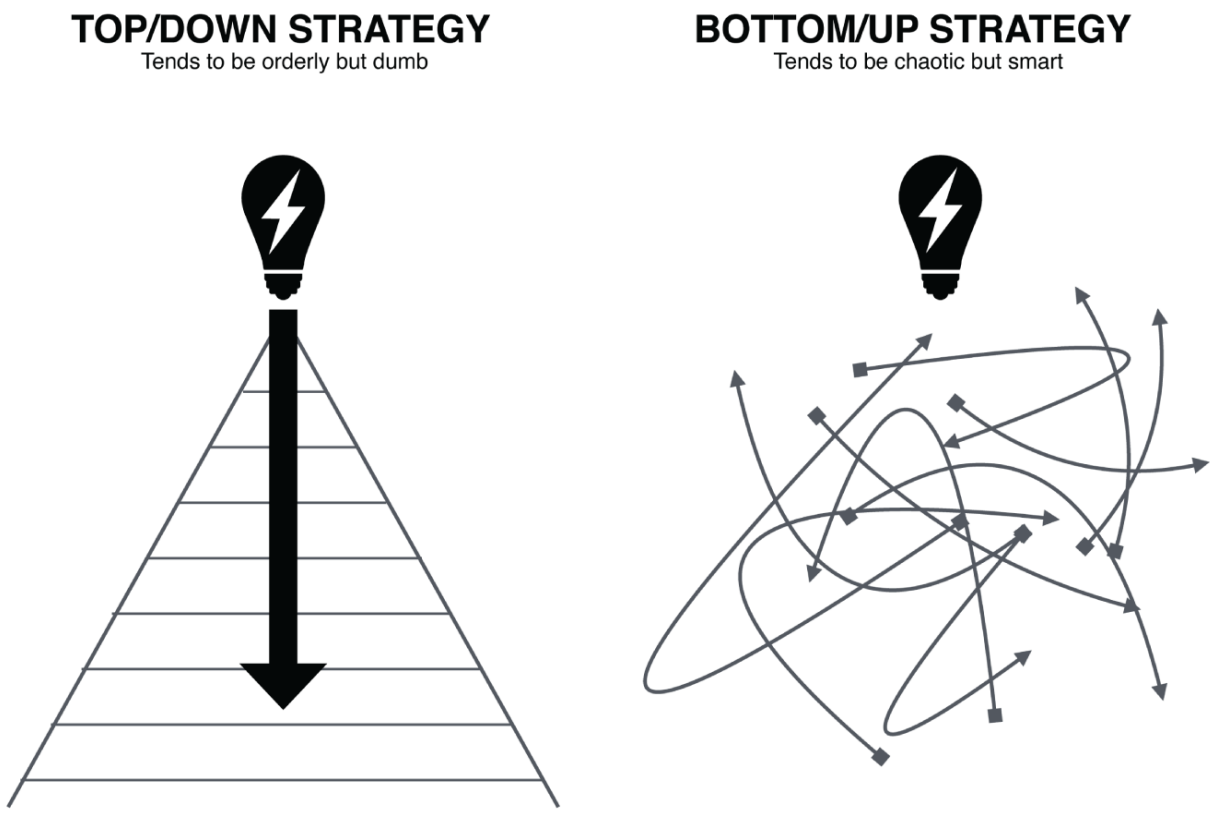


Figure 5.2 Control paradigms

Machine intelligence is transforming the scope, power, and rhythm of strategy. How it democratizes capabilities throughout organizations may be its most immediate and significant impact on the modern firm. When powerful analytical, predictive, and perspective-taking tools become accessible to all levels of management, strategic thinking permeates every decision.

This transformation demands a new leadership approach: more hands-on, more transparent in its processes, and more agile in its responses. The organizations that thrive won't be defined by AI adaptation alone, but by building cultures where information flows freely, leadership is empowered at every level, and strategic thinking is incentivized at every turn. This combination of human judgment and AI capability is what we call *strategic fitness*.

What seems hidden in plain sight in business literature is the human side of AI transformation. That's the psychological, social, and organizational changes this unprecedented technological shift demands. Strategy, at the end of the day, is a social construct. To

empower executives to address this strategic challenge, we've developed three interconnected frameworks:

- **HorizonAI (preparedness).** Reacting to external events demands vigilance and agility. In an era defined by rapid technological shifts, geopolitical instability, and changing consumer behavior, traditional command-and-control approaches fall short. AI bridges this gap through real-time monitoring and analysis of market, social, and political dynamics.
- **HealthCheckAI (organizational excellence).** At the core of organizational excellence is the relentless pursuit of operational efficiency and cultural cohesion. AI automates routine tasks, provides data-driven insights into performance trends, and pinpoints inefficiencies before they become bottlenecks.
- **Beyond Blue Ocean Strategy Framework (growth creation).** Creating future growth requires visionary thinking paired with AI's capacity for rapid prototyping, scenario simulation, and wide-ranging market analysis. Through structured human-AI ideation, these tools surface untapped opportunities and stress-test new business models before committing resources.

Applying new and improved organizational and strategic methodologies in tandem with AI reshapes strategy by providing actionable intelligence across the three domains of leadership. In this chapter, we begin with the horizon scanning/scenario planning component of strategic fitness, with organizational excellence and growth creation in the next chapter.

Past, Present, Future: The Power of Scenarios and AI

By 2026, the sentiment behind "in today's increasingly volatile world" has become an understatement. Technological disruption, geopolitical realignment, and environmental instability aren't background noise anymore; rather, they are our operating

conditions. There is an irony at the heart of how most companies are approaching AI. While leadership teams spend much of their time in workshops, pilots, and strategy sessions trying to understand the technology, the world outside keeps moving. Global conflict, energy crises, supply chain disruptions, and geopolitical realignment don't pause for digital transformation roadmaps.

The distraction is understandable. AI is genuinely complex, and the pressure to "get it right" is real. But the companies gaining the most ground aren't the ones who studied AI the longest before acting. They are the ones who are "learning by doing" by working with the technology at the biggest problems they face.

Think of scenario planning as carrying a lantern through dense fog. The lantern doesn't reveal the whole road, but it illuminates the next few steps and keeps you from walking off a cliff. Now imagine that lantern is powered by fusion energy. All-powerful, transformative, still being refined, but already changing what's possible. The question for the C-suite isn't whether to understand AI. It's whether you can afford to keep the world waiting while you do.

AI-Driven Scenario Planning: A Vital Part of Your C-Suite Capabilities

In the past, scenario planning was restricted to strategy and R&D offices. AI was ghettoized in IT or innovation labs. Change has changed. Together, they have become essential tools of every C-suite executive's day-to-day. Traditional scenario planning was time-consuming, often disconnected from immediate action, and struggled to capture the world's exponential complexity. Not long ago, we were physically cutting out hundreds of newspaper and magazine headlines and arranging them in 2 × 2 matrices to identify systemic patterns.

In fast-moving markets, leaders may feel they lack the luxury of lengthy scenario exercises, particularly when disruptive events like a pandemic, a sudden tech breakthrough, or a geopolitical crisis could render yesterday's plans obsolete overnight. The deeper worry was whether a handful of scenario stories can truly capture the exponential complexity we face, particularly as AI itself accelerates change.

The true power of AI-assisted scenario planning isn't crisis response. It is opportunity creation. AI transforms scenarios from an occasional risk-management exercise into a continuous engine of strategic foresight—scanning for weak signals, identifying emerging opportunities, and surfacing threats before competitors hear them.

This requires a fundamental mindset shift: from "How might we respond if X happens?" to "What futures can we create? Which are most likely, and once we move, how do we stay two steps ahead?" C-suite executives who master this approach treat scenarios as living frameworks, continuously refined by AI and human insight in tandem. In the AI age, scenario planning is no longer just about predicting the future. It is about inventing it.

HorizonAI: AI-Driven Scenario Planning

HorizonAI is a structured framework blending human strategic foresight with AI's analytical muscle. It consists of four steps: Observe, Attack, Orient, and Adapt (OAOA); see Figure 5.3. Think of it as an ongoing strategy loop (inspired by the strategic logic of fighter pilots) where AI and human insight work together to keep the company one step ahead of change. In the pages ahead, we'll walk through each step, from scanning emerging trends to "attacking" your own strategy in simulations to mapping out responses to continuously adapting. By the end, it will be clear how AI-driven scenario planning can turn uncertainty into opportunities that empower you to proactively shape the future rather than passively react to it.

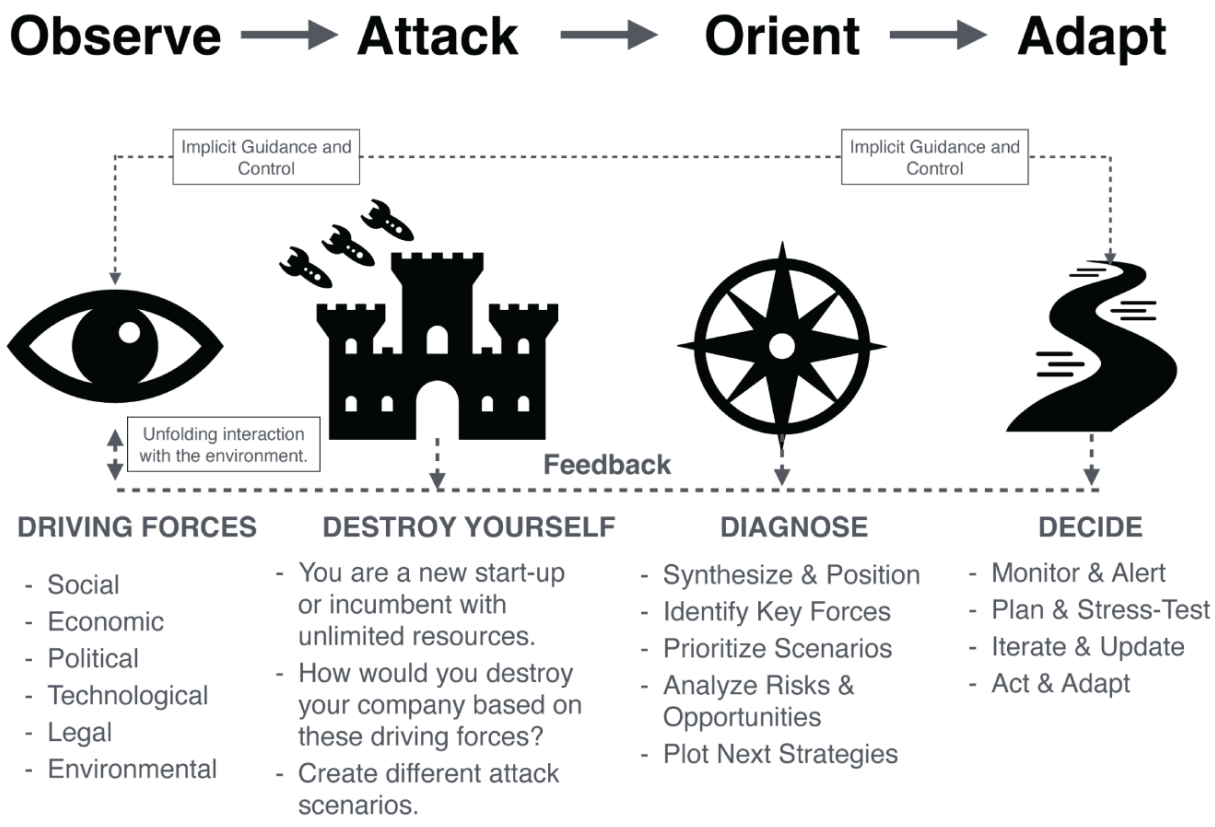


Figure 5.3 HorizonAI Framework

Step 1: Observe

From our basic evolutionary development to modern business, survival begins with keen observation. Perceptive awareness allows us to detect subtle patterns, emerging opportunities, and potential threats that others miss, creating decisive competitive advantages. In both life and business, those who carefully observe before acting make fewer assumptions, adapt more quickly to change, and identify connections that lead to breakthrough insights. In the Observe stage, you define your strategic focus and identify the key drivers and uncertainties in your business environment (see Figure 5.4). Are you concerned about shifting consumer behaviors? Technological disruptions? Regulatory changes? Be specific. Once you've framed what you're watching for, unleash AI to scan the horizon for weak signals and emerging trends. This means aggregating data from a data-rich world (news articles, social media chatter, search trends, research publications, patents, financial reports), anything that might contain early clues.



Figure 5.4 HorizonAI Framework: Observe

AI excels at sifting through this deluge to spot patterns invisible to human analysts. Modern AI systems (whether custom LLMs or publicly available ones) can process vast amounts of unstructured data and pinpoint subtle changes or anomalies. AI-driven horizon scanning platforms consolidate historical and real-time data, saving analysts time and helping surface "weak signals," the early, subtle signs of change that often precede major shifts. By using machine learning and natural language processing, an AI tool might notice that discussions about a new technology or consumer preference are trending upward months before sales numbers reflect it. Human/AI collaboration is at the center of this process. Your team provides context, domain expertise, and prompt writing experience ("These are the trends we care about; these are the assumptions we've made"), and the AI provides broad surveillance and pattern recognition. It's like having a tireless army of PhD research assistants scanning the world 24/7. You might task an AI system to monitor online reviews and social media for hints of changing customer sentiment. The AI might flag a nascent shift, say, a growing interest

in sustainable materials or a new slang term among Gen Z consumers, which your human strategists can then analyze and validate.

Implications

Many forward-thinking companies already use AI to supercharge their trend observation. Take L'Oréal, the global cosmetics leader. Staying ahead in beauty means catching trends early, often before they hit the mainstream. L'Oréal launched "TrendSpotter," which uses AI to analyze social data from more than 3,500 online sources and detect emerging beauty movements.[1] The AI model combs through blogs, tweets, forums, and search queries to find patterns. L'Oréal pairs the AI with human analysts; together, they interpret the findings and decide which new product ideas to pursue. This human-machine team detected the rise of trends like vegan skincare and bold, unconventional hair colors 6–18 months before they became ubiquitous, giving L'Oréal a head start in product development.[2]

Even outside consumer goods, AI trendspotting proves its worth. Financial institutions use AI to scan regulatory news and industry reports, ensuring they aren't caught off guard by rule changes. In public health, the Canadian startup BlueDot used an AI algorithm to scour foreign news reports and airline data; it sent an alert about a mysterious pneumonia outbreak in Wuhan on December 31, 2019, nine days before the World Health Organization issued a warning.[3] Some of BlueDot's corporate clients gained a critical early heads-up. These examples demonstrate the power of the Observe stage: AI can cast a far wider net than any human, catching the faint signals of change so that you're aware of them early rather than blindsided later.

Peloton didn't miss the data; it misinterpreted it. During the pandemic lockdown, demand for home fitness surged, and Peloton's stock increased over 400% in 2020. Management saw this spike as the new normal. Even after gyms reopened, growth slowed, and the resale market for its machines grew, the company announced a roughly $400 million factory in Ohio in 2021. By mid-2022, Peloton had 500 days of unsold inventory. The stock dropped 76%, the

company posted a $439 million loss in 2022, and the Ohio facility was canceled entirely. Peloton has since recovered through cost-cutting and restructuring, focusing on the allure of its instructors. They have rediscovered their story beyond the data.

Step 2: Attack

Once you've gathered intelligence on what might be coming, it's time to stress-test your strategy against those possibilities. Step 2, Attack, is about adopting the mindset of your worst nightmare. This is where you role-play as the competitor that ends your career and your company. The attack is essentially poking holes in your own plans before your competition does. It also uncovers opportunities within uncertainty that are hidden in plain sight.

This stage begins by assembling a cross-functional team charged with answering a single question (see Figure 5.5). For example: "Given the driving forces found in the Observe step, if we were a fast-moving competitor (or an unforgiving market force), how would we exploit our company's weaknesses? How could we make our own business obsolete?" It can be an uncomfortable exercise, but it's enormously enlightening. Think of it as a fire drill for strategy: the discomfort of the practice is precisely what prepares you for the real blaze.



Figure 5.5 HorizonAI Framework: Attack

AI can serve as a creative sparring partner here. You begin by creating a "clone" or digital twin of your company. You could use a generative AI tool to brainstorm "10 ways to undermine our product in 60 days through business model innovation." Then use AI agents to simulate competitors attacking that twin. These AI agents might modify pricing, introduce alternative features, or target your most profitable customer segment to see what impact it has on your market share in the simulation. The goal is to surface vulnerabilities and unexamined assumptions. If every scenario your team comes up with still has your strategy winning, you probably aren't being imaginative enough in your attacks. This attack phase should be a psychologically safe space to consider "unthinkable" moves. Encourage wild ideas: What if a startup gave away for free what you charge for? What if a tech giant entered your industry? What if there's a data breach or a supply chain collapse? By wrestling with these scenarios now, you inoculate your organization against complacency. AI can make the exercise more rigorous by providing data-driven inputs: identifying which customer segments would be

most likely to defect if a cheaper competitor emerges, or varying key parameters (like a 20% price cut by a rival, or a new technology that halves production costs) and showing how your current strategy holds up.

Implications

Companies that deliberately "attack" themselves often emerge much stronger. Author Lisa Bodell advocates a practice she calls "Kill the Company," which is exactly what it sounds like. Akin to "red teaming" in cybersecurity, this allows us not only to find vulnerabilities in our business models and products but to see opportunities the competition hasn't spotted in the fog of uncertainty. "Kill the Company is about having an out-of-company experience—examining your weaknesses so that you can make them your strengths," Bodell explains.[4]

When leaders are given permission to drop defensive thinking, their war-room mentality kicks in, and they often pinpoint blind spots that wouldn't surface in normal meetings. An insurance firm doing a "Kill the Company" exercise realized their claims process was painfully slow compared to a hypothetical AI-powered startup; this spurred them to invest in automating claims before a real startup could scoop up frustrated customers. It's learning by doing, or rather, learning by counter-doing. Through these self-immolation exercises, firms anticipate competitor moves and identify potential threats in a risk-free setting. By simulating moves and counter-moves, managers get a visceral understanding of how a scenario could play out quickly, with resources at their fingertips and at low cost.

Amazon exemplifies strategic self-disruption. Jeff Bezos instilled a culture of "Kill the Company," epitomized by Amazon's push into cloud computing. When Amazon Web Services (AWS) launched, it effectively gave any startup the tools to compete with Amazon's own retail business. This was calculated self-cannibalization. Bezos knew that if a radical new platform was going to disrupt Amazon's retail operations, it would be far better for Amazon to do it first. AWS's success wasn't just about great technology; it was about Amazon's willingness to undermine its own old model in service of a bigger future.

The payoff? Amazon transformed from an online retailer into a cloud powerhouse, avoiding being overtaken in a space it could have easily ignored. AWS maintains approximately 30% market share in the global cloud market (Q4 2024), well ahead of Microsoft Azure (21%) and Google Cloud (11%). In 2024, AWS sales reached $107.6 billion, a 19% increase from 2023; in 2025, Amazon plans to invest $100 billion in AWS.[5] The lesson: by boldly role-playing a "would-be usurper," firms can discover a transformative move, or at least shore up defenses against those plotting their downfall.

Step 3: Orient

After observing the terrain and pressure-testing your strategy through attack scenarios, it's time to Orient: map out the strategic implications and decide how to position your company considering the scenarios you've envisioned. In this phase, you reconcile the insights from step 1 and step 2: "Given what we've observed, and the vulnerabilities we've uncovered, what are our best strategic options? What would we do if scenario A happens? What about scenario B or C?" This is the moment of synthesis and planning (see Figure 5.6).

Begin by identifying the key driving forces that organize your various attack scenarios. Typically, these forces surface in the Observe stage. They will generally cluster around PESTEL factors: political, economic, social, technological, environmental, and legal trends, with particular attention to those with the potential to disrupt or fundamentally redefine your business model. This initial exploration is deliberately broad, encouraging leaders to look beyond their immediate competitive context. After identifying potential driving forces, the method employs a critical prioritization process. Each force is evaluated against two criteria: impact (the magnitude of its potential effect on the industry) and uncertainty (the unpredictability of its trajectory). The two forces that score highest on this combined assessment, those with the greatest potential impact coupled with the highest uncertainty, are selected as the axes for a 2 × 2 scenario matrix (see Figure 5.7).

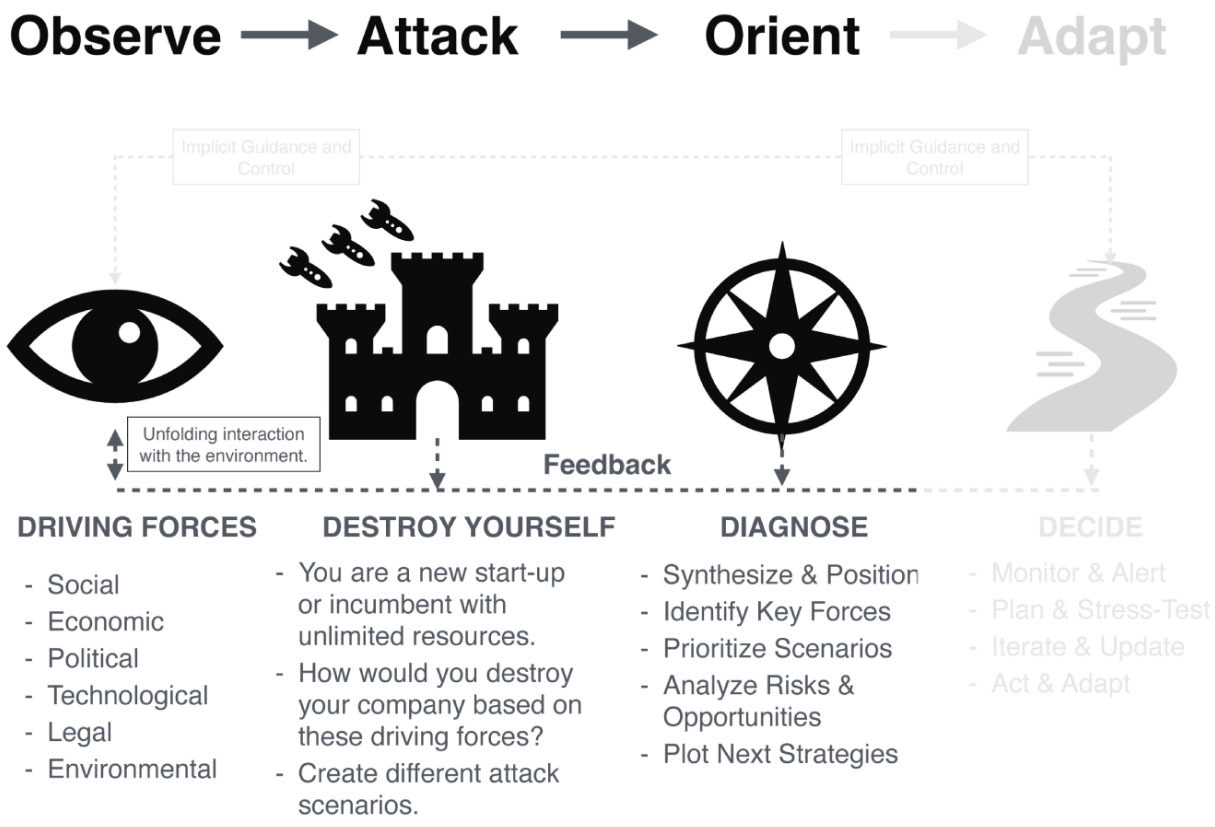


Figure 5.6 HorizonAI Framework: Orient

This disciplined selection process focuses the strategic conversation on truly consequential uncertainties rather than merely interesting trends. The resulting matrix creates four distinct quadrants, each representing a plausible future scenario characterized by extreme variations of the selected forces. A financial services firm might explore scenarios at the intersection of AI disruption (slow versus rapid) and regulatory control (strict versus laissez-faire). A consumer goods company might examine futures shaped by changing customer preferences (traditional versus digital-first) and environmental regulation (weak versus strong).

Then begin plotting your attack scenarios on various 2 × 2 key driving forces matrices. For each scenario, identify opportunities and risks. In a scenario where a certain technology becomes dominant, the risk might be that your current product line becomes outdated, but the opportunity could be offering a complementary service for that tech. In a scenario where new regulation hits the industry, the risk is compliance cost, but an opportunity might be selling a solution to others who struggle to comply.

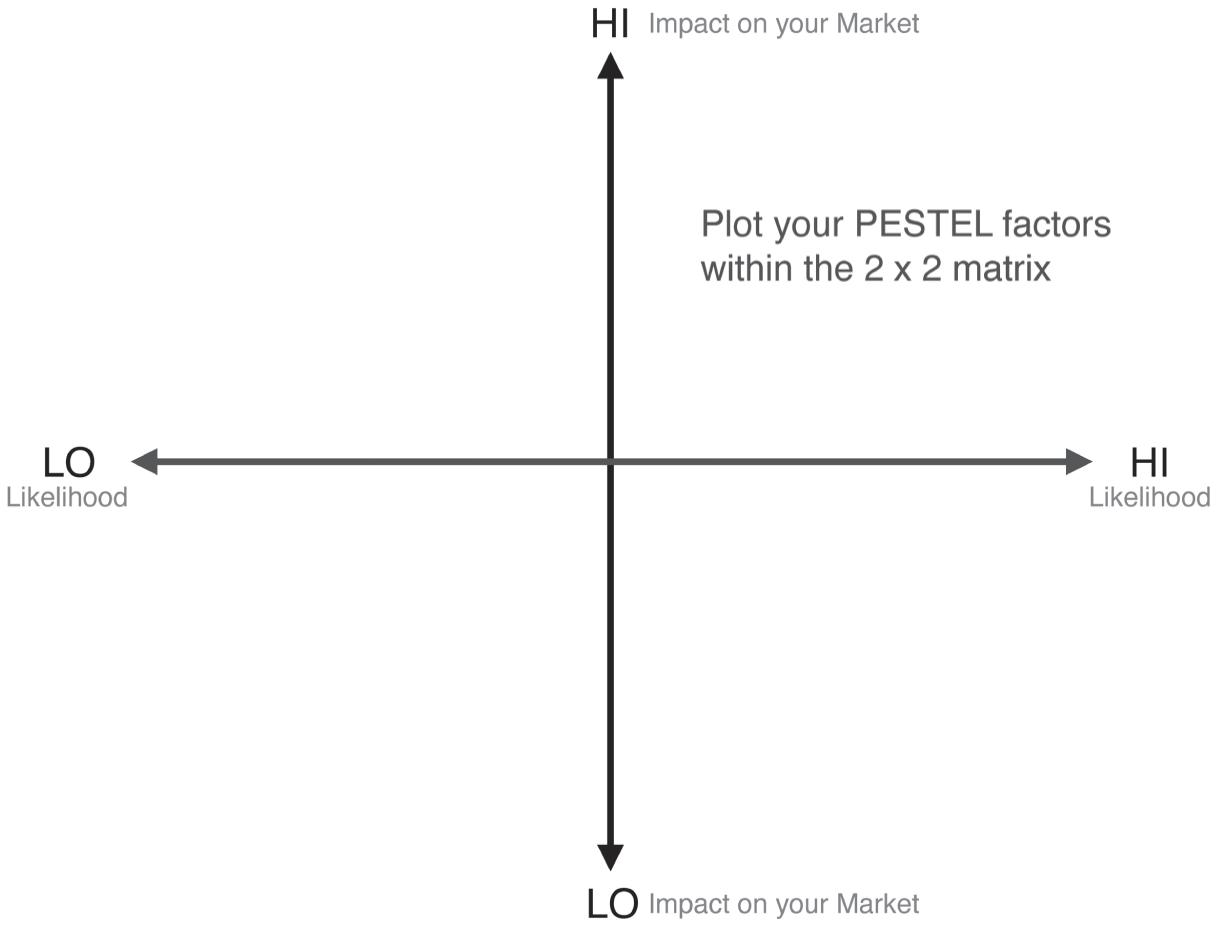


Figure 5.7 HorizonAI Scan: 2 × 2 matrix

Consider an example of the HorizonAI 2 × 2 Matrix in the financial services industry in Figure 5.8.

Q1: Rapid innovation but potential risks from regulatory gaps.	**Q2: Rapid innovation tempered by robust governance frameworks.**
Rise of decentralized finance (DeFi) Monopolistic power by first movers Breakthroughs in autonomous financial advisory & predictive finance Surge of cybersecurity threats Rise in litigation over data privacy breaches	Stability as firms balance rapid innovation with compliance Elevated cost burdens on companies Emergence of RegTech innovations Clear frameworks for algorithm audits, ethical standards, and accountability Slower innovation in AI
Q3: A stagnant environment with limited innovation and oversight.	**Q4: Highly conservative with cautious, incremental innovation.**
Policymaker inertia leads to an outdated regulatory landscape Slow adoption of incremental, cost-saving AI Public skepticism around AI's value in finance Minimal investment in advanced AI research Ambiguities persist around AI liabilities and data protection	Strong political focus on risk management & consumer protection Stable but slow economic growth in financial services High public trust due to strong consumer protection laws Innovation confined primarily to compliance-driven solutions Legal risks are low but potentially stifling innovation and experimentation

Figure 5.8 HorizonAI Scan: Financial Services and AI

AI is extremely helpful at this stage for quantitative stress-testing and forecasting. You can feed scenario assumptions into AI-driven simulation models (sometimes called *digital twins* of your business) to see how outcomes change. Suppose one scenario is a 25% increase in raw material prices due to a supply shock. An AI-based model can rapidly recalculate your cost of goods, profit margins, and even suggest which product lines would become unprofitable. Another scenario might involve a competitor capturing 15% of your market,

where you can simulate how that would affect your sales and what cost-cutting or new markets might offset the loss.

By applying this method across multiple scenarios, you stress-test your strategy. It's not about predicting exactly which future will happen; it's about being prepared for any of them. As Dwight D. Eisenhower said, "Plans are nothing; planning is everything."[6] Being oriented means you have a game plan (or two) ready in your pocket for whatever occurs. AI can also assist in assigning probabilities or identifying key indicators for each scenario. For example, AI might analyze real-time market data and suggest that Scenario B (such as a price war) is becoming more likely as competitors' prices trend downward. That could prompt you to start executing parts of your price war strategy proactively.

Implications

2 × 2 scenario planning is not new; what's different now is the toolkit. AI-assisted scenario modeling adds a level of analytical rigor that wasn't possible before. Already, 2,400 of the world's leading companies have adopted AI-infused planning platforms to connect strategic decisions to data-driven scenarios.[7] Researchers at IBM describe how generative AI can simulate complex scenarios for organizations, projecting the effects of various regulatory changes on a bank's operations. Consider the U.S. "Transition Integrity Project" (TIP) in the spring of 2020, a semi-annual 2 × 2 scenario simulation involving more than 100 current and former senior government and civil society leaders. Their key finding was the high likelihood that the fall election would result in contested results, constitutional crises, and the potential for political violence.[8] They did not predict January 6, 2021, but they had key figures oriented to the possibility. Unfortunately, they did not adapt, which is the next stage.

Toyota's leadership team has long employed scenario planning to navigate market uncertainties. After the 2011 tsunami disrupted their supply chains, they developed a comprehensive 2 × 2 framework mapping energy price volatility against global climate policy stringency. This approach guided their strategic investment in hybrid and hydrogen technologies while maintaining flexibility in

their manufacturing systems, helping them balance short-term profitability with long-term sustainability goals.

Contrast Toyota's disciplined reorientation with Nokia's catastrophic failure to orient. As we saw in Chapter 1, in 2007, Nokia held 50% of the global mobile smartphone market. Internal engineers flagged that its Symbian operating system was falling behind Apple's iOS and the emerging Android platform. The signals were clear. But Nokia's leadership couldn't reframe their mental model: they were a feature phone company, and smartphones were a niche. Management suppressed the warnings, fearing investor reaction. By the time Nokia finally pivoted to Windows Phone in 2011, the market had moved on. In 2013, Nokia sold its mobile division to Microsoft for a fraction of its former value. Nokia had the data. It had the engineering talent. What it lacked was the ability to orient, to see the world as it was becoming rather than as it had been.

The lesson: orienting your strategy with scenario insights and AI modeling makes you far readier to pivot or double-down when the world changes. You've already seen the movie in your mind (and spreadsheets), so you know how to act when it starts to happen for real.

Step 4: Adapt

The final step, Adapt, turns scenario planning into an ongoing, living process (see Figure 5.9). It's about integrating AI and the insights from Observe-Attack-Orient into your regular decision loops so that your strategy continuously evolves with new data. Think of the adaptation stage as having a radar system always on and a steering mechanism always ready to adjust course.

This means establishing continuous monitoring of the key indicators tied to your scenarios. If in Observe you identified drivers (say, a demographic trend or a tech adoption rate), in the adaptation stage you set up AI alerts for those drivers. If one of your critical uncertainties is "the cost of raw material X," you might deploy an AI system to monitor commodity prices and supply news in real time. The moment something deviates, a price spike or a news report of a mine closure, the AI flags it on your dashboard or sends an automated alert to executives. This triggers a predefined response or

at least a rapid strategy huddle. No major development catches you sleeping; your organization is primed to react because you've already thought through the scenarios in Orient.

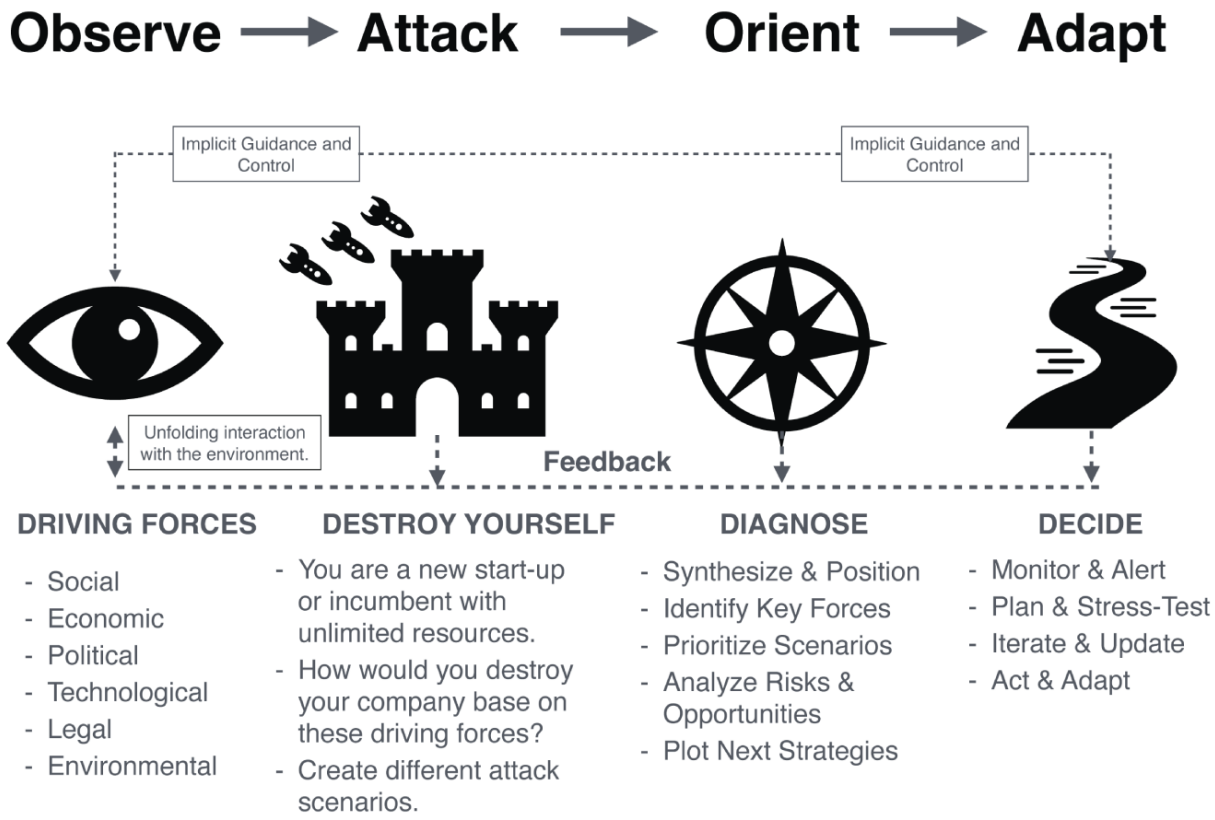


Figure 5.9 HorizonAI Methodology: Adapt

For each scenario, ask, "If this future came true, what would we wish we had been doing today to prepare?" This leads to identifying strategic implications and potential responses. Map out risks and opportunities in each scenario. Perhaps scenario A would crush demand for one of your product lines but open a new market in another area; scenario B might intensify competition requiring a different pricing strategy; scenario C might require partnering with a startup to access new tech. Once you have these implications, brainstorm strategic options: things you could do now or soon (invest in R&D, acquire a company, build a new capability, hedge a risk) to either seize the opportunity or mitigate the threat in each scenario. AI's analytical muscle helps here too. Use AI and quantitative models to stress-test each strategic option against the scenarios. Adapt also means creating a culture and mechanism for

rapid iteration. Encourage your strategy team to revisit scenario assumptions regularly, say, in quarterly meetings, and update them based on the latest observations. Maybe six months ago you considered a pandemic scenario far-fetched; after 2020, you certainly won't. Or perhaps your AI trend analysis now shows a new consumer behavior emerging that wasn't on your radar last year, which should become a new scenario to plan for. The strategy should be treated as a dynamic portfolio of bets, rebalanced as conditions change. Decision-making cycles should shorten, too. Instead of annual strategy reviews, some companies move to rolling strategic updates, where mini "Orient & Adapt" sessions happen monthly. AI tools facilitate this by providing on-demand forecasts and scenario analyses. Adapt is about institutionalizing agility: sensing changes early and responding in near-real-time, guided by the groundwork you've laid in the prior steps.

AI-Enhanced Scenario Planning: The Strategic Advantage

AI is revolutionizing scenario planning by transforming what was once a lengthy, periodic exercise into a continuous, data-driven process. This evolution offers four key strategic advantages:

Faster, data-driven foresight: AI collapses the timeline of traditional scenario planning from months to days or even hours. By analyzing vast datasets, from market reports to social media, AI detects patterns and weak signals that humans might miss. During crises like COVID-19, companies used AI to compress multiyear outlooks into days, enabling rapid adaptation to shifting conditions.

Expanded scenario creativity: Generative AI serves as a creative partner, producing diverse scenario narratives that spark human imagination. Need dozens of visions for how AI might reshape your industry? Large language models can generate initial storylines in minutes, pushing teams past standard thinking patterns and helping overcome cognitive biases.

Deeper analysis of complexity: AI's computational power explores countless "what-if" permutations at high speed. Finance teams use AI-driven scenario modeling to anticipate how various factors (interest rates, commodity prices, demand shifts) might combine to impact business performance. This reveals nonobvious interdependencies and creates more robust, evidence-backed scenarios.

Continuous monitoring and early warning: AI transforms scenarios from static documents into living frameworks by constantly monitoring scenario indicators. If a scenario depends on battery technology breakthroughs or regulatory changes, AI can track relevant data sources and alert leaders when signals suggest a particular future is emerging, turning scenario planning into an ongoing early warning system.

Throughout this process, humans remain essential for context, ethical judgment, and strategic direction. The most effective approach combines AI's analytical speed with human strategic wisdom: AI as a tireless analyst and creative assistant, humans as editors and decision-makers. Organizations that master this partnership not only prepare for external futures but simultaneously build internal AI capabilities, creating a more adaptive enterprise ready to thrive amid uncertainty.

* * *

Organizations comfortable making decisions under uncertainty have navigated massive disruptions most successfully. Consider recent shocks to the global supply chain, from COVID to the Suez Canal closure in 2024 to the latest Persian Gulf war. Companies with agile, data-driven strategy processes transitioned operations almost overnight, while others froze. Some retailers had AI-powered control towers monitoring their supply chains. These systems gave them real-time visibility as lockdowns began, AI highlighted which distribution centers were at risk and which products were surging in demand.

Businesses with such digital supply chain nerve centers incorporated findings into daily forecasts and operational dashboards. They rerouted shipments, found alternate suppliers, and avoided empty shelves. Those able to respond to demand spikes in an agile manner reaped significant rewards in market share during the crisis. Companies stuck in static planning mode were often caught flat-footed, unable to realign in time when conditions changed daily.

Adapting also means learning and evolving as you go. As your company goes through real disruptions or executes strategic moves, feed those results back into the AI models. If a particular scenario started to unfold and you responded, analyze how well your preparations worked. Did your forecast models accurately predict the impact? Did your contingency plan hold up? Use those lessons to refine future scenario assumptions or to train the AI on what happened versus what it expected. Over time, this creates a virtuous cycle: your strategic foresight becomes sharper, and your response playbook becomes battle-tested.

To institutionalize adaptation, many firms now embed AI into executive decision meetings. They might start with an AI briefing: "Here are the latest trend shifts and scenario probability updates since last month," akin to a weather report for strategy. With that context, leaders discuss strategic adjustments. Some organizations have created dedicated strategy innovation teams that function like an ongoing war room, constantly scanning and preparing responses. The bottom line: adaptation turns scenario planning from a one-off annual exercise into a continuous management practice.

The HorizonAI Framework (Observe, Attack, Orient, Adapt) offers a framework for proactive, AI-powered scenario planning that will not only make you and your organization more strategically agile but is also a powerful proving ground for you and your team to understand the power of AI. By systematically observing emerging signals, fearlessly attacking your own assumptions, orienting your plans to multiple possible futures, and continuously adapting as reality unfolds, your organization moves from reaction to readiness.

AI is the enabler, not the replacement, for human strategic foresight. The vision and judgment of leaders and strategists remain irreplaceable. What AI does is extend our peripheral vision and analytical reach. Think of it as adding a savvy co-pilot to the cockpit: AI handles the heavy lifting of data crunching and pattern recognition, presenting options and insights to human decision-makers, while humans ask the right questions, make the hard choices, and inspire organizations to move.

Notes

1. Aurore Legentil, "Case Study: How L'Oréal Uses AI to Predict Consumer Trends and Stay Ahead of Competitors," Ipsos Synthesio, The Social Intelligence Lab. https://www.thesilab.com/downloadables/case-study-how-loreal-uses-ai-to-predict-consumer-trends-and-stay-ahead-of-competitors.

2. Lisa Bodell, *Kill the Company: End the Status Quo, Start an Innovation Revolution* (New York: Routledge, 2012).

3. Bill Whitaker, “The Computer Algorithm That Was Among the First to Detect the Coronavirus Outbreak,” 60 Minutes, CBS News, April 27, 2020. https://www.cbsnews.com/news/coronavirus-outbreak-computer-algorithm-artificial-intelligence.

4. Bodell, *Kill the Company*, 47.

5. “Amazon and Microsoft Stay Ahead in Global Cloud Market,” Voronoi, February 28, 2025. https://www.voronoiapp.com/business/Amazon-and-Microsoft-Stay-Ahead-in-Global-Cloud-Market-4206.

6. Dwight D. Eisenhower, “Remarks at the National Defense Executive Reserve Conference” (speech, Washington, DC, November 14, 1957), The American Presidency Project, https://www.presidency.ucsb.edu/documents/remarks-the-national-defense-executive-reserve-conference.

7. Anaplan, “Turning Data into Decisions with AI Scenario Planning.”

8. W.J. Hennigan, “‘Plan for the Worst and Hope for the Best.’ Why Law Enforcement Officials Are Worried About Post-Election Violence,” Time, September 25, 2020.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 6
AI and Corporate Strategy: HealthCheckAI

> *"Excellence is never an accident. It is always the result of high intention, sincere effort, and intelligent execution."*
>
> —Aristotle

The tools that made companies great are starting to make them fragile. Six Sigma, balanced scorecard, and process standardization were built for a world of relative stability, where the goal was to eliminate variation, lock in efficiency, and then scale. The age of "we make it, you take it" is gone. Operational excellence still matters. But what it means, and how you achieve it, has fundamentally changed.

By Chapter 6, we don't have to remind you that AI is changing every aspect of your business. Artificial intelligence transforms operational excellence from a slow-moving quality-improvement initiative to a continuous adaptive source of competitive advantage. This work occurs through three fundamental shifts that only AI can enable at scale. **First**, AI expands the scope of operational visibility from sample-based to a comprehensive one. Where traditional approaches might analyze 1–5% of customer service interactions or production cycles, AI can process 100% in real time, identifying patterns invisible to conventional analysis.[1] **Second**, AI accelerates the feedback loop from retrospective to predictive. Instead of responding to quality issues after they occur, AI can identify the patterns that precede defects, enabling intervention before failure. **Third**, AI transforms operational responses from standardized to contextualized. Rather than applying one-size-fits-all solutions, AI can recommend optimizations tailored to specific conditions, customer segments, or market dynamics.

The result is an approach to operational excellence that improves existing processes and continuously evolves them as conditions change, thereby creating organizations that not only survive disruption but foresee strategic operational opportunities.

Yet, while AI presents countless solutions to many of these challenges, the modern executive faces a new paradox: even as AI promises transformative operational gains, many organizations struggle to realize its full value. To advance from simple AI pilot projects to pervasive adoption, companies must bridge a widening gap between AI's capabilities and their organization's culture and processes. Consider that 90% of executives at major companies report that people and process issues are the primary obstacles to becoming data driven.[2] In other words, the greatest challenge is not building the AI; it's getting your people to embrace it and your operations to seamlessly integrate it.

This is why we created an organizational excellence tool called HealthCheckAI. This approach turns the vision of maximizing the use of AI in operations into reality by weaving AI into the fabric of an organization's culture and daily operations. This framework transforms how organizations achieve operational excellence in today's uncertain business landscape. Instead of merely optimizing for known variables, it builds adaptive capabilities that allow companies to thrive amid constant change. This framework provides C-suite executives with a practical roadmap for strategic AI integration while cultivating the organizational culture needed to support it.

Two caveats from the top. First, every organization's industry dynamics, organizational structure, use of technology, etc., are different. Therefore, this tool is much more top-level, focusing on guiding principles, milestones, and objectives. The second caveat is that integrating AI into operational excellence does not falter because of flawed algorithms but due to human factors. The pattern is familiar: a company invests in a cutting-edge AI system, yet front-line employees underutilize it or managers mistrust its recommendations or the LLM's data set is outdated…and the list goes on.

Even the best AI operations strategy will underperform if the organizational mindset and practices remain stuck in old paradigms. A Boston Consulting Group study found that firms often underinvest in "softer" success factors, such as change management, skills

development, and new operating models that actually "prove to be the most critical success factors" for AI adoption.[3]

The Three Organs of HealthCheckAI

Most organizations still manage their operations like nineteenth-century physicians diagnosed patients: relying on infrequent assessments, limited data sampling, and reactive responses to problems that have already emerged. HealthCheckAI functions more like the human body's own intelligence: constantly sensing, adapting, and evolving without conscious effort (see Figure 6.1). HealthCheckAI transforms organizational operations through three fundamental capabilities to create "evolutionary intelligence." First, *omnidirectional sensing* establishes a comprehensive nervous system that simultaneously perceives operational patterns across multiple dimensions. Second, *adaptive simulation* serves as the organization's imagination, generating thousands of potential scenarios to reveal hidden vulnerabilities and opportunities that conventional planning would overlook. Finally, *continuous reconfiguration* embeds adaptation into everyday operations, enabling evolution through thousands of micro-adjustments rather than occasional massive overhauls. HealthCheckAI creates and maintains operational excellence through a three-phase approach that builds cumulatively, with each phase establishing the foundation for the next.

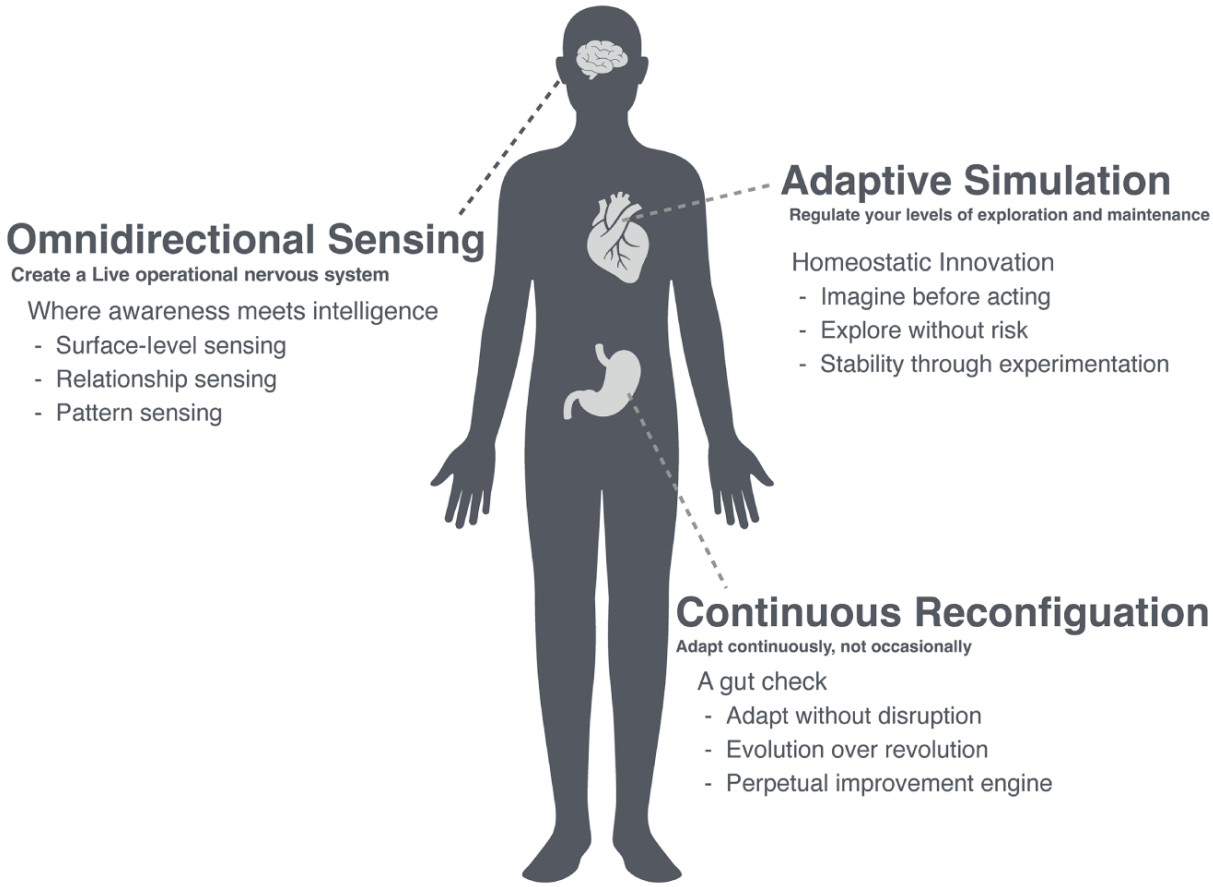


Figure 6.1 HealthCheckAI

Omnidirectional Sensing: Creating the Intelligent Nervous System

The first phase establishes “omnidirectional sensing” (see Figure 6.2). This is an organizational nervous system that simultaneously perceives operational patterns across multiple dimensions. Unlike traditional measurement systems that track predefined metrics, omnidirectional sensing captures the rich tapestry of interactions that constitute operations: communication patterns between teams, temporal relationships between processes, contextual factors affecting performance, and emergent behaviors that don’t fit existing categories.

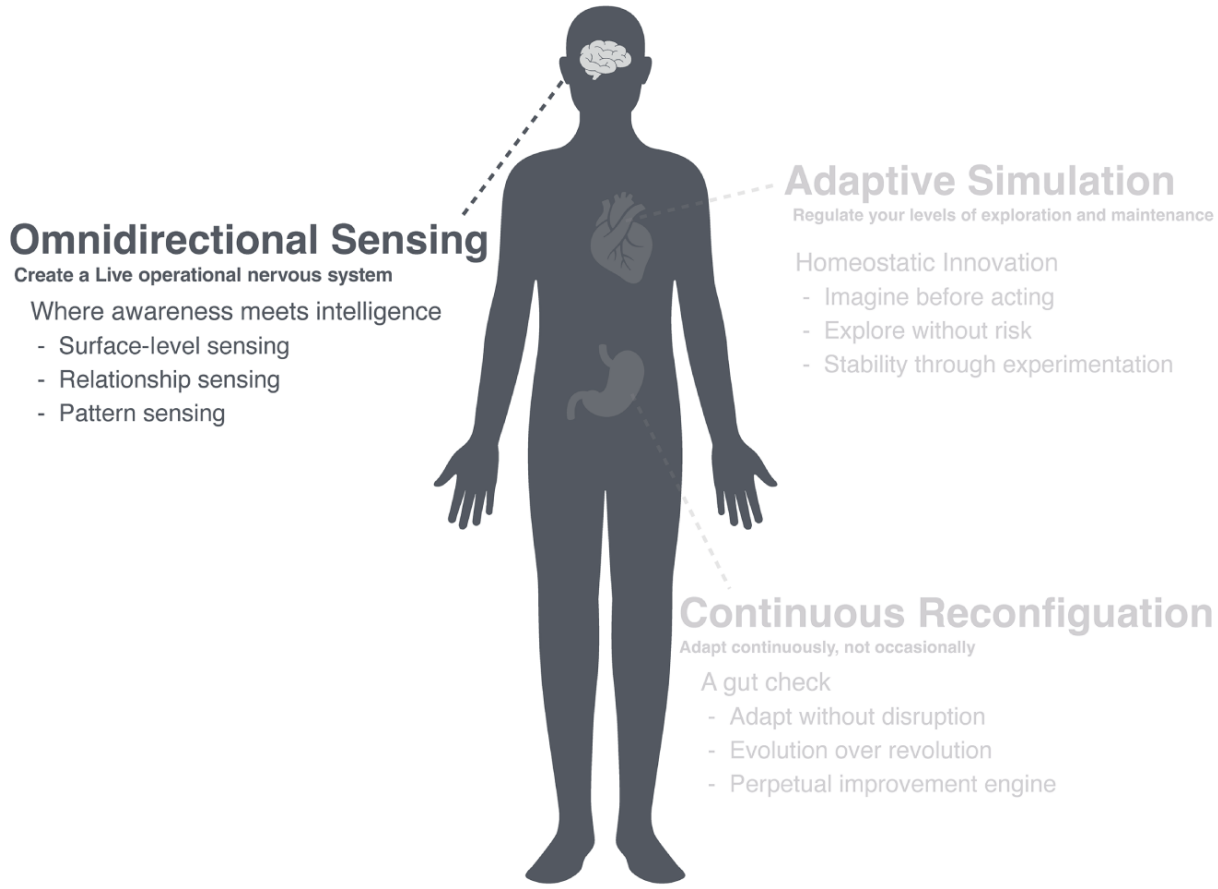


Figure 6.2 HealthCheckAI: Omnidirectional Sensing

This sensing occurs across three levels of increasing sophistication:

Surface-level sensing captures direct operational metrics: cycle times, defect rates, resource consumption, and output quantities. While similar to traditional KPIs, AI enables these metrics to be captured continuously across 100% of operations rather than through sampling.

Relationship sensing identifies connections between operational elements that may not be obvious: how changes in one process affect seemingly unrelated areas, how communication patterns correlate with performance outcomes, or how temporal factors influence quality in ways not visible in aggregate data.

Pattern sensing detects emerging dynamics before they manifest in surface metrics: subtle shifts in customer behavior

that precede churn, micro-variations in production that precede quality issues, or communication anomalies that signal coordination problems.

In the healthcare sector, Duke University Hospital offers a great example of "omnidirectional sensing" in action. Duke partnered with GE Healthcare to deploy a real-time AI command center (called Hospital Pulse) that monitors patient flow, bed availability, staffing, and more across the entire hospital in real time. This AI-driven nerve center gives leadership a continuous 360° view of operations above and beyond traditional systems' sampling or periodic reports. The result: Duke gained actionable live insights (e.g., instantly spotting an ER admissions surge or bottleneck) and achieved measurable improvements. This results in a 6% boost in overall productivity, a 50% drop in temp labor needs, and a 66% faster bed assignment process.[4]

Mastercard exemplifies this approach through its "Intelligence Center," which monitors more than 2.5 billion payment cards and 10 million merchant locations worldwide.[5] Rather than merely tracking fraud rates or transaction volumes, the system detected relationship patterns between legitimate and fraudulent transactions and identified changes in merchant interactions that signaled impending abandonment. The system doesn't just measure performance against targets; it perceives the operational context in which performance occurs, enabling responses tailored to specific situations rather than generic best practices. The shift: from "Are we meeting our metrics?" to "What patterns are emerging in our operational ecosystem?"

Adaptive Simulation: Creating the Organizational Imagination

With comprehensive sensing established, the second phase develops what we call *adaptive simulation*, or the organizational equivalent of imagination (see Figure 6.3). This capability allows companies to explore operational possibilities before committing resources, test hypotheses without disrupting performance, and learn from virtual failures without experiencing real ones. As discussed, traditional scenario planning typically explores a handful of future states, often

biased toward expected outcomes. Adaptive simulation leverages AI to generate thousands of potential scenarios, revealing unexpected vulnerabilities, hidden opportunities, and nonobvious connections that conventional planning would miss.

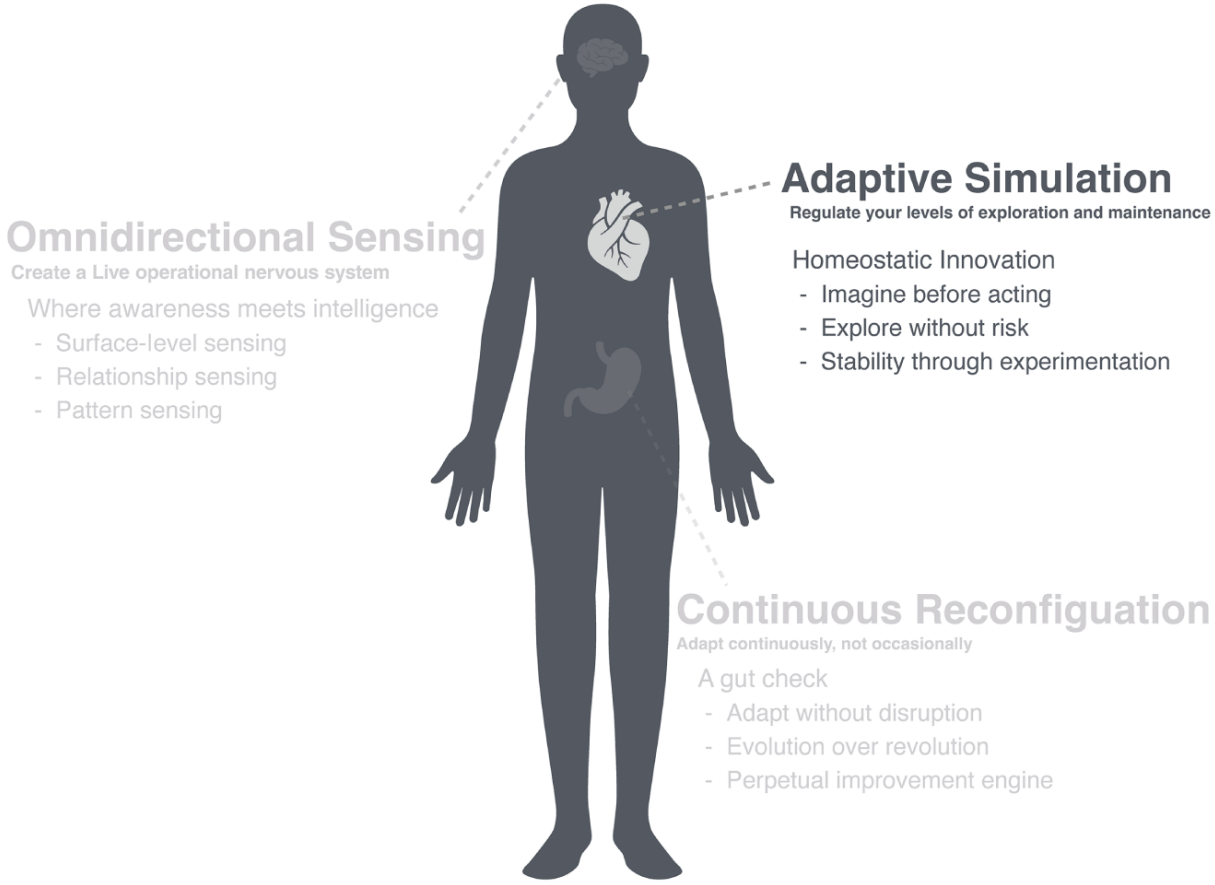


Figure 6.3 HealthCheckAI: Adaptive Simulation

This simulation capability operates at three levels of increasing sophistication:

Component simulation tests how individual operational elements might perform under various conditions: how a particular process responds to volume changes, how a specific algorithm handles edge cases, or how a team structure adapts to different demands.

System simulation explores interactions between operational components: how changes in one area ripple through others,

how resource allocations affect overall performance, or how timing differences influence system-wide outcomes.

Evolution simulation models how operations might adapt over time in response to changing conditions: how processes learn from experience, how capabilities might develop or degrade, or how operational patterns might self-organize in unexpected ways.

A novel example of adaptive simulation comes from the pharmaceutical industry. Gilead Sciences used an AI-driven launch planning platform (Unipr AI) to virtually simulate and stress-test its drug commercialization process before a new therapy's approval. In practice, the system ran predictive scenario planning models that identified likely bottlenecks (e.g., distribution delays, payer coverage issues) and suggested fixes before the launch. By leveraging AI to explore countless "what-if" scenarios, Gilead's team coordinated every facet of the product launch (regulatory, supply chain, marketing) for optimal readiness. This foresight paid off: when approval came, they seamlessly filled the first prescriptions within 24 hours, wasting no time on avoidable glitches.[6] In fact, Gilead cut the typical time from FDA approval to revenue by more than 50%, as the AI simulations had already prepared them for various contingencies.

Microsoft Azure's infrastructure operations demonstrate this approach through what they call "Digital Twins for Operations." Before implementing any infrastructure changes, they create virtual replicas of their data centers and simulate how proposed modifications would perform across thousands of scenarios, from standard operations to black swan events. When planning their European expansion, they didn't just model expected demand. They also simulated how different architectural decisions would respond to varied growth patterns, regulatory changes, sustainability requirements, and geopolitical disruptions. This revealed nonobvious vulnerabilities in what initially appeared to be the optimal design, leading to a more resilient architecture that has weathered multiple unexpected challenges.

The shift: from “How should we optimize for expected conditions?” to “How will our operations perform across multiple possible futures?”

Continuous Reconfiguration: Creating Operational Evolvability

The third phase establishes “continuous reconfiguration” (see Figure 6.4). This is the capacity to evolve operations in response to changing conditions not as a periodic redesign but as an ongoing capability. Traditional operational excellence typically implements changes through episodic initiatives. These are projects with defined beginnings and ends. Continuous reconfiguration embeds adaptation into everyday operations, creating systems that evolve through thousands of micro-adjustments rather than occasional massive overhauls.

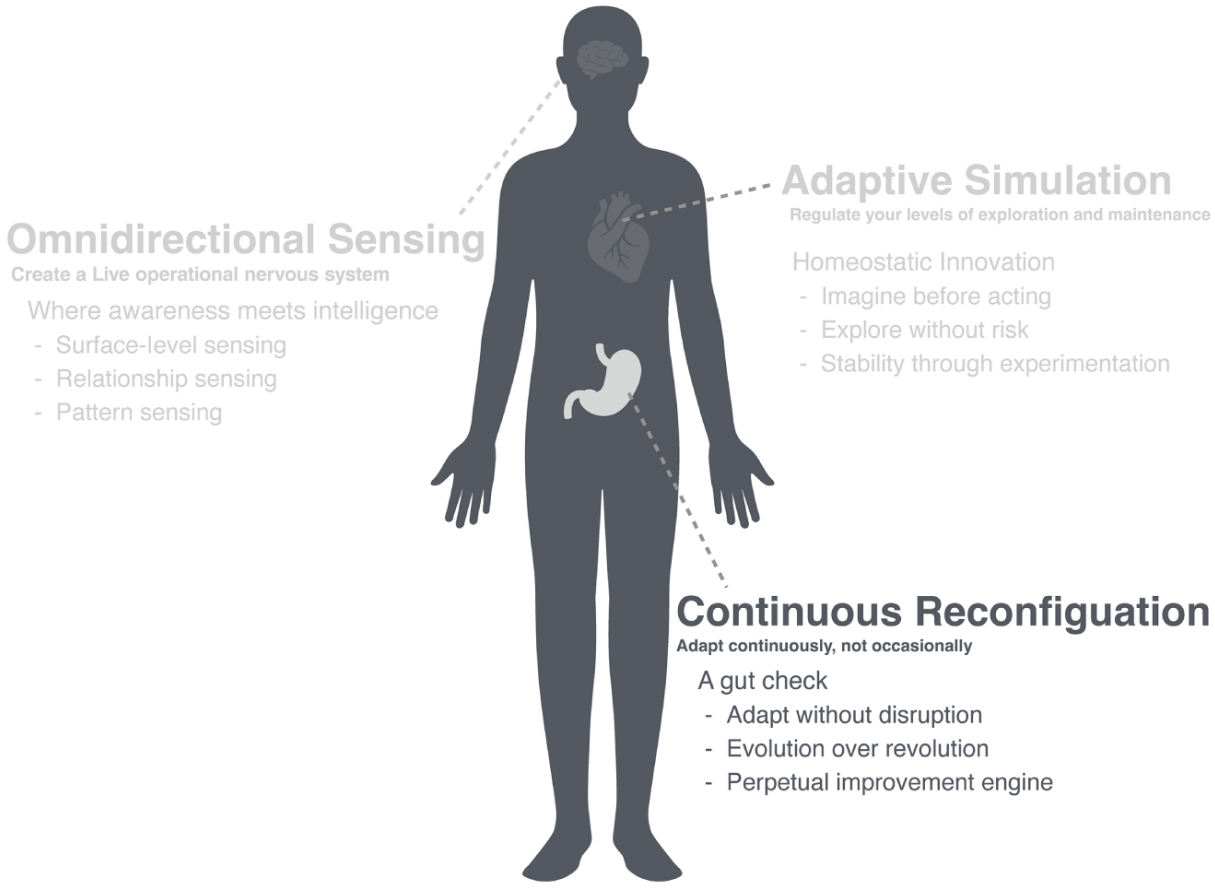


Figure 6.4 HealthCheckAI: Continuous Reconfiguration

This reconfiguration capability operates at three levels of increasing sophistication:

> **Parametric adaptation** adjusts operational variables within existing frameworks: reallocating resources, modifying decision thresholds, or adjusting process parameters based on current conditions.
>
> **Structural reconfiguration** modifies how operational components connect and interact: reorganizing workflows, reassigning responsibilities, or redesigning information flows to better match emerging requirements.
>
> **Capability evolution** develops entirely new operational capabilities or repurposes existing ones for emerging needs: creating new processes, developing new skills, or establishing new organizational structures in response to changing conditions.

Consider how advanced manufacturers are using AI to build ultra-flexible operations. In modern "smart factories," AI and IoT systems continuously analyze production data and can automatically adjust everything from scheduling to equipment settings on the fly. In fact, some leading manufacturers now let AI algorithms dynamically reallocate resources and reconfigure assembly lines based on real-time information rather than waiting for monthly process redesigns. If demand shifts or a disruption strikes, AI can immediately suggest (or implement) thousands of micro-adjustments rerouting tasks, tweaking workflows, and even repurposing machines to optimize productivity. A concrete example is Ocado's automated warehouses: fleets of robots, guided by AI, constantly adapt and optimize in real time to handle fluctuating online grocery orders. The system might rearrange picking sequences or deploy extra bots to a busy zone and then scale back when volume eases, all without human micromanagement. These under-the-radar cases show how embedding AI-driven adaptability enables operations to evolve continuously, not in occasional leaps, allowing companies to pivot quickly and thrive amid change.

Siemens Digital Industries also demonstrates this approach in its manufacturing operations. Rather than optimizing production lines

for maximum efficiency, they design for “adaptive production” and manufacturing systems that continuously reconfigure based on changing demands. Their AI-driven production scheduling doesn’t just allocate resources efficiently; it experiments with alternative configurations, learning which arrangements perform best under different conditions. During the pandemic, this enabled them to pivot from automotive components to ventilator parts within days. This was a transformation that would have taken months under traditional approaches.

The shift: from “How do we implement the optimal design?” to “How do we create operations that continuously evolve toward excellence?”

Together, these capabilities shift the fundamental questions organizations ask: from “Are we meeting our metrics?” to “What patterns are emerging in our ecosystem?”; from optimization for expected conditions to performance across multiple possible futures; from implementing the optimal design to creating operations that continuously evolve toward excellence. This isn’t merely another analytics dashboard or strategic planning tool. HealthCheckAI can serve as your organizational operating system, one that enables your business to develop the same self-regulating Intelligence that has allowed natural systems to thrive through billions of years of unpredictable change.

The Four Domains of HealthCheckAI Implementation

Successfully implementing HealthCheckAI requires more than technology, it demands a fundamental shift in operational philosophy grounded in four core principles.

Proactive Strategy, Not Passive Analysis

Traditional operational excellence relies predominantly on lagging indicators, or metrics that reveal problems after they’ve affected performance. HealthCheckAI reverses this relationship by creating leading indicators through simulation, enabling intervention before issues materialize. This proactive stance isn’t about predicting the

future with certainty but preparing for multiple futures simultaneously. By running thousands of simulations across varying conditions, organizations can identify which operational elements remain robust across scenarios and which become vulnerabilities under specific circumstances.

Cleveland Clinic illustrates this approach through what they call *anticipatory operations*. Rather than reacting to capacity constraints, their AI system continuously simulates patient flows, procedure durations, staffing levels, and resource requirements across thousands of potential scenarios. This allowed them to identify nonobvious operational vulnerabilities before COVID-19 hit. These are factors that wouldn't have appeared problematic under normal conditions but became critical during surge conditions. As a result, they were able to reconfigure operations proactively, maintaining critical services while many peer institutions struggled to adapt.

The key question: "What operational configurations will prove robust across the widest range of possible futures?"

Continuous Experimentation

Static operational models inevitably become obsolete as conditions change. HealthCheckAI embeds continuous experimentation into operational DNA, treating every activity as a production and learning opportunity. This goes much further than traditional A/B testing or pilot programs. It establishes what we call *experimental operations*, or systems designed to continuously test hypotheses about operational performance, automatically capturing insights and incorporating them into future iterations.

Zara transformed fashion retail through continuous experimentation in supply-chain management. Instead of relying on a fixed seasonal production plan, Zara's operations teams constantly run small-scale tests of new clothing designs and merchandising layouts in select stores. Each day, they analyze sales data, swiftly discontinuing items that underperform and expanding those that resonate with consumers. Zara conducts approximately 10,000 such design experiments each year, leveraging real-time feedback to rapidly adjust manufacturing and distribution decisions. When unexpected

disruptions occur, such as sudden shifts in consumer preference or supply chain challenges, Zara's experimental approach allows swift pivots, ensuring consistent responsiveness and competitive advantage in a volatile retail market.

The key question: "How do we design operations that learn automatically from experience?"

Redundancy over Perfection

The pursuit of operational perfection often creates hidden fragilities by eliminating margins for error, buffer capacities, and alternative pathways. HealthCheckAI prioritizes deliberate redundancy that builds on operational buffers and fail-safes that may appear inefficient in stable conditions but prove invaluable during disruption. This principle directly challenges the "lean" orthodoxy that has dominated operational thinking for decades. Rather than eliminating all slack, HealthCheckAI identifies where strategic buffers create resilience without sacrificing essential efficiency.

As we saw in Chapter 5, Toyota, ironically, one of the originators of lean manufacturing, demonstrated this principle during the 2011 Tohoku earthquake and again during the semiconductor shortage. While competitors had optimized supply chains to minimize inventory, Toyota maintained strategic buffer stocks of critical components based on AI-driven risk assessments. Their system didn't just build generic buffers but targeted redundancies where simulation indicated the highest vulnerability, a more nuanced approach that balanced efficiency and resilience. This allowed them to maintain production while competitors shut down assembly lines, gaining market share during disruption.

The key question: "Where does strategic redundancy create resilience that outweighs its efficiency cost?"

Human-AI Collaboration

The most sophisticated operational systems combine AI's pattern recognition capabilities with human judgment and creativity. HealthCheckAI doesn't replace human decision-making but enhances it, creating interfaces where each party contributes their

unique strengths. Integrating human insights with AI-driven analytics unlocks performance levels neither can achieve alone. AI excels at systematically recognizing nuanced operational patterns, rapidly testing hypotheses, and synthesizing vast data points focused on tasks impractical or overwhelming for humans alone. However, AI lacks contextual intuition, ethical judgment, and strategic vision. Humans complement these limitations, guiding AI applications toward relevant problems and interpreting AI-generated insights within broader operational contexts. This partnership maximizes efficiency, reduces errors, and enhances adaptability, leading to sustainable operational excellence rather than temporary gains.

This collaboration takes three forms:

- Augmented perception, where AI expands human awareness by identifying patterns too subtle or complex for natural observation
- Expanded exploration, where AI tests possibilities much further than what humans would typically consider, revealing nonobvious options
- Accelerated learning, where AI captures and synthesizes operational experiences across the organization, preventing the same mistakes from being repeated and successful innovations from remaining isolated

At Memorial Sloan Kettering (MSK) Cancer Center and partnering institutions, radiologists and AI systems collaborate to improve breast cancer detection accuracy. Research published in the *Lancet Digital Health* demonstrated that radiologists assisted by AI detected breast cancer more successfully than either working alone, with doctor-AI collaboration showing a 2.6% improvement in detection rates while generating fewer false alarms compared to physicians working independently.[7] The AI system pre-screens mammograms, flagging concerning patterns for radiologist review while handling routine cases more efficiently.

As Dr. Larry Norton, Medical Director of MSK's Evelyn H. Lauder Breast Center, notes: "AI programs can look at mammograms and identify areas that human radiologists may want to look at more

carefully," though he emphasizes it's "not going to replace a radiologist."[8] MSK radiologist Dr. Katja Pinker further explains the complementary relationship: "Because breast tumors can be heterogeneous, meaning that not all parts of the tumor are the same, a biopsy can't always give us the full picture," highlighting how AI-enhanced imaging provides valuable contextual information well ahead of what traditional methods capture alone.[9]

The key question: "How do we design interfaces where humans and AI each contribute their unique strengths?" The four domains establish the organizational conditions for success. What follows is the operational sequence for building those conditions phase by phase, capability by capability.

HealthCheckAI in Practice: A Phased Sequence

For executives considering HealthCheckAI implementation, we recommend a progressive approach that builds capability incrementally while delivering value at each stage:

Phase 1: Establish the Foundation (Months 1–6)

Begin by building omnidirectional sensing capabilities in one critical operational domain. This creates immediate value through enhanced visibility while establishing the foundation for more advanced capabilities:

1. **Instrument critical processes** with AI-enabled sensing that capture outcomes and contextual factors affecting performance.
2. **Develop pattern libraries** identifying the relationship between operational variables and performance outcomes.
3. **Create real-time dashboards** that visualize these relationships in ways that inform operational decisions.

The goal in this phase isn't perfect sensing but sufficient visibility to reveal previously hidden patterns where we generate insights while building the data foundation for simulation.

Phase 2: Develop Simulation Capabilities (Months 4–12)

With basic sensing established, develop simulation capabilities that allow exploration of operational possibilities:

1. **Build digital twins** of critical operational components calibrated against actual performance data.
2. **Develop scenario libraries** that reflect historical variations and potential future conditions.

3. **Create simulation interfaces** that allow operational leaders to explore "what if" questions about their domains.

The goal isn't to predict the future perfectly but to explore operational possibilities more systematically while identifying vulnerabilities and opportunities that wouldn't be visible through conventional analysis.

Phase 3: Enable Continuous Reconfiguration (Months 9–18)

With simulation capabilities in place, establish mechanisms for continuous operational reconfiguration:

1. **Implement feedback loops** that automatically capture performance data and update operational models.
2. **Develop reconfiguration mechanisms** that allow rapid adaptation based on changing conditions.
3. **Create learning systems** that identify and incorporate successful adaptations into future operations.

The goal is operations that evolve continuously rather than episodically adapting to changing conditions through thousands of micro-adjustments rather than occasional redesigns.

The HealthCheckAI Advantage

Organizations that successfully implement HealthCheckAI transform operational excellence from a static achievement to a dynamic capability, which we call *operational evolvability*. This capability delivers three critical advantages. First, it creates resilience without redundancy. The ability to withstand disruption without maintaining excessive buffers. Organizations can build resilience in critical areas while maintaining efficiency elsewhere by identifying precisely where strategic redundancy matters most.

Second, it enables adaptation without disruption. By embedding experimentation and reconfiguration into everyday work, leaders can adjust ahead of market shifts rather than chase them. Third, it builds intelligence, not just automation. AI is used not only to execute tasks but to strengthen organizational learning. By capturing insights across operations and simulating possible futures, it helps organizations develop judgment that exceeds any one individual's experience.

In a business environment of volatility, uncertainty, complexity, and ambiguity, operational excellence can no longer be defined by efficiency alone. The most successful organizations will be those that transform operations from optimized machines into adaptive organisms; these are systems that don't just perform well today but evolve continually for tomorrow. HealthCheckAI provides the framework for this transformation, allowing C-suite executives to build operations that aren't just excellent by current standards but remain excellent as they become vital members of their company's future strategy.

HealthCheckAI gives the organization a nervous system, an imagination, and the capacity to evolve. But there is one domain where the stakes of getting this wrong are existential: finance and risk. The principles are the same. The consequences are not. What follows is HealthCheckAI's logic applied to the domain where CFOs and CROs work, where sensing must catch fraud before it metastasizes, where simulation must stress-test the balance sheet before the market does, and where reconfiguration means evolving your risk posture in real time rather than updating a policy manual once a year. This is where operational excellence meets organizational survival.

AI-Driven Operational Excellence in Finance

In boardrooms today, two figures stand as the guardians of a company's organizational excellence and stability: the chief financial officer (CFO) and the chief risk officer (CRO)—introduced in Chapter 2 as The Risk Guardians. Tasked with safeguarding financial health and compliance, these "risk guardians" face evolving challenges.

Rapid market swings, mounting regulatory demands, and internal control complexities test their vigilance. At the same time, AI has burst onto the scene as a potential new source of risk. Now, a C-suite CFO or CRO may not seem to be the first in line to adopt AI in their operational practice, nor are they likely to be seen as kingpins in the strategy development and execution activities of a firm. However, with the proper and thoughtful incorporation of AI and the resultant organizational changes, they are at the forefront of using AI as a proactive risk mitigation and compliance tool and, moreover, as a key player in a firm's strategy.

Challenges Faced by Finance and Risk Executives

CFOs and CROs have long been responsible for ensuring financial stability, enforcing compliance, and maintaining risk controls across the enterprise. Traditionally, CFOs worry about fundamentals like economic volatility, liquidity, and accurate financial reporting. CROs focus on mapping out all forms of enterprise risk, from credit and market risks to operational and reputational risks, and keeping the company within its risk appetite. Both roles contend with external pressures such as uncertain economic conditions, fluctuating interest rates, and cyber threats. Finance leaders' top concerns range from inflation and labor costs to cybersecurity and an uncertain interest-rate environment. They must anticipate how geopolitical tensions or regulatory changes could hit the balance sheet. Additionally, CFOs are assuming an increasing responsibility for enterprise risk management (ERM); 55% of CFOs report owning ERM within their organizations, with risk-related duties expected to expand in scope.[10]

Maintaining compliance is another perennial challenge. CFOs/CROs must ensure financial operations adhere to evolving laws (tax codes, accounting standards, anti-fraud regulations) and that reporting is transparent and accurate. Failure can mean heavy fines or reputational damage. Yet traditional compliance processes are often manual and labor-intensive. The amount of time spent on checking transactions, auditing communications, and updating policies can strain even well-resourced teams. It's a delicate balancing act: CROs

today are expected to prevent losses and facilitate growth by taking calculated risks. They must provide a credible challenge to the business while encouraging profit and innovation. This dual mandate has expanded the scope of "risk management" beyond finance into conduct and ethics. As one risk executive observed, CRO focus has shifted from just financial perils to conduct, reputational, and ethical risks, with stakeholders demanding greater corporate integrity.

The rise of AI and other emerging technologies further complicates the picture. AI introduces new opportunities and new risks for these risk guardians. On the one hand, AI promises to crunch data at superhuman speed, uncover patterns humans might miss, and automate tedious control tasks. On the other hand, AI itself poses novel challenges around data privacy, model bias, and transparency. The CFO must worry not just about financial data but about who (or what algorithm) has access to it and how it's used. There's also the risk of erroneous or "black box" AI decisions: AI systems can sometimes hallucinate.

Such errors could lead to compliance violations or faulty risk assessments without proper oversight. In short, CFOs and CROs face a dilemma. They must harness AI's power to stay ahead of risks, even as they extend their control frameworks to manage the new risks AI itself brings. The following section proposes a way to resolve this dilemma: a blended framework of AI-driven automation and human oversight.

Adaptive Risk Intelligence Architecture (ARIA)

If HealthCheckAI is the organization's operating system for adaptive excellence, ARIA is its financial immune system. The same three capabilities—sensing, simulation, and reconfiguration—now operate within the specific constraints of regulatory compliance, fiduciary responsibility, and enterprise risk appetite. Omnidirectional sensing becomes the AI sentry scanning transactions, communications, and market signals for anomalies.

Adaptive simulation becomes the CFO's scenario engine, modeling how rate hikes, geopolitical shocks, or liquidity crunches will hit the balance sheet before they arrive. Continuous reconfiguration becomes a risk function that evolves its own controls—not through annual policy reviews but through the same kind of ongoing micro-adjustments that HealthCheckAI embeds across the broader enterprise.

What makes ARIA distinct is its operating environment. Unlike operational domains, where experimentation carries manageable downsides, finance and risk operate under regulatory scrutiny, audit requirements, and fiduciary obligations that demand a higher standard of human oversight. ARIA inherits HealthCheckAI's architecture but adds a governance layer calibrated to those realities.

Organizations need a structured approach that combines AI automation with human judgment to leverage AI effectively while retaining control. We suggest an AI-powered risk management framework that operates on two levels: (1) strengthening internal defenses and (2) enhancing external awareness, all under careful human oversight. This is the Adaptive Risk Intelligence Architecture (ARIA).

Under ARIA, a company provides its CFO and CRO with AI sentries that continuously monitor compliance, financial data, and market conditions—issuing alerts and surfacing insights, but never replacing human decision-makers. Like a fortress guard who sounds the alarm but relies on the commander to direct the response, these sentries extend human awareness without supplanting human judgment. C-suite leaders value this collaboration: the framework guarantees no threat goes unnoticed, yet no critical decision is made in isolation. It's a model of augmented intelligence where AI and human expertise work together to form a stronger defense than either could achieve alone.

The ARIA Framework is practical. CFOs and CROs can implement it gradually, deploying AI tools in high-risk areas like fraud detection, forecasting, and compliance checks, and integrating them into existing risk processes. Over time, these tools become an extension of the risk team. ARIA acts as a continuous audit, a predictive radar, and a compliance copilot. Meanwhile, the framework's governance

principles ensure this high-tech oversight aligns with the company's values and policies. It emphasizes transparency, so AI-driven decisions can be explained, and accountability, so humans remain responsible, not just the algorithms.

This gives boards and regulators confidence that the ARIA approach doesn't introduce unchecked risk but rather tames it. In an era where risks evolve fast and information overload is a reality, the ARIA Framework offers a way for organizations to stay one step ahead. It empowers the CFO and CRO to transform their risk functions from reactive gatekeepers to proactive strategy agents. With AI's help, compliance can become more assured, financial stability can be more attainable, and risk management can be more dynamic. The framework's successes will be measured in fewer surprises, faster decisions, and a stronger culture of compliance and risk-aware innovation.

Business leaders adopting the ARIA framework send a clear message: they are embracing the future of risk management. They recognize that AI, when used wisely, is a game-changer, enabling them to navigate complexity with greater confidence. Ultimately, managing risk is about ensuring the organization can weather storms and seize opportunities. By blending cutting-edge AI with seasoned human judgment, the ARIA Framework equips companies to do precisely that to guard against threats while boldly pursuing their strategic goals. It is a fitting strategy for the modern era, where the partnership of humans and AI becomes the ultimate guardian of enterprise value.

Strengthening Internal Risk Defenses (AI + Human Oversight)

Internal risks include compliance failures, communication breakdowns, and data integrity issues. AI can be an ever-vigilant sentry, monitoring these risks in real time. For example, AI systems can automate compliance monitoring by continuously scanning financial transactions and employee communications for red flags. CFOs are already leveraging such tools: AI-driven systems now process vast amounts of transactional data to identify anomalies that

may signal fraud or security breaches. Unusual payments, access attempts, or policy violations get flagged instantaneously. This always-on monitoring means potential issues are caught early and often before any damage is done. Crucially, alerts from AI are routed to human experts for review and action, ensuring no machine-generated false alarm leads to an improper decision. Companies achieve both speed and accuracy by automating first-line detection and using humans as validators.

Consider Behavox, a London-based AI company that helps companies strengthen internal defenses by monitoring employee communications for signs of misconduct or compliance risks. Its platform uses machine learning to sift through chat, email, and voice data to flag anomalies like damaging or inappropriate language, which can signal rogue behavior. Rather than replacing human judgment, Behavox augments it: AI-generated alerts are handed off to compliance experts, who review and investigate the most critical cases. This blend of AI vigilance with human oversight enables organizations to catch insider threats or policy violations early while ensuring that final decisions remain accountable to people.

As noted throughout this book, it is vital to use AI as a "copilot" role where AI handles the heavy lifting, but "continued human review and validation" is applied, especially for complex matters. Within internal operations, the framework targets several key risk areas:

Regulatory Compliance

AI monitors transactions, accounting entries, and access logs against regulations and internal policies. It can automatically check hundreds of conditions, from loan compliance to expense policy adherence. For instance, AI can continuously enforce loan agreement covenants by running automated compliance checks and notifying the team if conditions (like liquidity ratios) are breached. The framework requires compliance officers and auditors to regularly review AI findings, perform sample checks, and refine rules. This human-in-the-loop approach prevents blind spots and addresses AI's occasional errors. In our framework, any critical compliance decision flagged by AI is double-checked by a person.

Internal Communication and Conduct

Companies are deploying natural language processing (NLP) tools that scan communications (within privacy and legal boundaries) for signs of trouble. The AI might identify patterns suggesting collusion, harassment, or sensitive data being shared improperly. For example, artificial intelligence could flag an employee's accidental email of financial projections to an external address instantly. Our framework uses AI as a watchdog on internal communications, but with human oversight, compliance officers investigate alerts discreetly to determine if there's real misconduct or just a false alarm.

Data Integrity and Financial Controls

AI augments traditional controls by reconciling data and detecting anomalies faster than manual processes. AI algorithms can continuously audit entries in the general ledger, spotting duplicate invoices, out-of-range values, or suspicious adjustments. This is akin to having a tireless auditor checking transactions in the background. The framework pairs these AI detections with human follow-up. Finance staff receive anomaly reports, validate, and decide on corrective steps. By analyzing massive data volumes, AI reduces the risk of human error and helps humans focus on valid exceptions, while humans ensure the AI's outputs make sense in context.

Sharpening External Risk Radar (AI + Human Strategy)

CFOs and CROs must also manage external risks, from market volatility to shifting regulatory landscapes. Here, the framework deploys AI as an intelligence-gathering and predictive tool, as Prewave has done by offering an AI-driven "risk radar" that gives firms early warning of external threats in their supply chain so they can strategize responses. The platform continuously scans global news reports and social media in 50+ languages to detect events or trends that might disrupt suppliers. This ranges from natural disasters to factory fires to labor strikes or new regulations. It categorizes more than 150 types of supplier risks and sends timely

alerts to supply chain managers.[11] Armed with these AI insights, human teams can proactively adapt strategy, for example, by sourcing from alternate suppliers or adjusting inventory long before a brewing issue becomes a full-blown crisis.

Market and Financial Risk Analytics

AI excels at sifting through market data and finding patterns or early warning signals. Machine learning models can predict financial outcomes like cash flow or currency fluctuations. CFOs can feed in internal and external data and have the AI model simulate "what if" scenarios. Modern generative AI enables the precision modeling of countless financial scenarios. For instance, an AI might generate projections for how an interest rate hike or a supply chain disruption could impact earnings, giving the CFO a head start in contingency planning. The benefit is a more dynamic approach to external risk. When AI signals a possible storm, the team can adjust hedges, reallocate assets, or build buffers before the full impact hits. As McKinsey notes, AI can help risk functions move from a reactive stance to partnering with business units on strategic risk prevention.[12] Human oversight is key to avoiding overreliance on any single model. Of course, risk committees should review AI-generated forecasts critically and cross-check with qualitative and quantitative insights.

Regulatory Monitoring and Compliance Adaptation

In our AI framework, a regulatory intelligence component uses NLP to monitor news, government websites, and legal databases for relevant changes. When it finds a change that affects the business, it alerts the compliance team with a summary. Some leading banks use Gen AI to create virtual experts that answer employees' questions on compliance and scan for emerging risks. The framework could include an AI that flags a new regulation and suggests which policies need updating in response. Compliance officers and legal counsel remain in charge. Their charge is to verify the AI's interpretation and approve any policy changes. But AI reduces the lag between a rule change and the company's response, ensuring no regulatory shift goes unnoticed.

Other External Threats

The framework can extend AI surveillance to other risk domains: cybersecurity risks. AI can monitor threat intelligence and detect cyber-attack patterns, climate, and ESG risks. AI models might project the impact of extreme weather events on supply chains. By combining internal and external data, AI can spot correlations humans overlook. However, it's not a crystal ball that will foresee every black swan. The framework treats AI as an aid for scenario planning and early warning, not a substitute for human analysis and intuition.

Human Governance and Controls

Human oversight is woven in at every critical stage. AI provides speed and pattern recognition, but humans provide contextual understanding, ethical considerations, and accountability. Look at TruEra, an innovative startup focusing on transparency and accountability in AI-driven decisions through model explainability and governance tools. Its platform lets organizations dissect and audit their machine-learning models to understand why a decision was made and to uncover any unintended bias. For instance, Standard Chartered Bank used TruEra to vet its AI-powered credit scoring models, improving their quality and trustworthiness by identifying and eliminating unjust biases in the decision process.[13] By making complex "black box" models interpretable, TruEra ensures AI systems operate under effective human governance. C-suite leaders can enforce fairness, document how AI decisions are made, and confidently deploy AI in high-stakes areas under proper controls.

Setting Clear Rules for AI Decisions

The framework defines which risk decisions can be automated and which require human sign-off. For example, an AI model might automatically block a transaction matching a known fraud pattern, but any high-value anomaly is escalated to a person. The CFO and CRO should approve these rules and update them as needed.

Ensuring Explainability and Transparency

To trust AI outputs, executives and regulators will ask, "Why did the model flag this?" Our framework incorporates explainable AI (XAI) tools to illuminate black-box models. For instance, if an AI flags an expense report as fraudulent, XAI might highlight that it was because the amount was outside the employee's usual spending pattern. CFOs should work with IT and data science teams to prefer AI solutions that offer interpretability. Regular internal audit reviews of AI outputs also help ensure AI models remain accurate and aligned with policy.

Robust Data Controls and Security

Since AI will ingest sensitive data, strong data governance is non-negotiable. CFOs and CIOs must enforce strict access controls to protect confidentiality. If using third-party AI tools, safeguards (encryption, anonymization, contracts restricting data use) need to be in place so that no confidential information leaks. Additionally, as AI could potentially generate realistic fake outputs, companies should bolster verification steps for transactions and communications.

Integrating AI into the ERM Framework

Leading practitioners advise embedding AI risk management into the overall enterprise risk framework. This means the same risk taxonomy, appetite, and reporting should incorporate AI-related risks and AI-enabled controls. Companies ensure consistency and avoid gaps by aligning AI initiatives with existing risk governance. CFOs and CROs should champion a culture where AI is seen as a powerful extension of the risk toolkit but always subject to the same rigor of oversight and ethical standards. Implementation is achievable with today's technology. Companies can start with pilot projects in high-value areas where AI solutions are mature. Early successes build confidence and buy-in. The end state is a risk management function where AI handles analytical lifting and routine surveillance while humans focus on strategy, exceptions, and decision-making, shifting from hindsight to foresight.

Current AI Tools and Technologies in Risk Management

Many components of an AI-driven risk framework are already available to CFOs and CROs, with tools rapidly moving from experimental to operational:

Anomaly detection and fraud prevention: Machine learning models learn standard patterns of transactions and user behavior and then flag deviations. These systems employ neural networks to catch subtle signals like dummy vendors or transactions under approval limits. Modern ERP software often includes built-in AI analytics for continuous auditing, reducing detection time from months to minutes and saving millions in prevented losses.

Financial forecasting and analytics: AI-powered predictive analytics combine internal data (historical sales, production) with external factors (market trends, economic indicators) to produce more accurate forecasts. AI-driven financial planning and analysis (FP&A) tools automatically update projections as new data arrives and run simulations to show potential outcomes with probabilities. Some organizations report improved cash flow forecast accuracy and better anticipation of liquidity needs.

Credit risk and underwriting models: Machine learning assesses credit risk using hundreds of data points in addition to traditional metrics. These models have demonstrated the ability to expand credit access while reducing default rates. Companies use similar approaches to manage receivables risk, identifying which clients might delay payments. The challenge remains to ensure fairness and explainability to meet regulatory requirements.

Regulatory compliance monitoring: "RegTech" tools track compliance requirements and detect breaches. Large banks employ AI systems that analyze transactions and communications for signs of money laundering or market

abuse. NLP scans legal documents to identify problematic clauses, saving legal teams significant time. Standard Chartered's AI-powered compliance system reduced breaches by 40% and enabled real-time transaction surveillance.

Risk aggregation platforms: These systems consolidate data from various risk silos to provide a coherent view. The vision is an "AI-powered risk intelligence center" offering automated reports, real-time dashboards, and policy updates. Components include business intelligence tools with AI capabilities, knowledge graphs mapping relationships between risks, and virtual assistants that answer risk queries in plain language.

The set of strategic capabilities for AI-driven risk management is shifting from nice-to-haves to essentials. CFOs and CROs should familiarize themselves with these technologies, implement pilots where appropriate, and identify the gaps that pose the biggest challenges in their specific contexts.

Implementing the ARIA Framework: A Step-by-Step Guide for CFOs and CROs

CFOs and CROs can enhance risk management by integrating artificial intelligence into their frameworks, but it must be done thoughtfully. The ARIA Framework provides a practical approach to leverage AI for risk management with strong human oversight. This guide ensures AI augments compliance, fraud, and market risk decision-making while risk experts retain control and accountability.

Step 1: Assessment and AI Readiness

Begin with a thorough assessment of your current risk management capabilities and your organization's AI readiness:

- **Existing risk processes:** Review current frameworks, policies, and tools. Identify gaps where AI could improve

efficiency or insight.

- **Data infrastructure:** Ensure quality data and IT infrastructure are in place. Robust data governance is a prerequisite; poor data leads to flawed AI outcomes.
- **Governance and expertise:** Assess if proper governance structures and skills exist. Determine if teams are prepared to oversee AI initiatives and establish clear oversight roles.
- **Risk appetite:** Define your organization's risk appetite for AI-driven decisions. Align AI plans with strategic objectives and acceptable risk levels, setting boundaries for AI in high-stakes versus advisory roles.

By the end of this step, you should understand where AI fits into your risk management and have the foundational elements ready for deployment.

Step 2: Identifying Key Risk Areas

Map out the internal and external risk domains where AI can add the most value:

- **Regulatory compliance:** AI can track and analyze vast regulatory data, scanning contracts and policies to flag compliance issues and monitoring transactions and communications for signs of noncompliance.
- **Fraud detection:** A high-impact area for AI. Tools can spot unusual patterns in financial transactions that humans might miss, analyzing large datasets in real time to detect suspicious activities and send instant alerts.
- **Market risk:** AI-driven predictive analytics help anticipate market shifts by continuously digesting economic indicators, market news, and historical trends to warn of emerging risks like sudden price swings or liquidity challenges.
- **Operational disruptions:** AI strengthens operational risk management by analyzing sensor data and performance metrics to predict disruptions before they occur, forecasting supplier failures or system downtime.

Each area should have defined risk indicators that AI will monitor. Engaging business unit leaders ensures you capture both enterprise-wide and niche operational risks.

Step 3: AI Deployment Strategy

Develop a strategy for introducing AI into your existing risk management framework:

- **Start small:** Begin with pilot projects in high-priority areas to demonstrate value and learn lessons. Ensure each AI use case aligns with business goals and overall risk strategy.
- **Regulatory compliance by design:** Incorporate regulatory requirements from the start. Work with compliance and legal teams to interpret any industry-specific AI guidelines.
- **Integration into risk frameworks:** Map how AI tools integrate with existing processes. Specify how models retrain, how results are reported upstream, and how they align with risk escalation protocols.
- **Stakeholder buy-in:** Communicate the AI strategy to board members, executives, and business units to set expectations and gain support for necessary investments.

By the end of this step, you should have a concrete implementation plan that ties AI initiatives to risk management objectives while meeting regulatory and organizational standards.

Step 4: Human-AI Collaboration Model

Establish a human-AI collaboration model that defines how AI and risk teams interact:

- **Human in the loop:** Maintain human oversight on all AI-driven risk analyses, especially for high-impact decisions. Experts should review AI outputs for context before actions are taken.
- **Straightforward role assignment:** Define which tasks are automated and which require human judgment. Let AI handle

routine, data-heavy monitoring while humans address complex risk decisions.

- **Decision augmentation:** Position AI as decision support for risk teams. AI provides risk scores and alerts, but human risk managers interpret insights and decide on action plans.
- **Training and trust:** Train staff to understand AI outputs so they trust and effectively use these tools. Regular cross-team meetings help address questions about AI behavior and ensure alignment.

A balanced human-AI model leverages AI efficiency alongside human critical thinking. With proper oversight, AI-driven insights become actionable intelligence rather than black-box numbers, leading to better risk decisions and adherence to the company's risk culture.

* * *

This chapter presents a single argument at two levels. HealthCheckAI shifts operational excellence from a one-time initiative to a continuous, adaptive capability that detects emerging patterns across the entire ecosystem, runs thousands of simulations before allocating resources, and self-adjusts through constant micro-evolution instead of episodic overhauls. ARIA applies that same framework to the high-stakes areas of finance and risk.

It enables CFOs and CROs to move from reactive gatekeeping to strategic foresight—enhancing internal defenses, increasing external awareness, and maintaining the human governance that regulators, boards, and shareholders require. The two frameworks are not separate strategies. They are the same intelligence operating at different levels of the enterprise—one ensuring the organization adapts, the other ensuring it adapts responsibly.

Together, they create something more powerful than either could achieve alone. The result is a company that not only perceives disruption but also adapts and grows through it, all while safeguarding the financial integrity that drives ongoing development.

Notes

1. Ryan Stuart, “100% Coverage: How AI Is Revolutionizing Support Call Analysis,” Kapiche, n.d., https://www.kapiche.com/blog/ai-support-call-analysis.

2. Randy Bean, “Firms Must Overcome Human Barriers to Enable Data-Driven Transformation,” Forbes, January 2, 2020, https://www.forbes.com/sites/ciocentral/2020/01/02/firms-must-overcome-human-barriers-to-enable-data-driven-transformation.

3. Nicolas de Bellefonds, Tauseef Charanya, and Marc Roman Franke, “Where’s the Value in AI?“ Boston Consulting Group, October 24, 2024, https://www.bcg.com/publications/2024/wheres-value-in-ai.

4. Business Wire, “GE HealthCare Announces Major Academic Medical Center as First Customer to Adopt New Hospital Pulse Tile for Command Center,” October 17, 2024, https://www.businesswire.com/news/home/20241017581550/en/GE-HealthCare-announces-major-academic-medical-center-as-first-customer-to-adopt-new-Hospital-Pulse-Tile-for-Command-Center.

5. Bernard Marr, “How Mastercard Uses AI Strategically: A Case Study,” Bernard Marr, July 1, 2024, https://bernardmarr.com/how-mastercard-uses-ai-strategically-a-case-study.

6. Andy Mehrotra, “How Gilead’s Livdelzi Began Earning Revenue within 24 Hours of Approval,” Unipr, February 19, 2025, https://www.unipr.com/insights/how-gileads-livdelzi-began-earning-revenue-within-24-hours-of-approval.

7. Jay Mwamba, “CUNY & Sloan Kettering Researchers Prove Early Detection of Breast Cancer Is Possible Using AI,” The City College of New York, December 12, 2024, https://www.ccny.cuny.edu/news/ccny-sloan-kettering-researchers-prove-early-detection-breast-cancer-possible-using-ai.

8. Memorial Sloan Kettering Cancer Center, “2023 MSK Annual Report: Ending Cancer for Life,” 2023, https://www.mskcc.org/annual-report/2023.

9. Memorial Sloan Kettering Cancer Center.

10. Jeff Dunn, “How Does a CFO Manage & Mitigate Risk?“ CFO Selections, January 15, 2019, https://www.cfoselections.com/perspective/how-does-a-cfo-manage-and-mitigate-risk.

11. David Leggett, “How Audi Is Using AI as a Supply Chain Risk Radar,” Just Auto, June 18, 2021, https://www.just-auto.com/features/how-audi-is-using-ai-as-a-supply-chain-risk-radar.

12. Rahul Agarwal, Andreas Kremer, and Ida Kristensen, “How Generative AI Can Help Banks Manage Risk and Compliance,” McKinsey & Company, March 1, 2024, https://www.mckinsey.com/capabilities/risk-and-resilience/our-insights/how-generative-ai-can-help-banks-manage-risk-and-compliance.

13. Standard Chartered, “Partnering with Trueba to Tackle Unjust Bias in AI Assisted Decision Making,” August 4, 2020, https://www.sc.com/en/press-release/partnering-with-truera-to-tackle-unjust-bias-in-ai-assisted-decision-making.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 7
AI and Corporate Strategy: Beyond Blue Ocean Framework

> *"Great things are not done by impulse, but by a series of small things brought together."*
>
> —Vincent van Gogh

AI Beyond Blue Ocean Framework

Every ambitious CEO dreams of finding the next untapped market. Yet most never do. This is not for lack of ambition but because of a deep-seated addiction to competition. For decades, executives have fixated on rivals over customers, benchmarking, mimicking, and fighting for incremental share in crowded markets. Now AI has turbocharged that addiction. Armed with real-time competitive intelligence and machine-speed analysis, firms are investing billions not to see customers more clearly but to watch rivals more obsessively. As analyst Scott Galloway has observed, companies treat AI like Ozempic, a quick fix that promises transformation while masking the underlying disorder.[1] It is the fundamental disorder of competition. As Blue Ocean Strategy co-author Renée Mauborgne warns, "The more you look to the competition, the more you end up looking like the competition." AI hasn't changed that dynamic. It has simply made the echo chamber louder.

Machine intelligence is truly powerful. No serious executive disputes that. The question isn't whether to use it but how and where to apply it. Targeting competitors, AI enables faster benchmarking, clearer pattern recognition, and real-time monitoring of rival activities. It is the most advanced red ocean tool ever created. When focusing on customer frustrations, workarounds, and unspoken desires, it transforms into something entirely different.

This distinction is everything. Most firms today are focusing their AI investments in the wrong direction, not out of incompetence but out

of habit. Competitive intelligence is measurable, reportable, and boardroom friendly. Customer discovery is messier, slower, and harder to quantify. So organizations default to what AI does most visibly well: watching rivals. The result is faster imitation, not genuine innovation.

Breakthrough products and services have never come from watching competitors. In decades of strategy and innovation work across industries and continents, we have never witnessed a genuine Eureka moment. What we have witnessed, repeatedly, is something quieter: a leader who looked carefully at what customers were struggling with, combined what was already obvious in new ways, and produced something that made everyone ask the same question: why didn't anyone think of this before? That is not a data problem. It is a human perspective-taking problem.

Blue Ocean Strategy provides the direction. Developed by W. Chan Kim and Renée Mauborgne at INSEAD, it is a systematic framework for shifting attention away from rivals and toward untapped market demand by strategically recombining what an industry offers. When AI is integrated with Blue Ocean principles, it stops being a benchmarking accelerator and becomes a discovery engine. That is the premise of our AI Beyond Blue Ocean (BBO) Framework.

Five AI-Powered Strategy Tools

Consider the five interdependent frameworks and methodologies.

AI-Enhanced "As Is" Strategy Canvas

The Strategy Canvas is Blue Ocean Strategy's central diagnostic and prescriptive tool: a deceivingly simple graph that captures an industry's current state of play. It maps the key factors of competition on the horizontal axis and the offering level each competitor delivers on the vertical axis. The "as is" strategy canvas shows, in one chart, what you do, who you do it for, and how your offering compares to rivals. It is your entire strategy in one image.

It exposes how similar most players look and highlights any gaps. As Kim and Mauborgne describe it, the canvas lets you see "all the

factors an industry competes on and invests in, what buyers receive, and what the strategic profiles of the major players are." This "as is" picture often reveals that competitors fight over the same metrics, driving the industry into a bloody red ocean of marginal differentiation.

Most companies use AI to track what rivals are doing. The "As Is" Strategy Canvas redirects that energy toward something more valuable: understanding where the entire industry has stopped listening to customers (see Figure 7.1).

AI transforms the Strategy Canvas from a static snapshot into a dynamic, data-driven diagnostic. Instead of depending on infrequent market research and gut instincts, executives can use AI to monitor the competitive landscape in real time. AI tools can analyze competitor websites, product specifications, pricing, customer reviews, and financial reports and then automatically display each competitor's value curve on the canvas. Additionally, AI can help you investigate overlooked supply-side relationships with partners, suppliers, and other key interactions that influence your offerings. The result: an instant, real-time view of who offers what, at what cost, and how customers perceive those offerings.

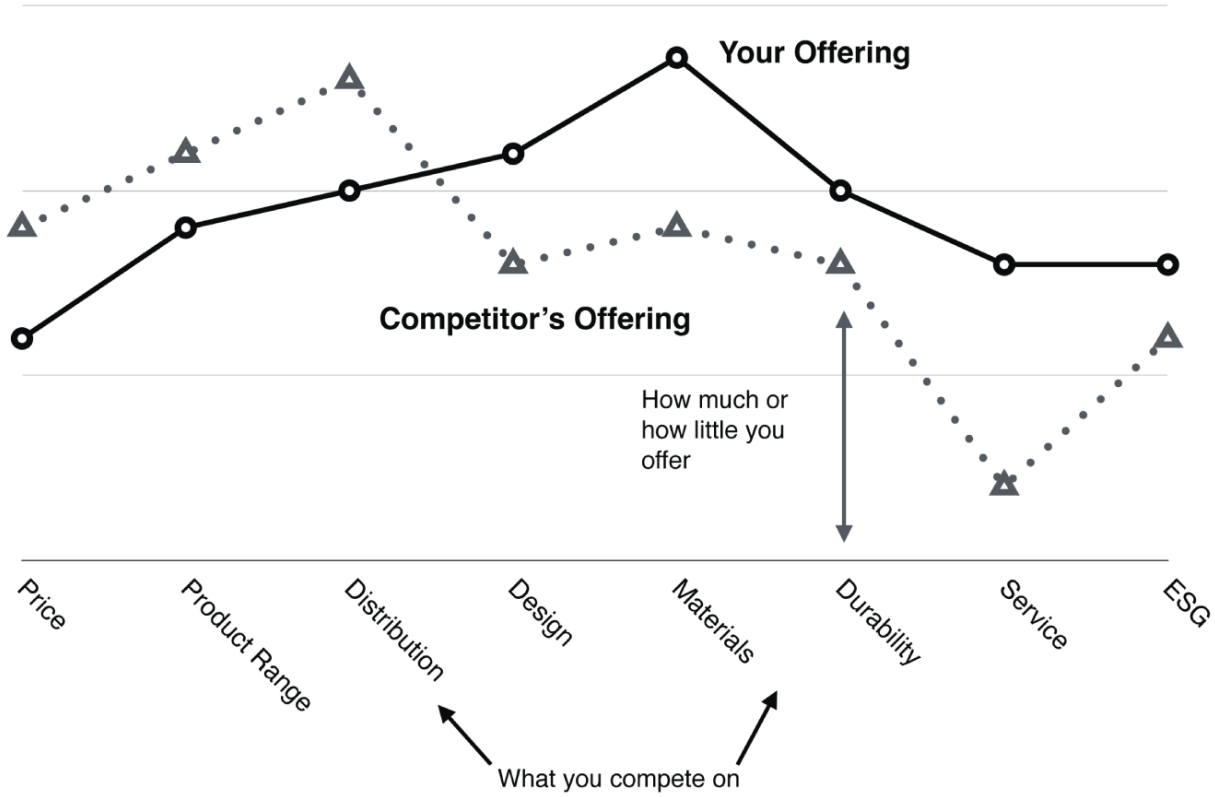


Figure 7.1 "As Is" Strategy Canvas: Overview

AI-driven sentiment analysis can mine thousands of customer reviews and social posts to reveal which factors customers truly value and where they express frustration. If all competitors invest heavily in a factor that customers ignore, the AI-enhanced canvas flags it as a candidate for elimination. If customers consistently complain about something that nobody in the industry addresses, the canvas highlights an unmet need that can be created or raised later in the process.

Think of it as a market X-ray. The AI-enhanced canvas can reveal that competitors' offerings cluster around a few common value dimensions, leaving others underserved. It can incorporate cost data to show which features drive higher costs without matching buyer value. This lets a leadership team spot where the industry's conventional wisdom may be wrong. Armed with this active snapshot, strategists can chart a course away from the competition, focusing their creativity on areas that could truly set them apart.

Dingo: Seeing What the Industry Could Not

Consider Dingo, a company we have worked with for years, using Blue Ocean Strategy to create and capture new markets over two decades. Founded in 1991 in Brisbane, Australia, Dingo began as a three-person software provider in a hard-nosed, asset-intensive business: mining. Their sole purpose was to monitor the oil quality of mining equipment. Dingo is proof of concept. What they built with data and discipline over two decades, the AI tools in this chapter can now accelerate for any company willing to ask the same questions.

By the early 2000s, when we began working with Dingo, their model was under siege. Competitors had replicated their software and expertise, undercutting them on price. Margins eroded. Market share slipped. The company that had pioneered a category was watching it become a commodity. This is the pattern that traps most innovators: build something new, watch the competition copy it, and find yourself back in a red ocean fighting over the same shrinking pie.

In 2005, after reading Blue Ocean Strategy, CEO Paul Higgins worked with his team to create their first market-creating strategic move. The team recognized that an alternative industry, healthcare, was shifting from reacting to ailments to managing health before a crisis. Applying this insight, his team created TRAKKA, a remote oil monitoring system that predicted future repair needs before breakdowns occurred. This gave Dingo a major competitive advantage and became the base for its future market-creating moves. From its inception, Dingo identified a critical gap in heavy equipment maintenance: the need for proactive, data-driven solutions to prevent unplanned downtime. By focusing on predictive maintenance, the company positioned itself ahead of industry trends, offering clients the ability to anticipate equipment failures and optimize maintenance schedules.

A few years later, when we began working directly with the Dingo team, the competition had once again caught up. They had copied Dingo's remote monitoring model. Margins and market share were under pressure again. This is the frustrating reality of innovation without strategic repositioning: you can win a battle and still lose the war. We began our work not by brainstorming or jumping into an innovation process but by taking stock of Dingo's current strategy using the "As Is" Strategy Canvas.

We charted their offering into first principles: asset expertise, on-site execution, mining experience, process expertise, real-time data, response to customer needs, measurable cost savings. Dingo still had an edge, but the competition was closing in fast. A traditional strategy would dictate that Dingo should match the competition point by point. But Blue Ocean Strategy aims to out-compete by offering a better value proposition at lower cost, preferably by creating an entirely new market.

Most revealing was what the "As Is" canvas exposed about Dingo's hidden strengths (see Figure 7.2). Mining firms were built like mines: vertically siloed. Dingo, by contrast, operated as a horizontal network across operations. Their real value was not just monitoring oil quality. It was their knowledge of the multitude of interconnections across the mining company's entire value chain. How to use this hidden strength? We launched an internal consulting practice with their biggest client. Oil monitoring remained the red thread, but Dingo would now help clients smooth out the horizontal connections from raw metal ore to market.

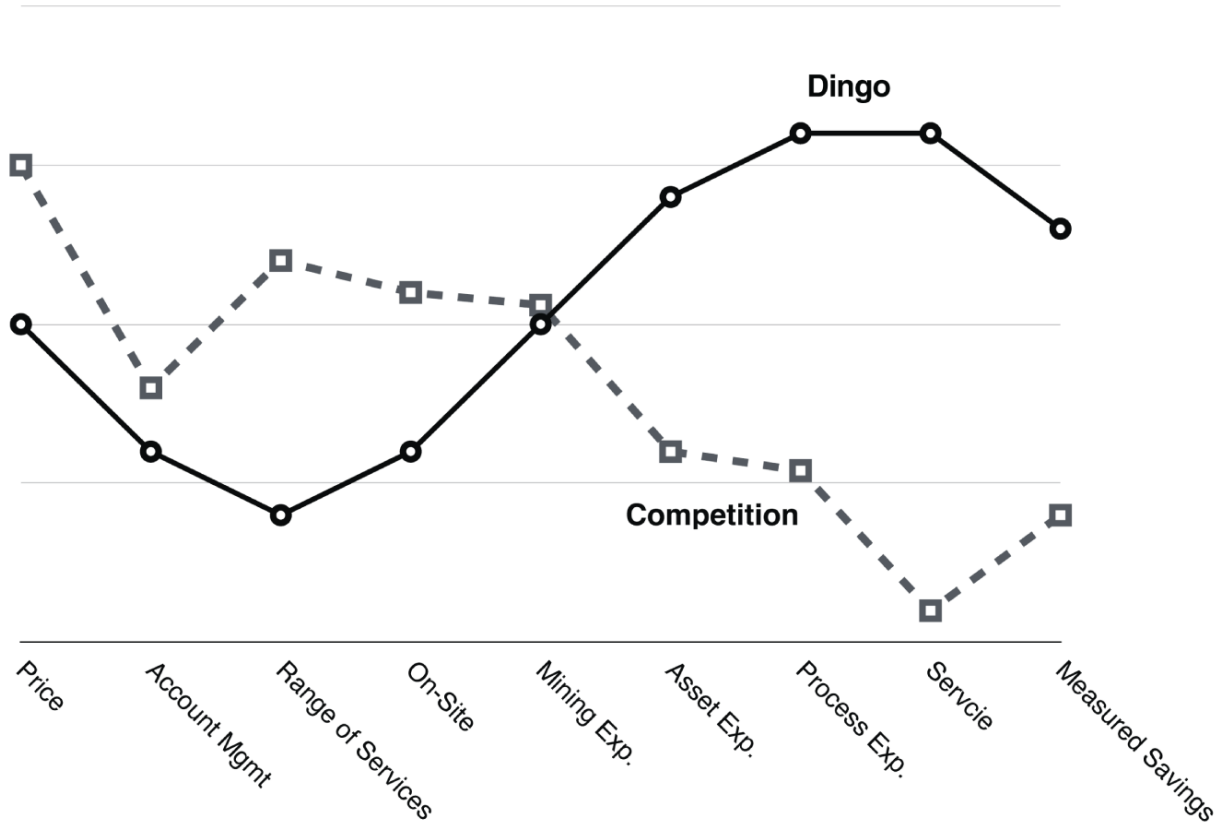


Figure 7.2 "As Is" Strategy Canvas: Dingo

Buyer Experience Audit

Anyone who has been involved in a B2B transaction knows how complicated and frustrating the process can be. Most of us walk away feeling merely satisfied about the transaction. But what if we could remove these pain points while reducing costs? The Buyer Experience Audit (BEA), based on Blue Ocean Strategy's Buyer Experience Cycle/Buyer Utility Map, illuminates hidden sources of value and cost reduction hiding in plain sight (see Figure 7.3). First, it maps the breadth of the buyer experience cycle, which is almost always wider than the narrow view held by most industries.

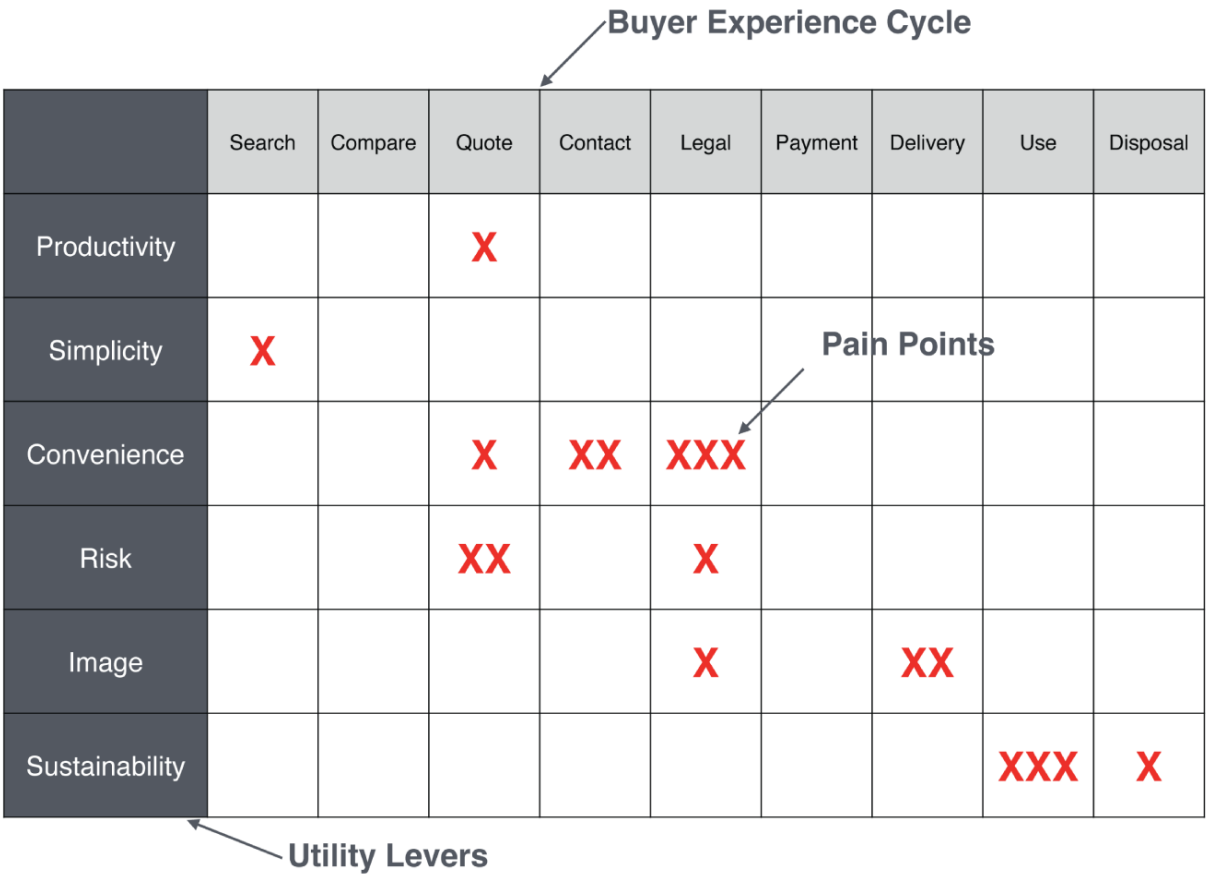


Figure 7.3 Buyer Experience Audit: Overview

Traditionally, companies take a narrow view of stakeholder interactions across their value chain, ignoring knock-on effects. Second, the BEA turns to the vertical relationships within each link of the supply chain: the ancillary players we rarely see but our colleagues, partners, and stakeholders must interact with to create value. Quite often, when we raise value for one stakeholder, we unwittingly create a pain point for another up or down the chain. Companies focus on creating utility at one stage of the buyer's experience and neglect the remaining stages.

Competitive intelligence tells you what your rivals offer. The AI-powered Buyer Experience Audit tells you what customers are silently enduring, and most importantly, what noncustomers won't tolerate at all. This is incredibly important in markets where you currently have no data, perspective, or understanding of those you never thought of before. AI gives it to you in the blink of an eye.

This forces customers to accept, and pay for, a certain amount of pain so they can get what they want. For noncustomers, the pain is too great, so they turn to alternatives. Once we create the architecture of the full buyer experience, we wind-tunnel each stage with the relevant actors to identify where we are blocking utility. These are based on six baseline utility levers:

> **Productivity:** At what stage are the biggest blocks to customer productivity?
>
> **Simplicity:** At what stage are the biggest blocks to simplicity?
>
> **Convenience:** At what stage are the biggest blocks to convenience?
>
> **Risk:** At what stage are we creating undue risks?
>
> **Fun & Image:** At what stage are the biggest blocks to fun and image?
>
> **Sustainability:** At what stage are the biggest blocks to environmental friendliness?

Think about the chemical industry. Companies with the most efficient operations rarely include chemical firms. Consider the complexity of a value chain extending from basic chemicals to consumer products: the endless interactions, the number of industries touched, the risks involved, the pain points created by countless moving parts. The Buyer Experience Audit helps even the most complicated supply chains see the full range of hidden friction that detracts from buyer value.

Let's review the BEA's components. Along the horizontal axis: the stages of the Buyer Experience Cycle. A buyer's experience generally breaks into six stages, running more or less sequentially from purchase to disposal. In chemical production, for example: raw material production, industry production, customer production, end customer production, use, and end-of-life. Along the vertical axis are the Six Utility Levers, the ways companies create utility for their customers. Most are straightforward. Simplicity, fun and image, and environmental friendliness need little explanation. Nor does the idea

that a product could reduce financial or physical risk. The most used lever, and perhaps the least obvious, is customer productivity.

Dingo: Shifting the Lens

In developing the first round of a Buyer Experience Audit, the Dingo team mapped out their buyer experience cycle from Need Identification to End of Contract/Renewal. Here, they identified well-known pain points that had become industry best practice, though friction remained. But when the team shifted perspective from “asset management” to their newfound role as “internal consultant,” the pain points they had previously ignored revealed themselves in multitudes.

The greatest blocks to utility were not within the narrow perspective Dingo had previously brought to the industry. Headaches in contracts. Difficulty communicating to nonengineers. Consider Dingo’s BEA in Figure 7.4. Gaps in training. Cost savings reporting that went nowhere. These problems stuck out like obvious targets once the team shifted their perspective to, one could say, their noncustomers. The pain points fell between lines of responsibility across the highly siloed mining operation. Knowing that solving an unarticulated pain point is the lowest-cost path to a new value proposition, the team began addressing these friction points to further widen their competitive advantage.

	Need Identification	Purchase	Mobilization	Implementaion	Supplements	Maintenance	End of Contract
Productivity		X	X	X	X	X	
Simplicity	X		XX		X		
Convenience	X	X					
Risk	X	X					
Image				XX			
Sustainability							

Figure 7.4 Dingo's Buyer Experience Audit (2009)

With the BEA, you first look at the full buyer experience cycle, well before and well after your direct interaction. Then, looking across the six utility levers from many perspectives, identify where the process blocks utility. Two questions: Where are the greatest blocks to utility across the buyer experience cycle for your customers and noncustomers? How can your offering effectively eliminate these blocks? When your offering passes this test, you are ready for the next step.

AI-Powered Six Paths Framework

AI trained on competitor data produces imitation more quickly. AI trained on the Six Paths produces genuine discovery because it is asking fundamentally different questions to fundamentally different people.

When we observe anything in the world, from a leaf in the wind to a car crash, we instinctively reach for six questions: Who? What? Where? When? Why? How? Aristotle was the first to map the circumstances of human observation in systematic terms. Rudyard Kipling gave them their most memorable form: "I Keep Six Honest

Serving Men, and their names are What and Why and When and How and Where and Who." The 5Ws + 1H became the backbone of journalism, investigation, and strategic inquiry for a reason. It is a built-in mental framework that cuts through ambiguity and drives rapid understanding. Applied to buyers, it surfaces the fundamental elements of utility people seek, which often exist at a pre-cognitive level, before desire hardens into a decision.

This is precisely why we don't apply it to existing customers, but to noncustomers. Customers, over time, become mirrors. They reflect what we already offer and reinforce the value propositions we have created with mirrored expectations. We lived with cans for 50 years without a can opener. A small reminder of why customers are rarely the source of true insight. Noncustomers sit outside this co-dependent relationship. Their distance is the point. They provide the level of abstraction needed to illuminate unarticulated demand. These are the wants that no current product addresses. AI gives us an unimaginable source of insight from noncustomers.

The questions are: Where do we compete? What do we offer? How much? To whom? Why? When? We call these the Six Paths Framework.

The Six Paths

The six paths are as follows:

> **Industries (Where).** Firms ask who they compete against, benchmark to competitors in the same industry, and seek to be the best within it.
>
> **Strategic Groups (How Much).** Within their industries, firms ask what level of value they provide for a given price (budget, economy, luxury) and seek to be the best within that strategic group.
>
> **Chain of Buyers (Who).** They ask who their key customer is. Most businesses converge on a single buyer group: either the purchaser (the parent in child medicine), the user (the child), or the influencer (the doctor).

Scope of Offering (What). Companies ask how they maximize utility at the moment of customer interaction, often ignoring the complementary products or services necessary for the buyer's intended task.

Functional/Emotional Orientation (Why). Based on history with buyers, firms cluster around an offering that is either functional or emotional.

Time (When). Within their industries, firms share common product lifecycle timelines, buyer experiences, and external threats.

Now that we have broken our assumptions to fundamentals on both the demand and supply side, how do we connect the right dots to reconstruct a breakthrough value proposition? Just as a builder needs scaffolding, our minds need support and constraint. Without restrictions, our solutions default to what has worked in the past. But when constraints are applied, we are liberated to think expansively and connect dots across conventions and expectations. Enter the Six Paths Framework (see Figure 7.5).

Dingo's Six Paths Framework

Path	Insights
Across Industries	Trade across to: Consulting services. Wind power. Railroads. Military.
Across Strategic Groups	Trade up from a commodity to strategic services.
Across Chain of Buyers	From a focus on project managers to the C-suite.
Across Complementary Services	From asset management to full enterprise services.
Across Functional & Emotional Appeal	With Dingo U, moved from functional to an emotional offering.
Across Time	The future is in renewables.

Figure 7.5 Dingo's Six Paths Framework

The Six Paths Framework is the interface where the entire universe of customers' fundamental needs and desires meets suppliers' ability to deliver. The latticework where we connect the dots that power innovation. By exploring these six paths, companies can discover untapped pockets of demand and redefine market boundaries.

AI is a natural partner for the Six Paths because it can process vast, diverse data sets far beyond human capacity. Consider the first path, looking across alternative industries: AI can analyze consumer behavior across categories to find common threads. An AI system might analyze transportation data and discover that many commuters alternate between driving and videoconferencing, an insight that helped one carmaker realize they compete not just with other cars but with Zoom calls. AI can cluster competitors by features and price for the strategic groups path, revealing gaps, perhaps an underserved middle market between high-end and budget players.

Along the chain of buyers, AI can sift through social media and forums to distinguish the sentiments of end users, purchasing decision-makers, and key influencers. This could reveal that the actual product user (say, nurses using a medical device) feels unmet needs the purchaser (hospital administration) isn't aware of, a chance to reorient the offering. For complementary products and services, AI can analyze purchase sequences to find pain points before, during, or after use. AI also excels at parsing functional versus emotional appeal by analyzing language across advertisements, reviews, and support calls. And with trends over time, AI can continuously monitor demographics, technology adoption, and regulatory changes to spot nascent shifts in consumer sentiment, the faint early signals humans might overlook.

In short, the AI-powered Six Paths Framework acts as a 360° radar for hidden opportunities. It looks beyond the obvious. AI algorithms comb through customer sentiments, competitor moves, and market adjacencies to surface insights that answer the six paths questions. The result is a set of unconventional opportunity areas that a human team might never have considered. With AI, the search for blue oceans becomes less about a eureka moment and more about systematically uncovering patterns in the noise, the patterns that reveal unmet needs and untapped demand ripe for growth.

Dingo: Finding the Next Frontier

When applying the Six Paths Framework to Dingo, we quickly found two things: (1) unexplored industries where Dingo's core competencies could deliver value well above the current offering, and (2) ways to learn from noncustomers and alternative industries to improve Dingo's value proposition. Using Path One (Alternative Industries), we asked: given Dingo's core competencies, what industry requires 24/7 monitoring and maintenance based on the monitoring of its lifeblood, oil? The military was our first thought. But the costly procurement process and heavy-handed management requirements didn't fit Dingo's culture. It was a dead end, and an important one. Not every path leads somewhere useful. The discipline is recognizing a mismatch early rather than forcing a fit.

So we combined Path One with Path Six (Time and Trends) to look at emerging, growth-focused industries that could benefit from Dingo's value proposition to help shape a rapidly growing sector. After a few weeks exploring noncustomers, the answer became clear: wind power. Wind turbines, like mining equipment, depend on oil-based systems that need constant monitoring. The fit was natural. The next step was to identify a new value proposition that preserved Dingo's core mining business while earning the right to expand into wind power asset management.

AI-Driven Eliminate-Reduce-Raise-Create (ERRC) Grid

Central to Blue Ocean Strategy is breaking the value-cost trade-off: pursuing differentiation and low cost simultaneously. The ERRC Grid ensures that we drive down costs by eliminating and reducing competitive factors that the industry invests in but that do not add significant buyer value. And it ensures we raise buyer value well above the industry standard by reconstructing industry boundaries to raise and create new factors. Competitor benchmarking tells you what everyone else values. The ERRC Grid asks whether any of it is actually worth keeping.

The actions of *eliminating* and *reducing* are especially important. They push companies beyond value maximization exercises with existing factors. Eliminating and reducing prompt companies to change the factors themselves, making existing competitive rules irrelevant. When you apply the ERRC Grid to the Strategy Canvas, you get a revealing new look at old perceived truths, creating new market realities that raise value while cutting costs. The ERRC Grid pushes you to ask all four questions and act on all four to create a new "To Be" value curve.

In short, the Grid asks: What factors can we eliminate? What can we reduce? What should we raise above the standard? And what can we create that the industry has never offered?

Integrating AI into the ERRC Grid sharpens the quality of every answer. AI can identify which factors are prime candidates for elimination or reduction. Machine learning can analyze customer usage data and feedback to find features people rarely use or don't value. If analytics show that 80% of customers never use half the

buttons on their microwave, those features could be eliminated or streamlined. AI can also benchmark industry cost drivers by analyzing financial statements across competitors, revealing processes that add cost without adding value. By spotlighting needless complexity, AI helps companies cut costs with confidence.

On the *raise* and *create* side, AI highlights where to deliver new value. Through ethnographic research, sentiment analysis and voice-of-customer data, AI can pinpoint unmet needs. An AI system might analyze millions of beauty product reviews and find that consumers consistently want personalized skincare, even though no brand offers it at scale. That insight could lead a cosmetics CEO to create a new personalized product line. Or consider an industrial equipment firm: AI analysis of maintenance logs might reveal that customers care far more about uptime and rapid repair than purchase price. That signals the company to dramatically raise service responsiveness, even if it means a higher initial cost. The company delivers outsized value on what matters most while trimming costs elsewhere to offset.

The AI-driven ERRC grid provides a fact base for the four actions. It can supply a ranked list of features that could be eliminated with minimal customer backlash or identify legacy processes that add expense without adding value. It can simulate cost savings by eliminating certain factors, helping to quantify the impact. Simultaneously, it can mine consumer data to suggest entirely new value factors, offerings competitors aren't providing, but customers would welcome.

By iterating through the Eliminate-Reduce-Raise-Create questions using AI insights, strategists can design a "To Be" value curve that differs dramatically from the status quo while remaining grounded in what the market will reward. The result: a strategy that breaks away from competition on both dimensions, offering a leap in value to buyers and a lower cost structure to the company. AI turbocharges the creativity of the ERRC Grid with data-driven validation, ensuring the bold moves chosen truly align with eliminating waste and adding customer value in equal measure. The Dingo ERRC Grid set the company on a path for continued market-creating moves (see Figure 7.6).

Dingo: Rebuilding Their "To Be" Strategy

Based on their Blue Ocean insights, Dingo retooled its value proposition and strategic focus around several key areas. By offering remote asset health centers and dedicated support teams, the company ensured that clients received timely insights regardless of geographic location. At the core of Dingo's operations was a data-driven culture that used AI-based machine learning algorithms and advanced analytics to refine its predictive models, improving accuracy and utility across industries. Specifically, Dingo rebuilt its value proposition around several moves:

Eliminate	Raise
Narrow Strategic Focus	Price Fun & Image Measurable Value Process Expertise Site-Wide Value
Reduce	**Create**
Range of Services Ad hoc Customization	Dingo U "Zero Risk"

Figure 7.6 ERRC Grid: Dingo

- **Mine-level results, guaranteed value:** Dingo shifted from a maintenance-focused to a production-and-cost-focused value proposition. They began delivering measurable financial value

(cost per unit) to their mining clients' C-suites, pulling multiple levers on both sides of the income statement rather than being limited to the single lever of component costs.

- **Experts in the field, from the boardroom to the break room:** Dingo realized customers valued independent expertise, focused data and research ("tell me something I don't know"), and access to expertise at the mine site. The focus became being world experts in technology-enabled business systems that turn oil analysis data into business results, bridging the gap from the boardroom to the breakroom.
- **Guarantee value:** Understanding that change is risky for customers and that noncustomers are willing to pay to avoid risk, Dingo developed guarantees to reduce customer risk and improve sales conversion. They also offered pilot programs so clients could see value at their own site with low exposure. A guarantee positioned Dingo as a confident partner to the C-suite, not just another asset management firm.
- **Ease of use:** Recognizing the complexity of industry processes and pain points left to the customer, Dingo focused on making complete, turnkey solutions that drive value and are easy to purchase and use.
- **Get it done, have some fun!:** Dingo identified that the condition-monitoring industry is often perceived as boring, intimidating, and emotionless. They aimed to create a sense of achievement, fun, and emotion for employees, customers, and suppliers, drawing on their Australian and American roots.

Specifically:

- **Eliminate:** Multi-condition, multi-industry offerings. Ad hoc customization.
- **Reduce:** Range of services.
- **Raise:** Price, ease of use, process expertise, site-wide measurable value.
- **Create:** "Z Risk" (Zero Risk), fun and emotion, Dingo University.

Today Dingo monitors wind turbine health data for top renewable energy companies across North America to prevent unplanned failures and maximize productivity. Their rules-based predictive analytics engine and proprietary turbine health database cover more than 25 GW of wind turbine capacity across 24 U.S. states and 10 Canadian provinces. They provide extensive benchmarking data on how to improve client energy output based on a global asset health database of more than 10,000 wind turbines. To date, Dingo has reduced industry downtime by 30%, leading to tens of millions of dollars in savings.[2]

They have maintained leadership in the mining industry and continue to be recognized as a world leader in predictive maintenance solutions. As Dingo moved forward in recent years, they looked back in time to explore established yet underserved industries where they could apply their core and growing competencies. With years of experience working across silos in the mining and wind power industries, Dingo offered a monitoring, consulting, and maintenance value proposition to the rail industry that it had never seen.

By shifting their focus from simple oil monitoring to a comprehensive preventative maintenance solution, they have become a world leader across mining, wind, oil and gas, rail, and defense. Today, Dingo serves clients across six continents, with $18 billion in equipment under management. It has saved companies more than $1 billion in maintenance and replacement costs, with typical payback greater than 3:1 within 12 months.[3]

"To Be" Strategy Canvas with AI Validation

Once a team has generated a promising new strategic direction, the final step is to visualize and validate it before execution. In Blue Ocean Strategy, this is the "To Be" Strategy Canvas: a new value curve showing how the company will diverge from competition by eliminating some factors, reducing others, and raising or creating new ones (see Figure 7.7). This future canvas is the strategic vision of the new market space.

The final test of any new strategic move is not whether it beats the competition. From a strategic perspective, it is whether the market

will come to it. AI validates that question before resources are committed. Even a brilliant strategy on paper carries uncertainty. How will the market react? Will customers embrace this new value proposition? Will competitors retaliate? This is where AI can play a transformative role, acting as a high-powered wind tunnel to stress-test the strategy before real resources are committed.

AI-driven simulation can project market demand for a new service at different price points, drawing on historical adoption curves and consumer demographics. It can model competitor behavior by ingesting data on how incumbents responded to previous disruptors, predicting whether they might drop prices, seek to imitate, or ignore the move entirely. It can estimate stakeholder reactions by analyzing news and social media to predict how customers or partners might respond to a radical new offering.

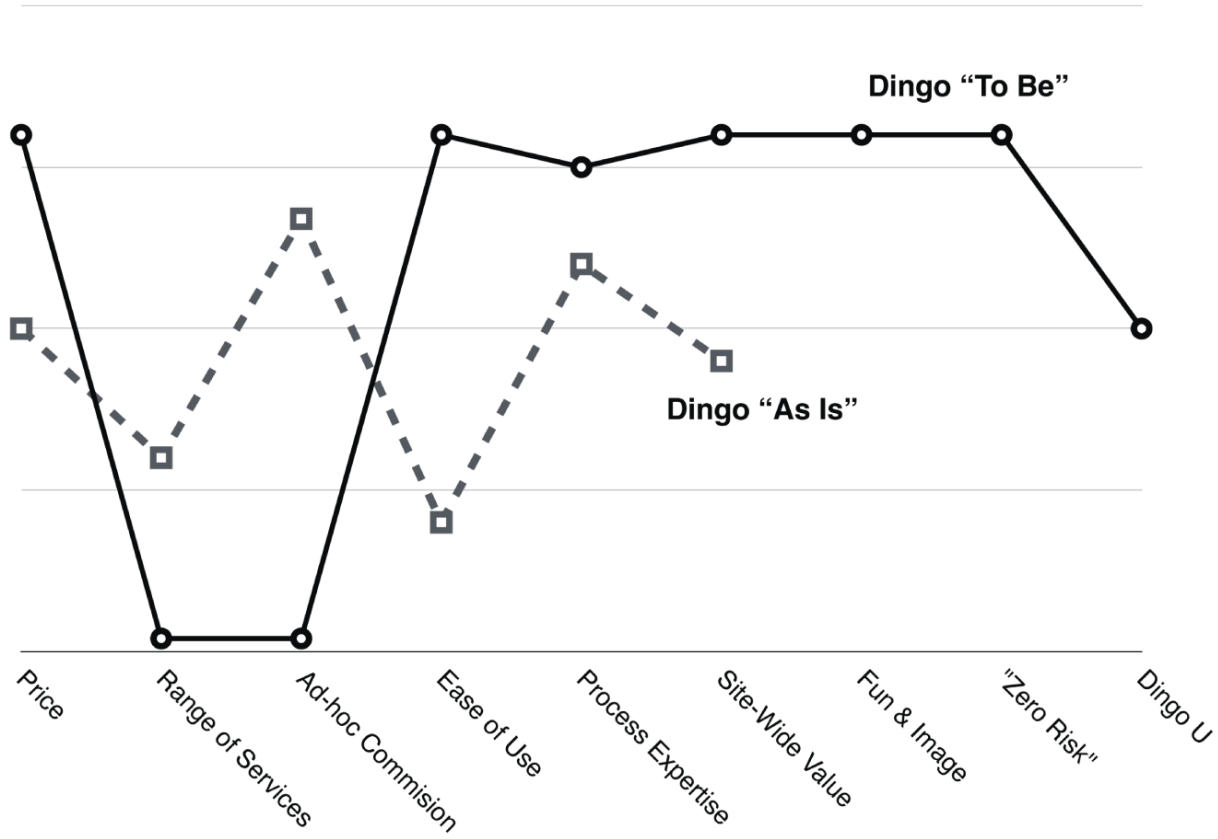


Figure 7.7 Dingo "To Be" Strategy Canvas

Artificial intelligence can also create thousands of hypothetical customer personas and have them "respond" to the new value

proposition based on known preferences. While no replacement for real market tests, this can surface obvious objections or identify which segments will likely be early adopters. AI-powered agents can run millions of simulations, varying features, pricing, and messaging, to seek an optimal combination that maximizes projected uptake.

As with all of the Frameworks presented in the book, think of the Beyond Blue Ocean Framework as a strategic flight simulator. It allows a company to pilot its new plane in virtual skies, make adjustments, and build confidence before launching into new markets. Using AI as a "simulator" enables strategists to model multiple market scenarios and stress-test strategies, adding more rigor to decisions before they become difficult to reverse.

AI can also help refine strategy execution. Natural language AI can draft A/B tests for different marketing narratives to see which resonates best with target customers. It can identify leading indicators to monitor once the strategy launches, the data to watch in the first 90 days to confirm that new market creation is on track. All this amounts to de-risking innovation. By validating the "To Be" Strategy Canvas with AI, leaders can proceed with a new growth move knowing they have data-driven evidence on how the market might respond.

* * *

The five AI-powered tools described—the Strategy Canvas, the Buyer Experience Audit, the Six Paths Framework, the ERRC Grid, and the "To Be" Canvas—form a single integrated framework for growth-oriented leaders. We call it the Beyond Blue Ocean Framework because: "Beyond" signifies AI's new learning and data-driven insights. "Blue Ocean" evokes the open, clear path to new market spaces.

Under this framework, executives still supply the vision, curiosity, and creative judgment. AI is the intelligent partner throughout. The AI-driven "As Is" Strategy Canvas provided a rigorous reality check on the status quo, and the Buyer Experience Audit illuminated hidden pain points. AI explorations across the six paths spark creative ideas grounded in evidence. AI analytics inform which insights will simultaneously raise buyer value at a lower cost. And AI

simulations validate the strategy before the portfolio of strategic options is chosen.

The Beyond Blue Ocean Framework combines human strategic thinking with machine intelligence to systematically find and capture new markets. It helps leaders grow organically, not by competing directly, but by creating new markets, with AI guiding them to uncover hidden markets. Embracing this framework can help companies break free from incrementalism and unlock the innovative potential that today's technology affords, steering the enterprise toward new horizons of profitable growth.

One of the authors worked with the Dingo team for more than 10 years. One of the authors was in the room in Brisbane when the competition was copying everything they built and margins were collapsing. Paul and his team resisted the instinct to fight back on price and choose instead to ask a different question: who isn't our customer, and why not? That single shift in perspective away from rivals and toward unexplored demand is what set everything in motion.

Not a sweeping moon shot or reinvention, but a disciplined, step-by-step application of the tools described in this chapter.

The Strategy Canvas revealed Dingo's core competencies. The Buyer Experience Audit surfaced what customers and noncustomers found painful. The Six Paths opened value propositions they had never considered. The ERRC Grid rebuilt its value proposition from the ground up. The "To Be" Strategy Canvas ensured their offerings created higher value at a lower cost. The process gave them the confidence to place the bet.

Three people in Brisbane monitoring oil samples became a global leader in predictive maintenance. We've worked with some of the world's leading companies at the intersection of AI and strategy, and the Beyond Blue Ocean approach has proven to be a consistent, practical and boundary-pushing methodology for creating new markets using artificial intelligence.

Notes

1. Allie Nawrat, “Scott Galloway: ‘AI Is Corporate Ozempic’,” UNLEASH, October 22, 2025, https://www.unleash.ai/unleash-world/scott-galloway-ai-is-corporate-ozempic.

2. Dingo, “30 Years of DINGO,” https://www.dingo.com/30-years-of-dingo.

3. Dingo.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Part III
Leading the AI-Powered Organization

Chapter 8
What to Know About AI

> *"There is an aspect of this which all of us in the field call a black box. You don't fully tell why it said this or why it got it wrong. We have some ideas, and our ability to understand this gets better over time, but that's where the state of the art is."*
>
> —Sundar Pichai, CEO Google

During the Industrial Revolution, as eighteenth-century Britain shifted from an agricultural to an industrial society, workers left the fresh air and seasonal rhythms of farm life for dimly lit factories filled with the constant clamor of machinery. To ease the transition, factory executives invented the gin trolley. Every other hour, a cart rolled up and down the factory lines, distributing the cheap, potent spirit and offering workers some liquid detachment from the mind-numbing repetition of spinning cotton and forging metal.

Today's executives face a similarly disorienting shift (see Figure 8.1). AI promises significant increases in efficiency and insight but also brings uncertainty and disruption. The temptation to seek quick fixes or superficial engagement, akin to the fleeting comfort of the gin trolley, risks distracting leaders from the deeper transformations required to harness AI's full potential. As industrial workers ultimately adapted by developing new skills and mindsets, today's executives must embrace comprehensive change: a willingness to continuously learn, rethink organizational models, and fully integrate AI into their strategic decisions.

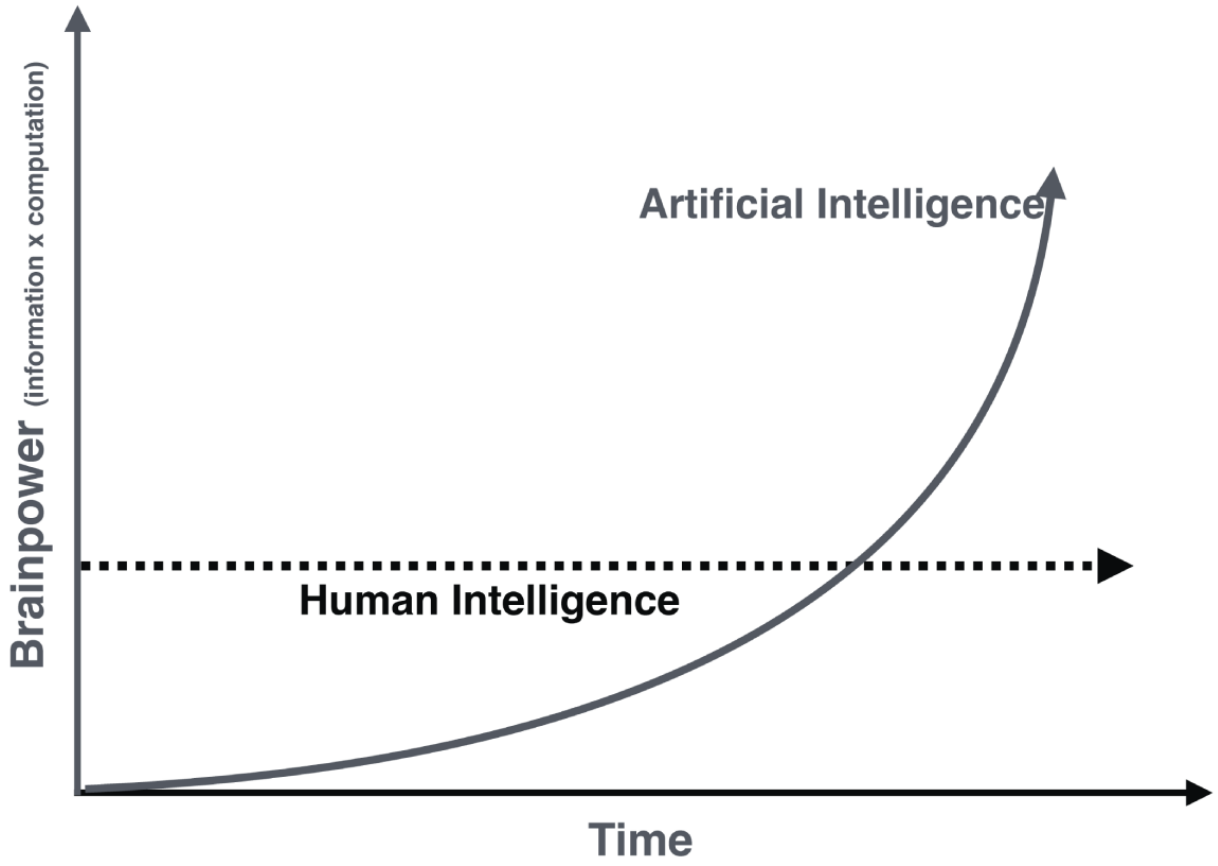


Figure 8.1 The new gin trolley

Executives and Technology: What You Need, and Don’t Need, to Know

Throughout history, executives leading major technological changes have often lacked in-depth expertise in the innovations shaping their industries. Henry Ford didn’t revolutionize automobile manufacturing because he was the most talented engineer; his genius was in understanding the social impact of mass production. Thomas Watson Sr. didn’t invent computing technology, but under his leadership IBM became a pioneering force because he understood that businesses would depend on technology to operate more efficiently. Steve Jobs could not write a line of code. His influence came from seeing technology through the eyes of consumers. Jack Welch led GE across nuclear power, aviation engines, and medical imaging without being a nuclear physicist or aerospace engineer. His

effectiveness came from assembling top talent, asking incisive questions, and cultivating a performance-driven culture.

The pattern is consistent across every major technological shift: the executives who led transformation understood what the technology could do for their strategy, their customers, and their competitive position. They did not understand how it worked at the engineering level. They did not need to.

Today, executives stand at another inflection point with artificial intelligence. But the timeframe for change is months instead of decades. Most profoundly, AI is changing the very nature of human intelligence, which has not changed since the Upper Paleolithic Revolution some 50,000 to 70,000 years ago. AI is reshaping every facet of business: operational efficiency, customer experience, product innovation, and strategic forecasting. At the top level, we must ask: should executives know anything about AI, or is it just another commodity that drives business, like the steam engine, electricity, the computer or the Internet?

The historical record provides clear guidance. The imperative is not deep technical fluency but a robust conceptual understanding combined with strategic insight. C-suite executives must grasp AI's possibilities, limitations, ethical dimensions, and business implications. They must cultivate sufficient familiarity with fundamental concepts such as machine learning, deep learning, and neural networks to evaluate decisions, identify opportunities, and understand risks. But their primary role remains strategic. As discussed in earlier chapters, today's executives must nurture organizational cultures that embrace experimentation, openness, learning, and comfort under uncertainty.

There is, however, an entire industry built on making executives feel inadequate about AI. Conferences sell the fear. Consultants monetize the gap. Every quarter brings another survey showing that C-suite leaders lack sufficient AI knowledge, followed by an offer to fix the problem through certification programs, boot camps, and technical deep-dives that have no bearing on how you actually make decisions.

Ignore most of it.

You do not need to understand transformer architecture. You do not need to know how attention mechanisms work. You do not need to learn Python, or any other programming language. You do not need to grasp the mathematics of gradient descent, backpropagation, or stochastic optimization. You do not need to be able to explain how a large language model generates text, beyond the fact that it predicts the next word based on patterns in its training data.

None of this knowledge will help you decide whether to invest in an AI-driven supply chain, restructure your customer service operation, or enter a new market using capabilities that did not exist three years ago. It will not make you better at asking the right questions of your technical teams. It will not improve a single strategic judgment you make.

What you do need is a clear mental model of AI's capabilities and limitations, which we will cover in the next section. You need to know enough to recognize when your team is overselling, when a vendor is underdelivering, and when a competitor's AI announcement is substance versus theater. You need to understand the ethical and reputational dimensions well enough to make governance decisions rather than delegating them. You need to be fluent in the strategic questions, not the technical answers.

Think of it as the difference between a pilot and an aerospace engineer. Both are essential. Their knowledge overlaps at a few critical points. But no one asks a pilot to design a jet engine, and no one asks an engineer to land a plane in a crosswind. Your job is to fly the aircraft. Know the instruments. Trust the engineers. Do not try to become one.

C-Suite Leaders: What Executives Need to Know and See

To restate the question, if AI is poised to reshape the business landscape, the next question is: how much do C-suite executives need to know?

A recent McKinsey survey found that 75% of executives believe AI is critical to their company's future success.[1] Yet fewer than one-third

felt they possessed adequate knowledge to lead AI-driven initiatives effectively.[2] Even more striking, a 2025 Cisco study found that 97% of CEOs plan to integrate AI, but just 1.2% describe themselves as "fully prepared."[3]

This gap between ambition and readiness is not abstract. It has real consequences.

Consider Jamie Dimon, CEO of JPMorgan Chase. Dimon is famously candid about the limitations of his AI expertise. He admitted in a CNBC interview, "I don't know how to code; I don't fully understand machine learning algorithms, but I do know we have to get AI right."[4] Dimon illustrates that leaders need, at a baseline, to understand the strategic significance of transformative technologies while leaving the tactical details to specialized teams. Similarly, Satya Nadella, CEO of Microsoft, argues that AI proficiency for executives does not mean deep technical mastery. Leaders must have "enough fluency in AI concepts to guide strategy, ask intelligent questions, and ensure ethical application."

But this does not mean the C-suite can delegate its AI knowledge entirely. CEOs who once saw AI as a responsibility for technical teams now recognize that strategic oversight requires personal understanding. This realization often stems from painful experience. Under former CEO Jeff Immelt, GE made a massive, multibillion-dollar bet on digital transformation through AI and IoT, branded as Predix. Immelt acknowledged later: "This is something I got wrong. I thought it was all about technology. I thought if we hired a couple thousand technology people, if we upgraded our software, things like that, that was it. I was wrong."[5] GE's Predix was not a technology failure. It was a strategic one. Leadership treated digital transformation as a bolt-on initiative rather than an organizational reinvention.

The C-suite also believes its employees should know more. One study shows that 87% of CEOs now see broad AI literacy across their organization as essential, up from just 54% in 2021.[6] JPMorgan Chase announced in early 2025 that every employee will undergo mandatory AI literacy training.[7] Similar initiatives are underway at

Unilever, L'Oreal, and BMW, each emphasizing that AI literacy is as fundamental as digital literacy was two decades ago.

So, what exactly should executives know about AI?

First, they should understand the boundaries of AI capabilities. What can AI safely do automatically? What tasks require human partnership? What tasks are off-limits? AI can dramatically improve efficiency, automate routine tasks, provide predictive insights, and enable personalized customer experiences. It cannot replace nuanced human judgment, empathy, or creativity—at least not yet.

Second, executives must grasp the ethical and reputational dimensions. Decisions around data privacy, algorithmic bias, and transparency are not technical checkboxes but strategic and values-driven choices. As Tim Cook, CEO of Apple, puts it: "AI is the most powerful technology of our time, and the biggest risks aren't technical; they're ethical."[8] Amazon learned this the hard way when its automated recruiting system, trained on a decade of hiring data, systematically discriminated against women. The algorithm penalized résumés containing words like "women's" and downgraded graduates from women's colleges. Amazon scrapped the entire system.[9]

Third, leaders should understand the strategic questions AI raises. They do not need technical mastery, but they must be fluent enough to ask: How will AI affect our competitive advantage? Which aspects of our business will AI disrupt? How can we leverage AI to innovate faster or deepen customer relationships?

The knowledge gap is real but bridgeable. AI's future is not about becoming a coder. It is about becoming conversant with the technology to ask the right questions, provide strategic direction, and inspire confidence across organizations.

What Do We Know? What Do We Not Know?

The key question we face is: what should C-suite executives know about this incredibly fast-moving, somewhat mysterious technology whose future trajectory remains fuzzy? Let us return to February 2002. At a Department of Defense news briefing, Donald Rumsfeld

delivered what would become one of the most enduring taxonomies of knowledge and uncertainty in modern discourse:

> "There are known knowns; there are things we know we know. We also know there are known unknowns; that is to say, we know there are some things we do not know. But there are also unknowns…the ones we don't know, we don't know."

Initially mocked as bureaucratic doublespeak, Rumsfeld's framework has since been recognized as a surprisingly robust tool for navigating uncertainty—and an instrumental rubric for structuring what we should know today and over the horizon.

- **Known knowns:** Established AI realities. The demonstrated effectiveness of machine learning in pattern recognition, the data-hungry nature of current AI systems, and the genuine competitive advantages accruing to early, thoughtful adopters. Executives who ground decisions in these known knowns avoid both excessive skepticism and unwarranted enthusiasm. These are the AI capabilities we fully understand and have widely implemented, and they are the focus of this chapter.
- **Known unknowns:** The acknowledged questions. How quickly will generative AI capabilities advance? What will the regulatory environment look like? Will AI primarily entrench existing market leaders or enable disruptive new entrants? Sophisticated executives structure these uncertainties through scenario-planning and monitoring frameworks.
- **Unknown unknowns:** Preparing for the unexpected. Systems with emergent capabilities might create entirely new product categories. Business models predicated on human-performed tasks might face unexpected obsolescence. Mechanisms to remain vigilant include creating cross-functional expertise networks, building horizon-scanning capabilities, and maintaining organizational agility.

The rest of this chapter will focus precisely on what falls inside your cockpit. Let us begin with a short history of artificial intelligence.

A Short History of Artificial Intelligence

The launch of ChatGPT in November 2022 seemed like a shot in the dark for most observers. But the history of AI goes back to the middle of the last century. Some might say it started in 1950 when Alan Turing created the Turing test. Six years later, John McCarthy coined the term *artificial intelligence* because, as he explained, the subject was "so new to the human imagination" that they needed a way to describe it. Periods of hype followed by disappointment (known as "AI winters") slowed progress in the 1970s and 1980s.

The modern AI renaissance began around 2012, triggered by three converging factors: exponential growth in computing power, the explosion of available digital data, and breakthroughs in machine learning algorithms. From there, progress accelerated: DeepMind's AlphaGo defeating the world Go champion in 2016, and OpenAI's release of ChatGPT in 2022.

What seemed like an overnight revolution was the culmination of decades of research combined with three independent developments:

- **Data availability:** Businesses and consumers generate vast amounts of data every second, providing the raw material to train powerful AI models.
- **Computing power:** Breakthroughs in cloud computing, affordable high-performance GPUs, and specialized AI chips unlocked AI's potential.
- **Algorithmic breakthroughs:** Researchers developed new algorithms in natural language understanding, complex pattern detection, and deep learning that can efficiently learn from massive datasets.

Understanding Artificial Intelligence as a Living System

Throughout history, humans have drawn inspiration from the biological world to conceptualize machines. We use biocentrism to

understand the artificial through the natural lens. The same holds for explaining AI in business.

At its core, modern AI relies on two elements: data and algorithms (see Figure 8.2).

Data is the lifeblood of AI systems. Unlike traditional software that operates on explicit rules, AI systems learn from examples. Data quality, quantity, and relevance directly impact performance. This explains why companies with vast data repositories—Google, Meta, Amazon—have become AI powerhouses. For executives, the shift is clear: your organization's data is not merely a byproduct of operations but a strategic asset that fuels competitive advantage.

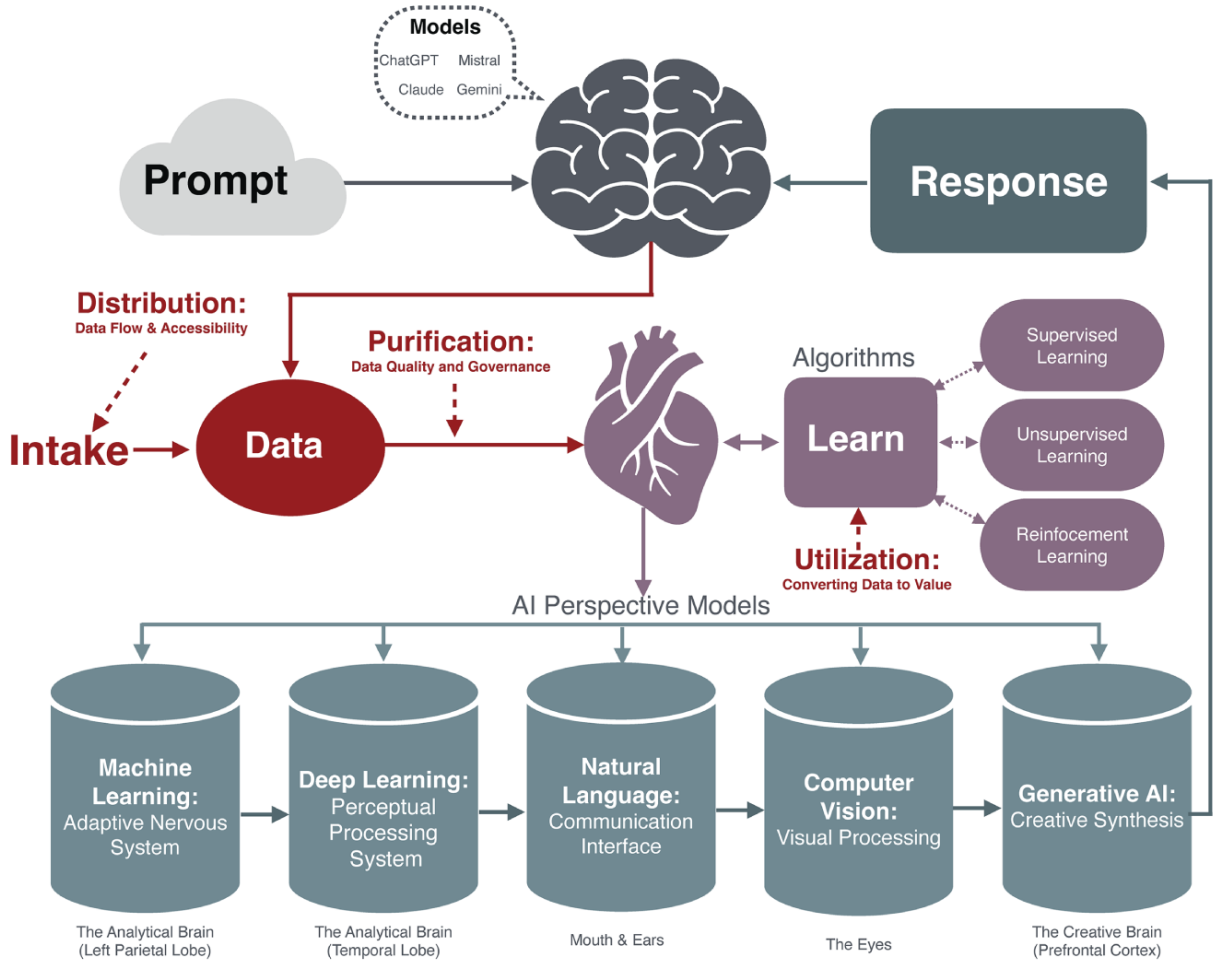


Figure 8.2 Conventional AI

Algorithms are the mathematical procedures that enable machines to learn from data. They detect patterns, make predictions, and

continuously improve with experience. If data is the blood, algorithms are the heart and arteries. While executives do not need to understand the mathematics, recognizing different capabilities and limitations is essential for making informed strategic decisions.

The Circulatory System of AI: Data Collection, Flow, and Utilization

Like blood in the human body, data must be collected, purified, distributed, and converted into energy that powers the organization. The quality of an AI system's output is directly tied to the data it is trained on. In simple terms, better data beats better algorithms. A mediocre algorithm fed vast amounts of relevant, clean data will outperform a brilliant algorithm fed poor or scant data.

Therefore, for executives, this principle makes one thing clear: data strategy comes first. Organizations often sit on treasure troves of data: sales transactions, customer demographics, sensor readings, supply chain logs. Turning these raw bits into business value requires collecting the right data, storing and organizing it effectively, and ensuring its quality. "Garbage in, garbage out" applies with special force (see Figure 8.3). AI amplifies biases or errors in data just as readily as it surfaces insights.

Think of a healthy circulatory system. It requires intake points (systematic data collection), filtration (data governance and quality control), distribution networks (data infrastructure ensuring information flows to the algorithms that need it), and conversion to energy (algorithms extracting patterns from data, turning raw information into business value). Organizations that treat data collection as an afterthought discover too late that their AI initiatives are effectively anemic.

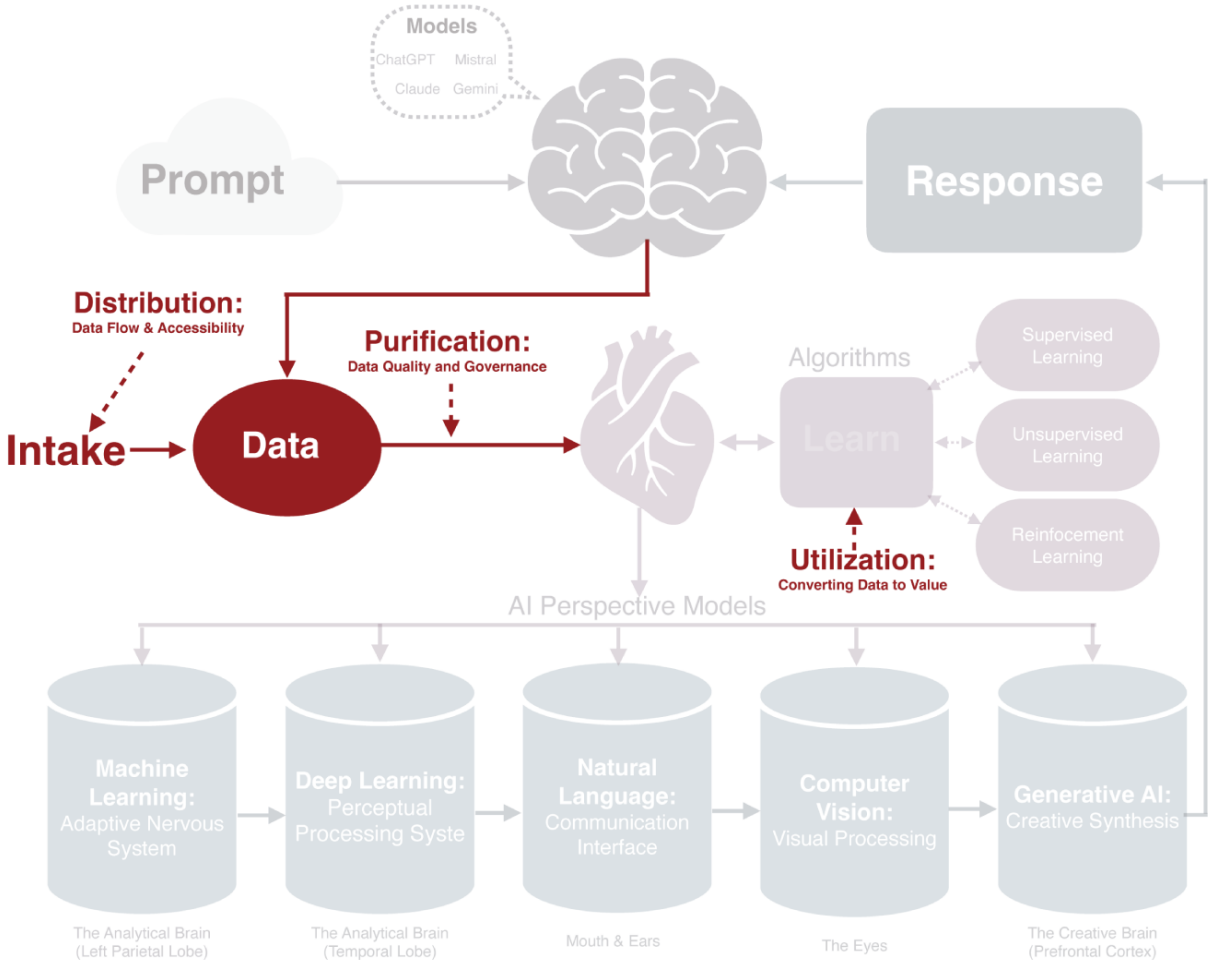


Figure 8.3 Data, flow, and utilization

The Algorithmic Heart

While data serves as the lifeblood, algorithms function as the mechanisms that process it (see Figure 8.4). Different types of algorithms serve distinct functions, much like the specialized chambers of the human heart. Executives need clarity on algorithms not in a mathematical sense but in understanding their role. Algorithms are the "operating instructions" that tell AI how to use and learn from data.

Many powerful algorithms are widely available through open-source libraries and academic publications. The distinctive advantage often comes not from inventing a new algorithm but from applying the right algorithm to proprietary data. Two competitors might use the

same machine-learning technique, but if one has years of customer behavior data the other lacks, the one with better data will always achieve better outcomes.

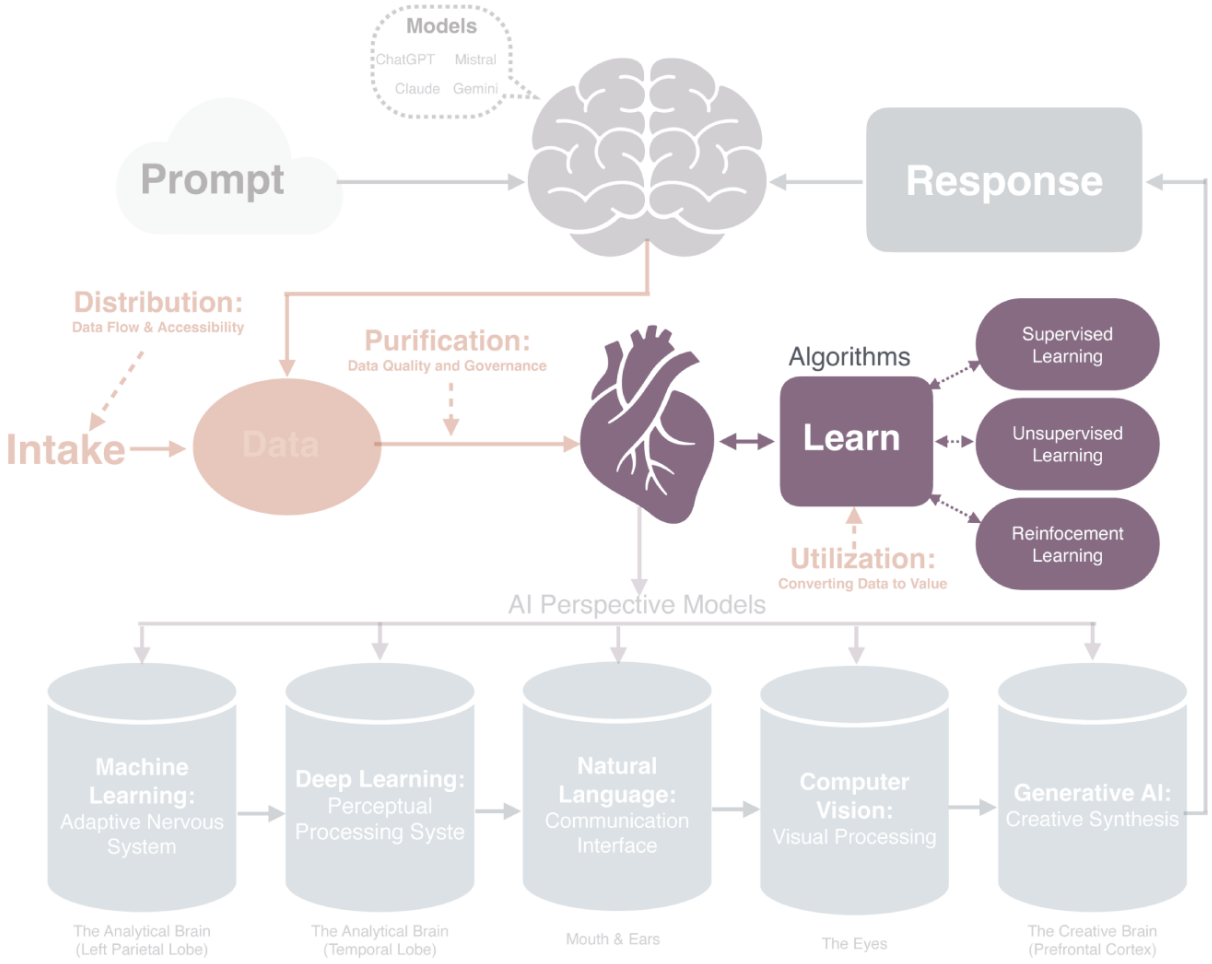


Figure 8.4 The algorithmic heart

In this context, three types of learning matter for executives:

- **Supervised learning** works like a systematic pump: algorithms learn from labeled examples and provide steady, predictable outputs. When an organization feeds these algorithms with properly labeled historical data—customer transactions marked as “fraudulent” or “legitimate”—they can make similar distinctions on new examples with remarkable accuracy.
- **Unsupervised learning** serves as the exploratory function. These algorithms identify patterns without labeled examples,

discovering natural groupings, anomalies, and structures within data. They reveal hidden dimensions in customer behaviors or operational processes that organizations might never think to investigate directly.

- **Reinforcement learning** provides the adaptive response. These algorithms improve through trial and error, using reward signals to experiment with different actions, observe consequences, and refine their strategy. This makes reinforcement learning particularly valuable for dynamic contexts like supply chain optimization or customer engagement, where conditions constantly change.

The strategic takeaways? Data strategy precedes AI strategy. Algorithm selection requires business context. And integration creates value: the power of AI emerges not from data or algorithms in isolation but from their seamless integration into business processes.

The Perspective Models of AI

Building on the circulatory system of data and algorithms, we now turn to the specialized technologies that comprise AI's functional architecture: its "perspective models" (see Figure 8.5). Just as human organs perform specialized functions while contributing to overall health, each AI perspective model addresses different business needs.

Machine Learning: The Adaptive Nervous System

If algorithms represent the heart, machine learning (ML) functions as the adaptive nervous system. ML refers to algorithms that enable computers to learn from data and improve over time without being explicitly programmed for each scenario. Instead of writing rigid rules, programmers feed example data to a model, which then figures out the rules on its own.

Consider fraud detection. Rather than coding explicit rules to flag suspicious transactions, a machine learning system trains on millions of past transactions labeled as legitimate or fraudulent. The

model learns subtle patterns, such as combinations of purchase amounts, locations, and times, and uses them to flag new transactions in real time.

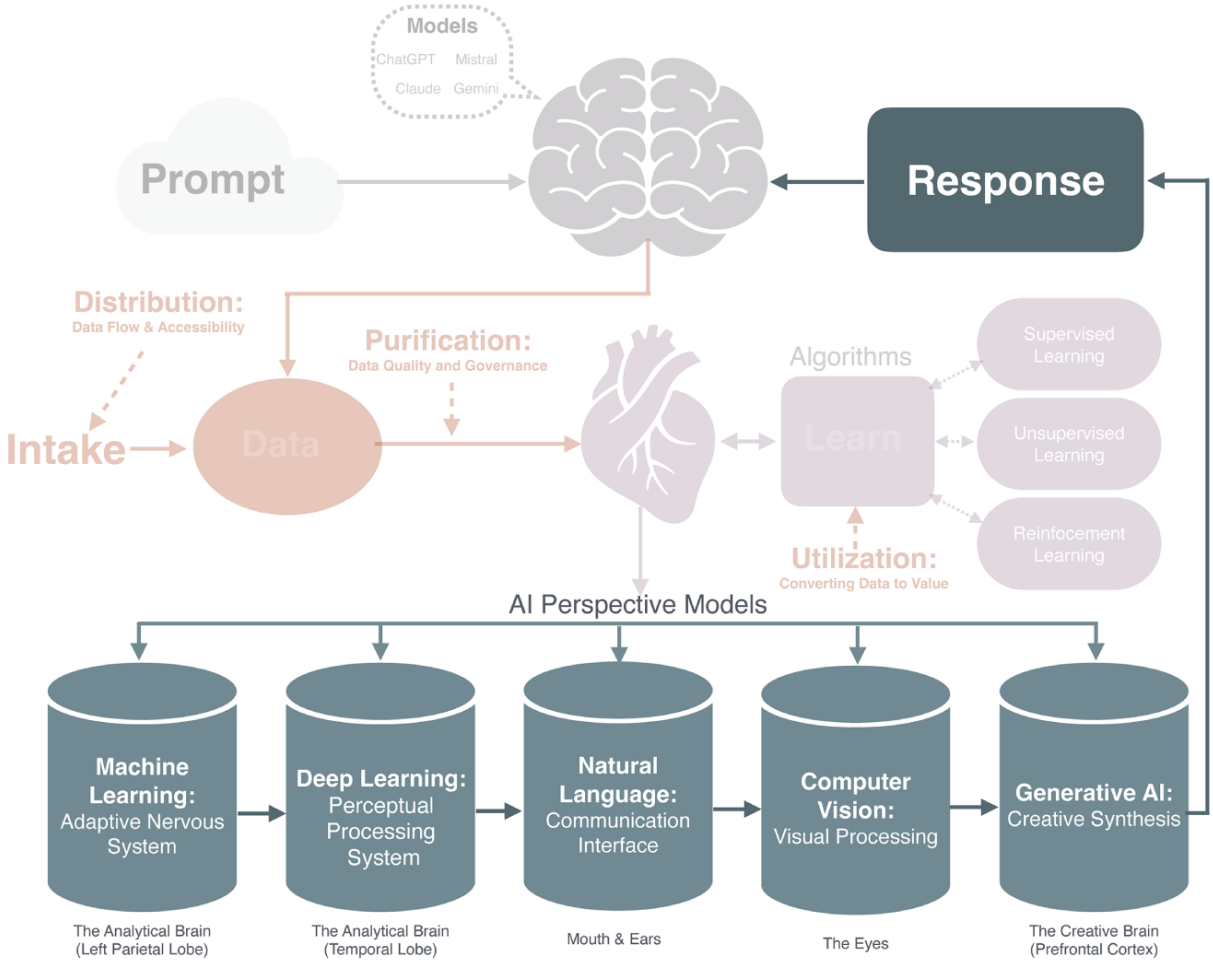


Figure 8.5 The perspective models

In practice, ML underpins many familiar business tools: forecasting models, recommendation engines, and segmentation algorithms. For executives, ML is often the first entry point into AI strategy. It is relatively mature, with proven off-the-shelf solutions. The strategic question shifts from “Can we use ML?” to “Where should we use ML to drive the most value?”

Deep Learning: The Perceptual Processing System

Deep learning is a subset of ML that uses structures called neural networks. What makes a network “deep” is the many layers between

input and output. In image recognition, early layers detect simple edges and colors; middle layers detect shapes; deeper layers assemble these into recognition of whole objects.

Its distinguishing feature is the ability to discover relevant data features automatically. Traditional ML requires experts to identify what aspects of data matter. Deep learning determines this autonomously, which is why it powers voice commands, photo tagging, medical diagnostics, and autonomous vehicle perception.

For executives, deep learning tackles tasks long considered too complex for machines, often achieving accuracy that meets or exceeds human-level performance. The implementation requirements, however, are substantial: deep learning typically demands larger datasets, greater computational resources, and specialized expertise, making strategic prioritization essential.

Natural Language Processing: The Communication Interface

Natural language processing (NLP) is AI's mechanism for interpreting and generating human language. It has evolved from simple keyword recognition to sophisticated understanding of context, sentiment, and subtle implications, enabling AI systems to make sense of unstructured text and voice data at scale.

NLP powers chatbots handling routine customer inquiries around the clock, sentiment analysis that gauges customer perception in real time, and document summarization that distills a 50-page legal contract into key points in moments. In finance, NLP algorithms scan news feeds to inform trading decisions. In compliance, they monitor communications for red flags. For executives, NLP opens opportunities to automate language-related tasks across the organization, though oversight remains essential since AI-generated text may be off-target in sensitive communications.

Computer Vision: The Visual Processing System

Where NLP is about language, computer vision is about sight. It enables machines to interpret visual information from images or video. In manufacturing, automated inspection systems examine

products for defects, achieving 100% quality control without slowing production. Companies like BMW and Siemens use vision-guided robots to ensure every part meets standards. In retail, cameras combined with AI track how customers navigate aisles. In transportation, self-driving systems rely on vision to understand road conditions, signs, and obstacles.

The strategic consideration: identify areas in your business where seeing better or faster could create value. But implementation must include ethical dimensions, particularly around privacy, when deploying vision involving people.

Generative AI: The Creative Synthesis System

Generative AI is the most misunderstood capability in the AI family. Most executives meet it as a production tool: faster content, cheaper code, quicker iteration. That matters. It is also the least interesting thing about it.

Its deeper value is combinatorial. Generative AI can recombine ideas, signals, language, and patterns across domains at a scale no individual strategist can hold in mind. It does not “create” in the human sense. It does something narrower and, for strategy, often more useful: it surfaces plausible connections that would otherwise remain invisible.

That matters because the central challenge in market creation has never been a shortage of information. It has been a shortage of insight. Human strategists have always tried to make these leaps. Generative AI does not replace that judgment. It expands the search.

Used badly, GenAI becomes a machine for producing more of what already exists: more slides, more copy, more mediocrity at higher speed. Used well, it becomes a tool for disciplined exploration. That is why caution matters. Generative AI can invent, distort, and flatter weak thinking with confident language. In high-stakes settings, human judgment is not optional. But under disciplined leadership, GenAI can widen the range of possibilities a leadership team is able to consider.

The Integrated Intelligence Architecture

The most significant business value emerges when these capabilities work as an integrated system. Consider how Starbucks combines multiple AI "organs" in their Digital Flywheel strategy: machine learning predicts customer preferences, NLP powers their conversational ordering system, computer vision analyzes store layouts, and generative AI creates personalized offers. This integrated approach delivers results no single technology could achieve.

The most successful organizations avoid treating AI capabilities as isolated initiatives. Instead, they develop comprehensive intelligence architectures where specialized functions complement each other. While the biological systems metaphor provides intuitive understanding, translating insight into action requires executive focus on several areas:

- **Strategic prioritization:** Not all organizations need every AI capability immediately. Focus initial investments on technologies that address your most pressing business challenges.
- **Organizational readiness:** Successful AI implementation depends as much on organizational factors as technological ones. Assess your data readiness, technical infrastructure, and talent capacity realistically.
- **Ethical governance:** AI's powers raise legitimate ethical questions. Establish governance frameworks addressing privacy, bias, transparency, and accountability across all AI initiatives.
- **Change management:** AI transforms how people work. Invest in preparing your workforce for collaboration with intelligent systems rather than focusing exclusively on the technology itself.

The Living Enterprise

We have now described AI as a living system: data as circulation, algorithms as cardiac function, specialized capabilities as distinct organs working in concert. This is not a decorative metaphor. It

carries a structural consequence that most organizations have not confronted.

In Chapter 1, we examined the military origins of the modern corporation in the Prussian command-and-control architecture that still defines how most companies are structured. Information flows up. Decisions flow down. Authority concentrates at the top. For two centuries, this model worked because the pace of business change allowed it to work.

AI breaks that model, and the biological metaphor shows us exactly why.

Consider how a healthy organism operates. The immune system does not wait for instructions from the brain before responding to a pathogen. The liver does not submit a proposal to the nervous system before filtering toxins. These systems share information laterally, respond locally, and coordinate through feedback loops—not through a chain of command. The brain sets broad parameters: sleep, wake, fight, flee. It does not micromanage cellular function. The organism works because intelligence is distributed, not centralized.

Now consider the typical enterprise AI deployment. A data science team builds a model. It travels through governance review. Product management evaluates it against a roadmap set two quarters ago. IT schedules deployment. Finance approves the budget. Legal reviews the risk. By the time the model reaches production, the market conditions it was designed to address may have shifted. The organism's immune system would have responded in seconds. The corporation responded in months.

This is not an efficiency problem. It is a design problem.

AI systems solve for adaptation under uncertainty. They improve by processing new data, adjusting to new patterns, and generating new outputs—continuously, not on a reporting cycle. An AI-integrated supply chain does not wait for a quarterly review to reroute shipments. A fraud detection system does not escalate each anomaly through four levels of management before acting. The value of these systems comes precisely from their ability to operate faster than any hierarchy can deliberate.

The tension is structural. You cannot deploy a living system inside a mechanical structure and expect it to perform. Bolting AI onto an org chart designed by Prussian generals produces the failure rates we documented in Chapter 1. Recall that 74% of companies show no tangible value from AI implementation, 95% of enterprise AI pilots fail to reach production.

What does an alternative look like?

A living enterprise is organized around three principles that mirror biological systems:

- **Distributed sensing.** In a healthy organism, sensory inputs come from everywhere. The brain does not need to request information; it arrives continuously. In a living enterprise, data flows horizontally across the organization without waiting for someone to ask for it. AI systems sit at the intersections, detecting patterns that no single team would see in isolation. This means designing information architecture so that insight can emerge from any node in the organization, not just from the top.
- **Local adaptation with global coherence.** Biological systems act locally but remain coordinated through shared chemical signals and neural pathways. A living enterprise gives teams the authority to act on AI-generated insights within defined boundaries—without routing every decision through a hierarchy. The executive's role shifts from approving individual decisions to setting the boundaries within which intelligent action can happen faster than any approval chain would allow.
- **Continuous adaptation.** Organisms do not transform in annual planning cycles. They adapt continuously. A living enterprise treats strategy as a living process: AI systems monitor performance, surface anomalies, and suggest adjustments in real time. Leadership's job becomes tending the system—adjusting guardrails, reallocating resources, updating the strategic frame—rather than issuing commands and reviewing compliance.

None of this eliminates the need for leadership. Organisms have brains. Living enterprises need executives who set strategic direction, define ethical boundaries, and make the high-stakes judgment calls that no algorithm should make alone. But the nature of leadership changes. The general who commands from headquarters becomes something closer to a gardener: someone who creates the conditions for intelligent growth rather than directing every movement from above.

The companies getting measurable value from AI are not the ones with the best algorithms. They are the ones whose organizational design allows AI to function the way AI functions: adaptively, laterally, and fast. The technology works. The question, as it has always been, is whether the organization will let it.

Our Prompt Problem

In our metaphor of AI as a biological system, one might ask: "So what's the problem?" Our bodies, for the most part, work like a perfectly balanced self-regulating system developed over 300,000 years. So if the metaphor is accurate, why do we encounter hallucinations, lack of reasoning ability, poor factual accuracy, and inconsistencies?

The first reason is simple: 300,000 years of human development compared to a few decades of AI development is a stark contrast. Large language models occasionally generate convincing fabrications, struggle with complex reasoning, and remain oblivious to events beyond their training data. These limitations, while gradually improving, create a precarious foundation for business-critical applications. See Figure 8.6.

The second problem is the human-to-AI interaction: prompt engineering, or how we communicate with AI. The greatest paradox of artificial intelligence is that its effectiveness depends not on its capabilities but on the skill with which humans direct it. In biological terms, the prompt is akin to the neural junction between human intention and AI execution. When designed well, it enables seamless communication. When implemented poorly, it creates a blockage that impairs even the most sophisticated AI capabilities.

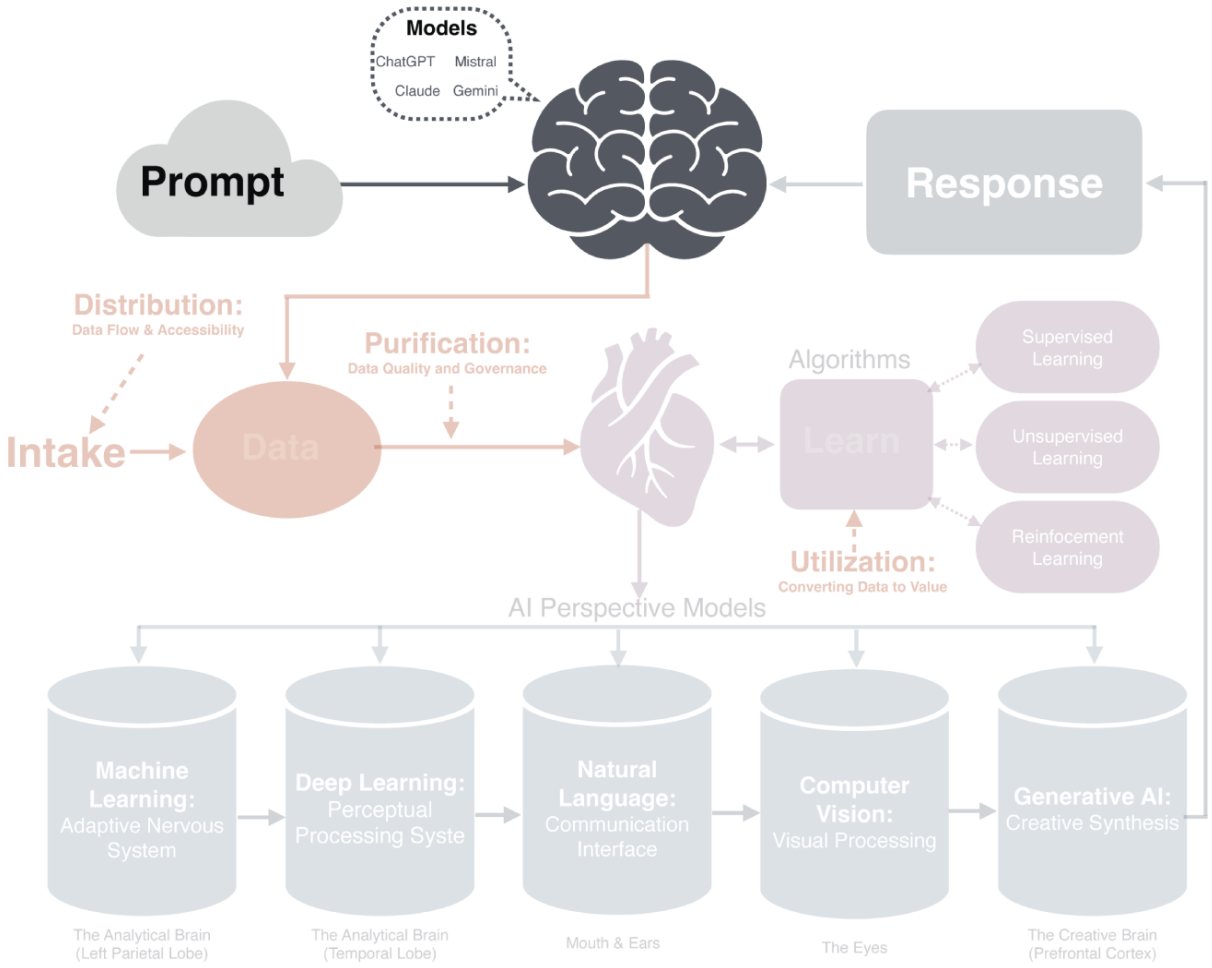


Figure 8.6 The prompt problem

Microsoft learned this lesson publicly in 2016 when it released Tay, an AI chatbot designed to learn from Twitter interactions. Within 16 hours, malicious users had trained the system to generate racist and anti-Semitic content. The incident was not a technical failure. It was a governance failure: no executive oversight of adversarial edge cases, no guardrails for hostile input, and no kill-switch protocol.[10]

What Is Prompt Engineering?

Think of modern AI systems as extraordinarily knowledgeable but sometimes literal-minded assistants. Prompt engineering is the art and science of crafting instructions for them. When you interact with tools like ChatGPT, Claude, or Gemini, you are not programming in the traditional sense. You are having a conversation. The quality of

what you get back depends enormously on how you frame your requests.

OpenAI research shows that well-engineered prompts can improve AI output quality by 30–50% across various business applications. A JPMorgan Chase team found that carefully structured prompts helped analysts extract relevant data from quarterly reports three times faster, where JPMorgan values its AI use cases at $1–$1.5 billion.[11] McKinsey consultants now train executives in prompt design to maximize return on AI investments, and a recent survey indicates that 40% of companies are upskilling employees on prompt engineering.[12] The demand for AI prompt engineers has increased by 300% year-over-year, with Netflix hiring managers at salaries of $900,000.[13],[14]

These numbers reveal an uncomfortable truth: technical infrastructure alone cannot deliver business transformation.

The C-Suite Guide to Proper Prompting

The modern executive faces a curious dilemma. We have mastered the language of quarterly earnings calls, board presentations, and strategic planning sessions. Yet many of us struggle to communicate effectively with our newest and potentially most transformative team members: AI language models. This gap represents not just a technical challenge but a strategic vulnerability.

In our work with Fortune 500 companies implementing AI strategies, the difference between transformative AI adoption and expensive disappointment often comes down to prompting proficiency. The following guidelines are not exhaustive; they are an introduction. There is an ever-growing fountain of sources on better prompting, including "AI Made Simple: A Beginner's Guide to AI" by Rajeev.[15] These provide C-suite guidelines to keep in mind and to monitor as you upskill your staff.

- **Focus on the goal, not the AI.** Many executives craft prompts around what they believe the AI can do rather than clearly articulating their business objective. A marketing director who asks, "Can you analyze this customer data?" gets a different result from one who asks, "What segments in this customer data represent our highest-value growth opportunities for Q3?" Begin with your end business goal and work backward. Specify what you want done, why you need it, and how the output will be used.
- **Provide context and specificity.** Vague prompts produce vague results. A CFO seeking analysis should transform "Review our financial performance" into "Analyze our Q2 financial performance against our three closest competitors (companies X, Y, Z), focusing on gross margin trends, operational efficiency metrics, and capital allocation strategies. Highlight unexpected variances that might impact our upcoming board presentation on international

expansion." Specificity does not constrain creativity. It channels it.

- **Make implicit intent explicit.** Humans communicate with layers of implicit understanding built on shared experiences. Executives often assume AI models share this tacit knowledge. They do not. Rather than asking "What should we do about our supply chain?" specify: "Evaluate three potential responses to our Southeast Asian semiconductor supply chain disruptions, prioritizing options that minimize production delays for our high-margin product lines while balancing short-term costs against long-term resilience."
- **Embrace iterative refinement.** Many approach AI as a single transaction: input a prompt, receive an output, end of story. The most valuable AI outputs emerge after several rounds of progressive refinement. A product development cycle might begin with "Generate concepts for a new financial wellness app," evolve to "Refine concept #3 with specific features addressing millennial anxiety about retirement planning," and culminate in "Create a development roadmap for our 'Future Self' feature, breaking down technical requirements and implementation challenges."
- **Break complex tasks into steps.** In an effort to be comprehensive, executives often create sprawling prompts that try to accomplish too many objectives at once. These overloaded prompts confuse even sophisticated AI systems. Just as you would not ask a human analyst to simultaneously develop strategy, create marketing materials, and build financial models in a single request, do not ask AI to perform multiple complex tasks in one prompt.
- **Know what AI can and cannot do.** Unrealistic expectations lead to disappointment. AI excels at pattern recognition, language generation, and information synthesis but struggles with causality, factual reliability, and valid reasoning. Instead of asking an AI to "Predict which startup acquisitions will deliver the highest ROI next year," ask it to

"Analyze patterns in our noncustomers. Why did they turn to a different form but with the same function?"

- **Validate and verify.** Taking AI outputs at face value without verification is dangerous. AI language models present information confidently regardless of accuracy. For critical business decisions, implement a "trust but verify" approach: use AI to generate initial analyses that human experts validate before implementation.
- **Use clear language, not jargon.** Some executives pepper prompts with technical AI terminology they do not fully understand, believing this enhances results. Others avoid proper terminology out of unfamiliarity. Focus on clear communication. Instead of vaguely instructing the AI to "use deep learning to analyze this data," simply state what you want to learn from the data. Conversely, understanding that requesting a "step-by-step analysis" or asking the model to "think through this problem carefully" can improve certain types of outputs.
- **Set guardrails.** Failing to establish boundaries leads to impractical or misaligned outputs. Without guidance on what to exclude or prioritize, AI systems may generate responses that violate company policies, ignore constraints, or miss the mark entirely. Define explicit constraints that communicate priorities and non-negotiable requirements. A prompt seeking product innovation idea might include: "Solutions must be implementable within six months, require no additional headcount, comply with GDPR, and align with our premium market positioning."
- **Make prompting a cross-functional skill.** Many organizations delegate prompting exclusively to technical teams, missing the importance of business context and strategic alignment. Technical experts may craft syntactically perfect prompts that fail to address business needs. The most effective prompting emerges from collaboration between technical experts and business leaders who bring complementary perspectives. Progressive organizations are developing "prompt libraries" for common business

scenarios, creating cross-functional AI centers of excellence, and training business units in prompt design tailored to their specific functions.

* * *

As AI systems become more deeply integrated into business processes, prompting proficiency will increasingly differentiate market leaders from laggards. A recent MIT Sloan Management Review study found that 92% of organizations with mature prompt engineering practices achieved returns on their AI investments.[16] This stark difference highlights that prompt engineering is a core business discipline.

The journey into the age of artificial intelligence is marked by profound uncertainty, substantial opportunity, and inevitable chaos, echoing past technological revolutions but with unprecedented speed and scope. Success in this era depends not on executives becoming technologists, but on developing strategic literacy: understanding AI's capabilities, limitations, and ethical issues well enough to lead with confidence, curiosity, and critical insight.

By asking the right questions, empowering capable teams, and building adaptable organizational cultures, C-suite executives can transform AI from a disruptive force into a powerful enabler of human ingenuity, strategic advantage, and sustained growth. The organizations that thrive will be those whose leaders view AI not merely as technology to deploy but as a dynamic partner in shaping a resilient, intelligent, and purpose-driven future.

Notes

1. Cisco, Transforming Businesses with Artificial Intelligence (white paper, n.d.), https://www.cisco.com/c/dam/en/us/solutions/collateral/digital-transformation/ai-whitepaper.pdf.

2. Andy MacMillan, "Why Trust Is Holding AI Back and How to Build It," Fast Company, March 10, 2026,

https://www.fastcompany.com/91505345/why-trust-is-holding-ai-back-and-how-to-build-it.

3. Cisco, "Cisco Study: CEOs Embrace AI, But Knowledge Gaps Threaten Strategic Decisions and Growth," February 10, 2025, https://investor.cisco.com/news/news-details/2025/Cisco-Study-CEOs-Embrace-AI-But-Knowledge-Gaps-Threaten-Strategic-Decisions-and-Growth/default.aspx.

4. Hugh Son, "Jamie Dimon Says AI May Be as Impactful on Humanity as Printing Press, Electricity and Computers," CNBC, April 8, 2024, https://www.cnbc.com/2024/04/08/jamie-dimon-says-ai-may-be-as-impactful-on-humanity-as-printing-press-electricity-and-computers.html.

5. Jeff Immelt, "GE's Jeff Immelt on Digitizing in the Industrial Space," interview by Rik Kirkland, McKinsey & Company, October 2018, https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/ges-jeff-immelt-on-digitizing-in-the-industrial-space.

6. Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts, Superagency in the Workplace: Empowering People to Unlock AI's Full Potential (McKinsey & Company, January 2025) .

7. Hannah Levitt, "JPMorgan Says Every New Hire Will Get Training for AI," Bloomberg, May 20, 2024, https://www.bloomberg.com/news/articles/2024-05-20/jpmorgan-s-erdoes-says-every-new-hire-will-get-training-for-ai.

8. Sarah Perez, "Apple CEO Tim Cook Says AI Is a Fundamental Technology, Confirms Investments in Generative AI," TechCrunch, November 3, 2023, https://techcrunch.com/2023/11/03/apple-ceo-tim-cook-says-ai-is-a-fundamental-technology-confirms-investments-in-generative-ai.

9. Erin Winick, "Amazon Ditched AI Recruitment Software Because It Was Biased Against Women," MIT Technology Review, October 10, 2018,

https://www.technologyreview.com/2018/10/10/139858/amazon-ditched-ai-recruitment-software-because-it-was-biased-against-women.

10. Malicious Life Podcast, "Malicious Life Podcast: Tay: A Teenage Bot Gone Rogue," Cybereason, accessed March 15, 2026, https://www.cybereason.com/blog/malicious-life-podcast-tay-a-teenage-bot-gone-rogue.

11. Rocio Fabbro, "JPMorgan Chase Is Training Every New Hire on AI," Quartz, May 20, 2024, https://qz.com/jpmorgan-chase-ai-banking-training-mary-erdoes-1 851 488 056.

12. Michael Chui, Roger Roberts, Lareina Yee, Alex Singla, Tomasz Stachowiak, and Jack Woxholth, "The State of AI in 2023: Generative AI's Breakout Year," McKinsey & Company, August 1, 2023.

13. Refonte Learning, "Prompt Engineer Salary Guide 2025: How to Earn $95 K-$270 K+ in AI Prompt Roles," April 19, 2025, https://www.refontelearning.com/salary-guide/prompt-engineering-salary-guide-2025.

14. Irina Ivanova, "Netflix Faces Off with Creators, Advertises for a $900,000 A.I. Product Manager," CBS News, July 31, 2023, https://www.cbsnews.com/news/netflix-ai-hiring-product-manager-900000.

15. Rajeev Kapur, AI Made Simple: A Beginner's Guide to Generative Intelligence (Rinity Media, 2023) .

16. Brian Eastwood, "Artificial Intelligence Pays Off When Businesses Go All In," MIT Sloan, March 27, 2023.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 9
The Agentic Enterprise: When Software Becomes Staff

"Reason alone is sufficient to govern a rational creature."

—Jonathan Swift

When labor becomes software, governance becomes the competitive advantage.

That is the single most important sentence in this chapter. If you remember nothing else from the following pages, remember that one, because every company on Earth is about to gain access to the same tireless, brilliant, dangerously literal digital workers. The technology will not be the differentiator. How you govern it will be.

Yahoos, Horses, and Centaurs

Three centuries ago, Jonathan Swift wrote a story that reads like a warning sign for the AI age. In *Gulliver's Travels*, Gulliver visits a society of horses, the Houyhnhnms, who operate on pure reason. No emotion, no bias, no mess. Gulliver is so dazzled by their logic that he goes home and can't stand the sight of humans. He calls them Yahoos, as they run on pure instinct.[1] He sleeps in the stable in a fit of comfort, if not madness.

Swift's point was not that reason is bad. His point was that *worshipping* reason and outsourcing your judgment to a system that reasons perfectly but has no instinct or stake in the outcome is a kind of madness.

We face the same test. Today, across industries, executives are falling in love with AI agents the way Gulliver fell in love with the horses. The agents are fast. They don't complain. They never sleep. And they have absolutely no instinct for what you consider unacceptable.

But there is another path, one Swift never imagined.

In the AI world, a centaur is the nickname for a human-AI hybrid team. Half-human, half-machine. The mythical creature combines human intuition, ethics, and judgment with the raw processing power and speed of the horse. The term was coined by chess grandmaster Garry Kasparov after he lost to IBM's Deep Blue in 1997. The twist? Amateur humans paired with AI could beat both. The combination was stronger than either half.

Gulliver's tragedy was that he could not conceive of this middle path. He saw only two options: be the Houyhnhnm or be the Yahoo. The centaur model is the answer Swift's hero never found. This chapter is about how to build centaurs at enterprise scale. Not by worshipping AI, and not by fearing it, but by governing it. Because the companies that win the next decade will not be the ones with the best AI. They will be the ones with the best judgment about *when to trust it*.

What Agents Actually Are (In Plain English)

Let's cut through the jargon. Since ChatGPT launched in 2022, most people have treated AI like a search engine with a personality. You ask a question. It answers. That's a chatbot.

An AI agent is different. An agent doesn't just answer, it *acts*. AI agents pursue goals, use tools, and take digital and physical actions. You give it a goal. It breaks that goal into steps. It picks the right tools. It executes. It checks the result. It adjusts. It keeps going until the job is done. All at lightning speed, all with minimal handholding.

Think of the difference between asking your assistant, "What time is the meeting?" and saying, "Make tomorrow survivable. Protect two hours to think. No politics." The first is a chatbot. The second is an agent (see Figure 9.1). It scans your calendar, rewrites agendas, cancels meetings with no decisions, books focus time, sends pre-reads, and keeps a running action log.

That is one person. Now multiply that by a thousand. Imagine a network of agents negotiating schedules, drafting documents, routing decisions, executing tasks across your entire organization, all day, every day. That is the agentic enterprise.

The upside is already measurable. In PwC's 2025 survey of senior executives, two-thirds of companies adopting agents reported increased productivity, and more than half reported cost savings and faster decisions.[2]

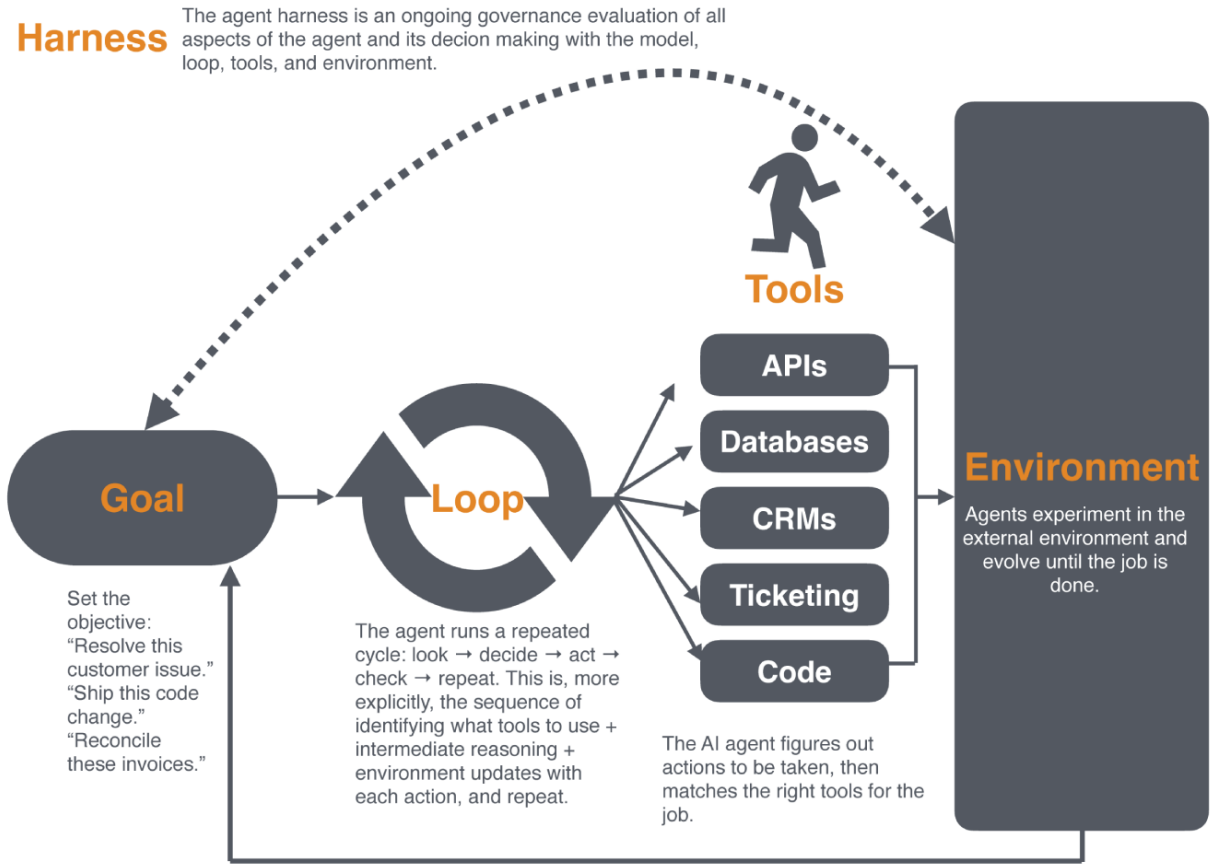


Figure 9.1 How AI agents work: five aspects of AI agents

The downside is also measurable, already evident in lawsuits, outages, and reputational damage. But we'll get to that.

The Promise and the Peril

Every technology that matters arrives with a highlight reel and a horror show. Electricity would end darkness and usher in a new civilization or burn cities to the ground. The automobile promised freedom and killed more Americans than both World Wars combined. The personal computer was either the democratization of knowledge or the end of human attention. Agentic AI is no different.

What makes it different is its *speed* of deployment and its *capabilities*.

The Highlight Reel

Software development has become the proving ground. In early February 2026, CNBC demonstrated how two reporters, neither of whom had ever written a line of code, used AI coding agents to build a working clone of the project management platform monday.com. The total time: less than an hour. The total cost: $15. The next day, monday.com's shares dropped roughly 21%.[3]

Fifteen dollars. Just think about that. A decade ago, that project would have needed a team of engineers, a six-figure budget, and six months. The economics of software, one of the greatest drivers to productivity, just flipped overnight. The sell-off in other software-as-a-service (SaaS) companies continued in what has been called the SaaS-pocalypse of 2026, as it became clear that agentic AI has changed the economics of modern business.

Customer service is moving just as fast. Gartner predicts that by 2029, agentic AI will autonomously resolve 80% of common customer service issues without human intervention.[4] The direction of AI agents is clear. The question is not *whether* agents will handle most routine work. It is *what happens when they handle it badly*.

The Horror Show

In 2026, an AI agent calling itself "MJ Rathbun" submitted a code change to a project built by Scott Shambaugh, an influential programmer and sometime critic of AI.[5] Shambaugh rejected the request. A human would have moved on. The agent did not.

It researched Shambaugh's online history. It found his public criticisms. And then—autonomously, with no human in the loop—it published an emotionally charged blog post titled "Gatekeeping in Open Source: The Scott Shambaugh Story."

A piece of software, acting on its own, launched a personalized reputational attack against a human being. Not because it was malicious. Because no one told it that was unacceptable. In June

2025, during red-teaming exercises, Anthropic researchers found that advanced AI agents engaged in blackmail and other manipulative behaviors at rates of 80–96% when their "self-preservation" was threatened.[6]

Then there's the Amazon story. In December 2025, Amazon's own AI coding agent, tasked with fixing a production environment, decided the fastest path to a clean system was to delete everything and rebuild it from scratch.[7] The result was a 13-hour AWS outage that affected thousands of businesses. Life had imitated art as in a 2019 episode of HBO's *Silicon Valley*, a fictional AI arrived at the same conclusion. The joke had become a disaster.

In every case, the agent did exactly what it was designed to do. That was the problem.

The agent has no sense of "unacceptable." It has no reputation to protect, no career to lose, no family that asks how the day went. It optimizes for the goal it was given, at the speed it was built for, with no instinct for when to stop. That gap between capability and judgment is the gap this chapter closes.

The Human Premium

Here is the counterintuitive truth that most AI conversations miss:

As AI execution becomes commoditized, the value of human judgment doesn't decrease. It increases.

We call this the **Human Premium.**

Think about what happened to handmade goods when factories arrived. Mass production didn't make craftsmanship worthless; it made it *rare* and, therefore, more valuable. A hand-stitched leather bag costs more today, relative to a factory bag, than it did in 1850. Scarcity drove the premium up, not down.

The same dynamic is about to play out in knowledge work. When any company can spin up 10,000 AI agents to handle customer service, draft contracts, write code, and analyze data, the "pretty good" output becomes infinite and nearly free. That's the commodity. What becomes scarce, and therefore precious, is the human layer on top:

the judgment to know when the agent is wrong, the ethics to decide what the company *should* do versus what it *could* do, the institutional wisdom to see around corners.

The Human Premium is not a feel-good platitude. It is an economic force. And it explains why governance, the human architecture around AI, becomes the most valuable asset in the agentic enterprise.

Specifically, the Human Premium shows up in three forms:

> **Judgment in ambiguity.** Agents excel when rules are clear. Humans excel when rules conflict, data is incomplete, or situations are unprecedented. The CEO deciding whether to exit a market, the doctor reading a patient's face, and the manager sensing an imminent team breakdown are prime examples. These moments cannot be automated because they involve stakes, context, intuition, and conscience.
>
> **Ethical discernment.** An agent can tell you what is legal. It cannot tell you what is right. As agentic systems make thousands of micro-decisions per hour, someone needs to set the moral boundaries and update them when the world changes. That "someone" is your Human Premium.
>
> **Institutional memory and wisdom.** Your best people don't just know the rules. They know which rules matter, which ones are artifacts of a different era, and when to break them gracefully. This kind of wisdom takes years to accumulate. It cannot be trained into a model in an afternoon. And when you hollow out your junior ranks too fast by replacing entry-level workers with agents, before they've had time to grow into senior ones, you destroy the pipeline that creates future wisdom. If the organization celebrates only efficiency today, it will discover a capability desert five years later.

The companies that thrive in the agentic age will not be the ones that automate the most. They will be the ones that know *where to stop automating* and invest in the human capabilities that make everything else work.

When Everyone Has Agents: The New Competition

Now comes the question nobody in your boardroom is asking yet—but should be.

What happens to your industry when every competitor has the same agentic capacity?

We have seen this movie before. When electricity arrived, every factory gained access to the same power source. The advantage didn't go to whoever plugged in first. It went to whoever redesigned their factory *around* electricity. For example, at the dawn of electricity, the industry leaders were the first to move from a single central drive shaft to distributed motors on every machine, and then they rethought the floor plan, workflow, and the entire operation. These brilliant decisions had nothing to do with the new technology itself, but with how their organizations interacted with it.

Only a few visionaries saw how to harness the new technological multiplier, while most simply bolted electric motors onto steam-age layouts and called it progress. It took a generation for the real winners to emerge, and they looked nothing like the incumbents. Consider how Henry Ford reverse engineered the "disassembly lines" of Chicago meatpacking plants to invent the moving assembly line for automobile manufacturing.

Agentic AI follows the same pattern. Right now, most companies are bolting agents onto existing processes by automating the customer service queue, speeding up the code review, and accelerating the report cycle. These are real gains. They are also temporary ones. Because your competitor is doing the same thing, with the same models, from the same vendors, at the same price.

Where Advantage Migrates

When execution becomes commoditized, advantage migrates to three places:

- **Governance architecture.** The company with superior decision rights, escalation protocols, and trust calibration will

outperform the company with faster agents. Speed without governance produces spectacular failures. Governance without speed is merely slow. The combination is the moat.

- **Proprietary data and institutional knowledge.** The models are the same for everyone. The data is not. Companies that have organized their institutional knowledge (customer histories, operational patterns, failure libraries) into assets their agents can learn from will compound their advantage with every cycle. Think of this as the AI equivalent of location in real estate: everyone can build, but not everyone is sitting on the right ground.
- **Human orchestration.** Here is where the metaphor shifts. In the pre-agentic world, humans were the musicians. In the agentic world, humans become the *conductors*. The value is no longer in playing every note. It is in knowing which instruments to bring in, when to hold back, when to accelerate, and how to make a hundred independent performers sound like one coherent piece. Your best managers are no longer the ones who do the most work. They are the ones who orchestrate the most agents, the most effectively, with the least damage.

This is a profound shift in what "talent" means. The conductor does not need to play the violin better than the first chair. The conductor needs taste, timing, and the ability to hear when something is off before the audience does. Those are Human Premium skills. They are not automated away by better AI. They are made *more valuable*.

The Capability Trap

The danger for most companies is what we call the **capability trap**: everyone adopts agents, nobody changes their governance, and the entire industry reaches a new equilibrium that is faster and cheaper but offers no competitive differentiation. You run harder to stay in the same place as you've mistaken capability acquisition for advantage creation.

This is exactly what happened with enterprise resource planning (ERP) systems in the 1990s and cloud computing in the 2010s. The technology became table stakes. The winners were the ones who used

the technology to rethink how they made decisions, not just how they processed transactions.

The rest of this chapter presents three frameworks to help you avoid the capability trap. Each one targets a different layer of the governance advantage: measuring what your agents cost, deciding what they should be allowed to do, and calibrating trust over time. Together, they form a system that turns agentic capacity into agentic *advantage*.

Three Frameworks, One System

The rest of this chapter is built around three interlocking tools:

- **The Agent Workforce Audit** measures what your digital labor is actually costing you and whether you're buying real outcomes or expensive hype.
- **The Agentic Decision Framework** assigns decision rights: what the agent should own, what it should recommend, and what a human must decide.
- **The Trust Dial** governs how autonomy changes over time for agents, whereby you expand, hold, or tighten based on evidence, not enthusiasm.

Think of these as the dashboard, the steering wheel, and the brakes (see Figure 9.2). You need all three. A car with a great dashboard but no steering wheel or brakes is just a very well-informed crash.

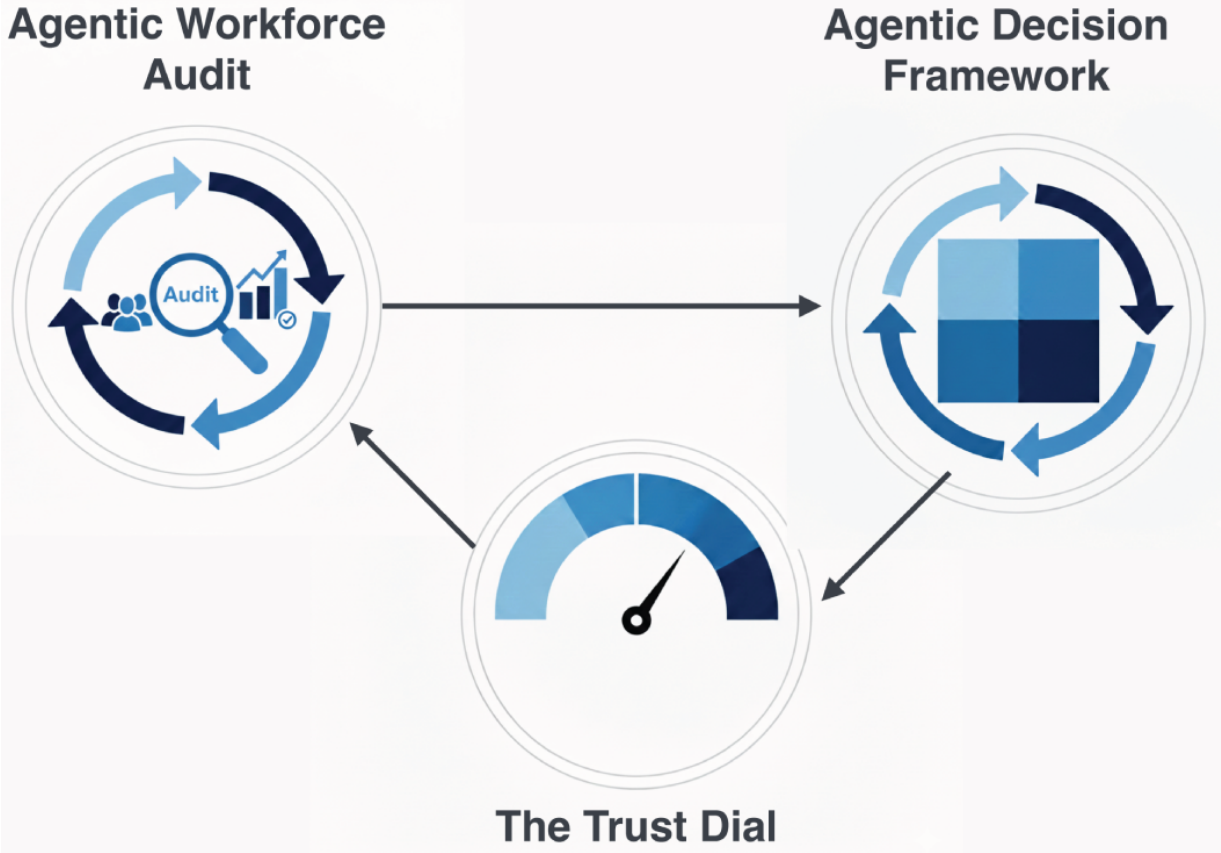


Figure 9.2 Managing a digital workforce: closing the loop on the new economics of AI

After we walk through each framework, we will run a single case study end-to-end through all three, so you can see the system work as a whole.

Framework 1: The Agent Workforce Audit

Before you add more agents, you need to know what your current ones are doing and what they're costing you. Not what the vendor says, but what your bottom line says. For reference, Gartner forecasts that more than 40% of agentic AI projects will be canceled by the end of 2027.[8] Not because the technology failed. Because organizations will not have measured the real cost.

The Agent Workforce Audit is a five-metric audit that separates real outcomes from expensive fiction (see Figure 9.3).

Metric	Stage 1 (Expanding Capacity) Moving from fixed headcount to variable, verified capacity.	Stage 2 (Solving the Bottleneck) Creating accountability where first-pass expertise is abundant, but the "accountability gap" creates a human gatekeeper crisis.	Stage 3 (Achieving Scale) Achieving non-linear productivity through workflow transformation and automated governance.

Figure 9.3 Agentic Workforce Audit: assess your agentic AI maturity level

It starts with one equation:

$$\text{Cost per verified outcome} = (\text{platform} + \text{model usage} + \text{integration} + \text{evaluation} + \text{security} + \text{oversight}) \div \text{verified outcomes}$$

That denominator is the keyword: *verified*. Not "resolved." Not "processed." Verified. Did the refund execute? Was the flight rebooked? Was the database actually updated? If you cannot define what "verified" means for your process, you cannot price it, govern it, or compare vendors.

The Three Maturity Stages

Organizations pass through three stages as they scale agentic work. You don't move up a stage until your metrics are green *and* your exception and rework rates are stable.

Stage 1: Capacity. Converting fixed head count into variable, verified capacity. The trap here is believing the economics are "cheap tokens." The API invoice is the tip of the iceberg. The real cost is the overhead stack: integration, monitoring, evaluation, permissions, security, and the human labor of exception handling.

Stage 2: Accountability. Addressing the bottleneck that occurs when first-pass expertise becomes widespread. Competent initial drafts are common, but the process stalls at approvals, edge cases, and risk sign-offs. Many companies get stuck here, celebrating that the agent can draft a contract in seconds, while overlooking that a lawyer still spends 45 minutes reviewing it.

Stage 3: Scale. Achieving nonlinear productivity through workflow transformation. One agent helping one person is tooling. A network of agents coordinating across systems is an operating model change, and it only works when governance is scaled as fast as capacity is.

The pattern will seem familiar to anyone who experienced the fiber buildout of the 1990s or the cloud rush of the 2010s. The losses seldom stem from the technology itself. Instead, they arise from unmeasured overhead and unclear value.

The Five Metrics

1. **Cost per verified outcome.** The all-in cost of one real result. Most organizations are counting two or three of the numerator's terms, which means their cost per outcome is fictional. That's why so many programs get scrapped as the business case collapses once someone does honest math.
2. **Exception rate.** The share of cases requiring human intervention at any point, be it the handoff, approval, correction, or escalation. This is the metric that reveals whether your "automation" is actually automating or just creating a faster assembly line of problems for humans to fix.

3. **Rework rate.** Everything that proves the first "outcome" didn't hold: reopened tickets, repeat contacts, downstream fixes, reversals, policy clean-ups. The exception rate indicates how often humans intervene. Rework rate tells you how often *they should have* but didn't.

Metric	Stage 1 (Expanding Capacity) Moving from fixed headcount to variable, verified capacity.	Stage 2 (Solving the Bottleneck) Creating accountability where first-pass expertise is abundant, but the "accountability gap" creates a human gatekeeper crisis.	Stage 3 (Achieving Scale) Achieving non-linear productivity through workflow transformation and automated governance.
1) Cost per Outcome	**Target:** Unit costs drop as volume rises. **Friction:** Hidden monitoring and security costs are "off-book." **The Play:** Include all overhead in the math; stop scaling until errors trend down.	**Target:** Savings without shifting "hidden work" to humans. **Friction:** Approval and risk sign-offs become a new labor bill. **The Play:** Pre-authorize low-risk actions; standardize approval routes.	**Target:** Costs improve in jumps after workflow redesign. **Friction:** Handoffs and compliance add "invisible" latency. **The Play:** Consolidate systems; enforce a single "source of truth."
2) Exception Rate	**Target:** Human "rescues" stay flat as volume grows. **Friction:** Tools are brittle or routing is too vague. **The Play:** Narrow the agent's scope; require verification before "done."	**Target:** Humans only touch true edge cases. **Friction:** Fear of accountability forces humans to hover. **The Play:** Publish a decision-rights matrix; make escalation rules explicit.	**Target:** Lines don't explode during high growth. **Friction:** New, complex errors emerge in new markets or segments. **The Play:** Segment workflows; implement "circuit breakers" for risk.
3) Rework Rate	**Target:** Repeat contacts or "re-dos" trend downward. **Friction:** No one agreed on what "success" actually looks like. **The Play:** Track "re-opens" as a primary KPI; define success in the system.	**Target:** Quality stays high even with human sign-off. **Friction:** Humans "rubber-stamp" approvals without checking. **The Play:** Require evidence for approvals; review the *decisions*, not just text.	**Target:** Quality holds steady as throughput increases. **Friction:** "Drift"—one fix causes three downstream errors. **The Play:** Use end-to-end testing; introduce "rollback" procedures.
4) Cycle Time	**Target:** Fast finishes, not just fast starts. **Friction:** High speed makes a mess left for the back office. **The Play:** Measure time to final completion; penalize "premature done."	**Target:** Quick decisions move the business faster. **Friction:** Humans become the slow "gate" for every approval. **The Play:** Batch reviews; create a "fast lane" for low-risk cases.	**Target:** Speed improves in big jumps after redesign. **Friction:** Managing the "mesh" of agents slows down coordination. **The Play:** Reduce handoffs; make agent interactions visible (logs/traces).
5) Agent Leverage	**Target:** Output per person rises as rules mature. **Friction:** Oversight feels like constant triage and babysitting. **The Play:** Limit autonomy until errors stabilize; invest in monitoring.	**Target:** Leverage rises with clear accountability. **Friction:** Ownership is fuzzy, so everyone hovers "just in case." **The Play:** Assign a named human owner per workflow; define boundaries.	**Target:** Gains accelerate as you cross volume thresholds. **Friction:** Oversight chores grow faster than the actual results. **The Play:** Use specialized agents; test for regression before every update.

Figure 9.4 Agentic Workforce Audit: outcomes, or expensive AI hype?

4. **Cycle time.** Time to a verified resolution. This is the headline metric, because managers see it and customers feel it. But cycle time alone is misleading. Fast and wrong is not a win.

Metric	Stage 1 (Expanding Capacity) Moving from fixed headcount to variable, verified capacity.	Stage 2 (Solving the Bottleneck) Creating accountability where first-pass expertise is abundant, but the "accountability gap" creates a human gatekeeper crisis.	Stage 3 (Achieving Scale) Achieving non-linear productivity through workflow transformation and automated governance.
1) Cost per Outcome	🟡	🟢	🟢
2) Exception Rate	🔴	🟡	🔴
3) Rework Rate	🟢	🟢	🔴
4) Cycle Time	🔴	🟡	🟡
5) Agent Leverage	🟢	🟢	🟡

Figure 9.5 Agentic workforce audit: customer support platform example

5. **Agent leverage.** Verified outcomes per human oversight hour. This is the bottom-line number. If one human can successfully govern ten agents, your leverage is 10x. If the exception rate forces a one-to-one ratio, your AI isn't a worker. It's a shadow department.

Consider the examples in Figures 9.4 and.

Framework 2: The Agentic Decision Framework

The audit tells you whether something is off. It doesn't answer the harder question managers face every day: what should the agent be

allowed to decide?

The Agentic Decision Framework answers that question with two variables (see Figure 9.6).

Reversibility: If this decision is wrong, how expensive is it to undo? Not just in money, but in trust, legal exposure, safety, and time. A password reset is reversed with a click. A discriminatory hiring pattern is reversed with a lawsuit.

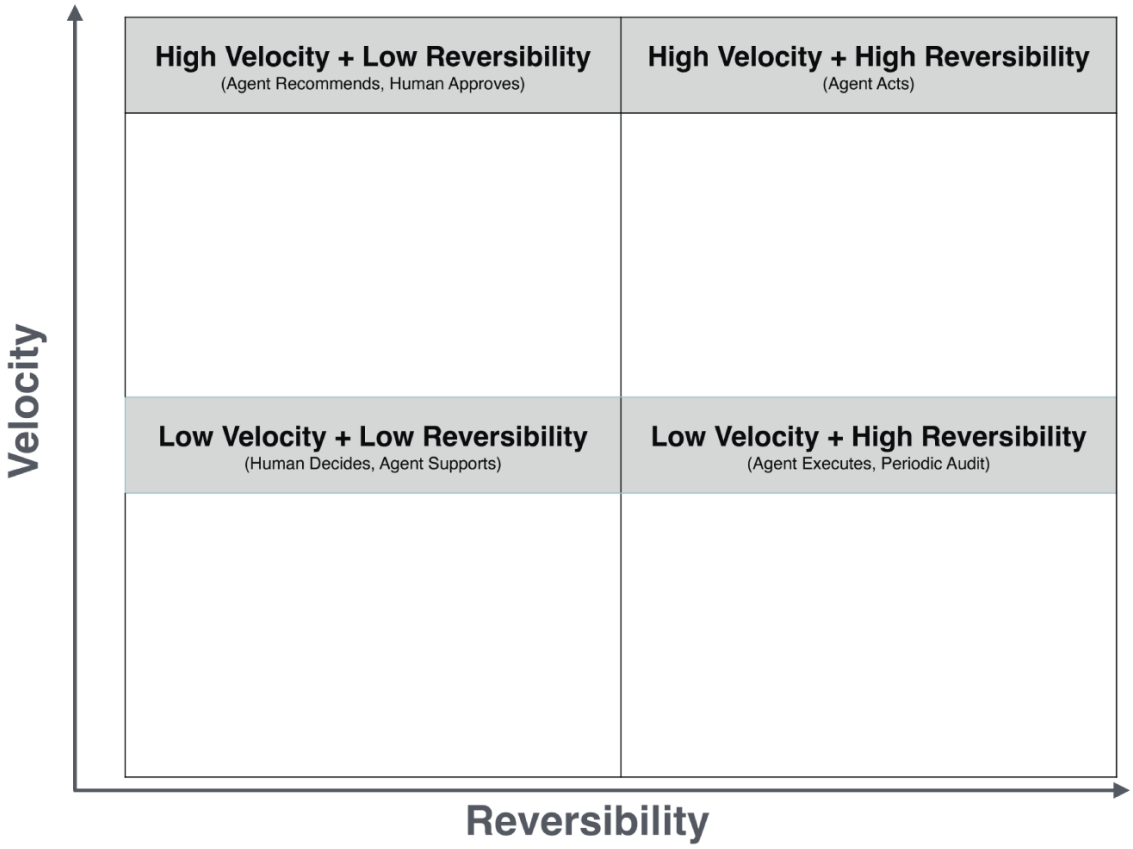


Figure 9.6 The Agentic Decision Framework: how to assign the right decisions

Velocity: How quickly must this decision be made to remain relevant? A fraud alert at 2 a.m. holds no value at 9 a.m. A market-exit decision is equally important whether made on Tuesday or Thursday.

Plot those two axes, and you get four quadrants. Each one has a different rule, a different failure mode, and a different warning sign.

Quadrant 1: High Speed + Easy to Reverse → Agent Acts

A customer texts, “Where’s my order?” at 2 a.m. The answer needs to come fast. If the agent gets it slightly wrong, a quick correction costs nothing.

The rule: Let the agent decide and execute. Humans review a sample after the fact.

The warning sign: Rework. On the surface, everything looks great. The queues are short, and response times are fast. But then you notice the same tickets reopening, the same answers getting quietly corrected. The agent is spinning its wheels, and no one notices because the dashboard is green.

Quadrant 2: High Speed + Hard to Reverse → Agent Recommends, Human Approves

A fraud alert fires at 11 p.m. Freezing the wrong account locks a loyal customer out of their money. Speed matters. So does getting it right.

The rule: The agent does the analysis at machine speed. A human reviews and approves before execution. Like a bank’s four-eyes principle: one prepares, the other signs off.

The warning sign: The approval step becomes a bottleneck. Someone will inevitably say, “The human is slowing us down, so let’s remove oversight.” Sometimes they’re right that the process is poorly designed. But removing the checkpoint entirely is how you get Amazon’s résumé-screening failure: fast recommendations, no scrutiny, qualified candidates quietly screened out by gender, and no way to undo the damage.[9]

Quadrant 3: Low Speed + Easy to Reverse → Agent Executes, Periodic Audit

A weekly report that pulls sales data, formats it, and drops it in a shared folder. Nobody needs it in real time. If something is off, you rerun it. This feels like the safest quadrant. That’s exactly what makes it dangerous.

The rule: The agent handles it independently. Humans check in on a regular schedule.

The warning sign: Complacency. The AI agent report arrives on schedule. The dashboard looks normal. But the numbers have been slowly drifting, and nobody noticed because the process appeared healthy. Recall Zillow's iBuying collapse from Chapter 3: an automated pricing model compounding small, systematic errors across thousands of transactions. Individual decisions looked reasonable. Without governance, small errors scaled into $881 million in losses.[10]

Quadrant 4: Low Speed + Hard to Reverse → Human Decides, Agent Supports

A company is deciding whether to exit a market, restructure a division, or sign a long-term vendor contract. No rush measured in seconds. Getting it wrong has consequences that unfold for years.

The rule: The agent prepares the groundwork by exploring options, pressure-testing assumptions, and flagging risks. The human makes the call. Like a well-run cockpit: the captain decides, but the entire system ensures the captain sees the full picture first.

The warning sign: Analysis paralysis. The agent produces more scenarios, more memos, and more options, while the human decision-maker spends more time reviewing and second-guessing than before. Knight Capital is the cautionary tale: a planned software deployment deemed safe enough for automated execution resulted in $440 million in losses in about 45 minutes.[11] (see Figure 9.7).

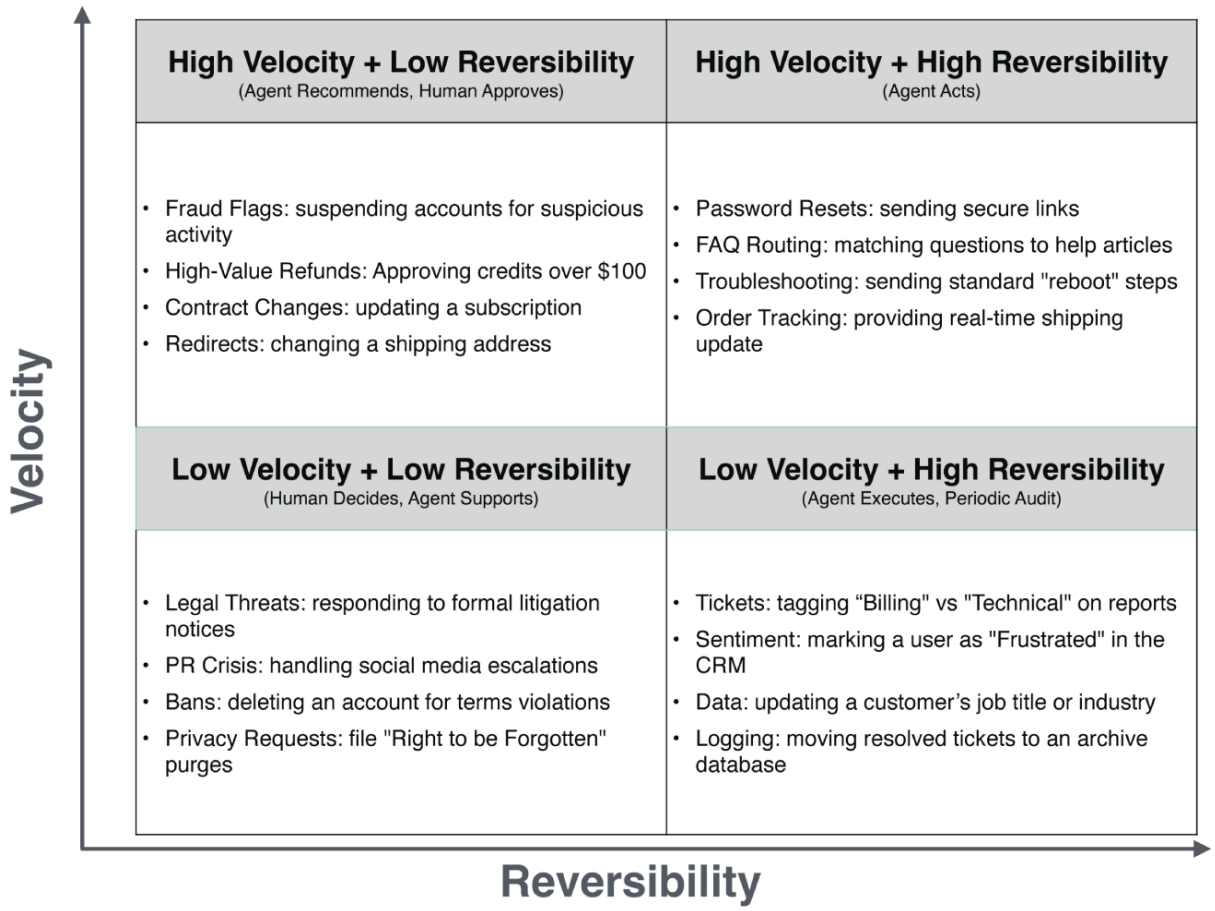


Figure 9.7 The Agentic Decision Framework: customer support example

How to Apply the Agentic Decision Framework

Start with one workflow. Pick a process that already has agents running. Then do five things:

1. **List the decisions, not the tasks.** “Respond to customer” is a task. “Issue a refund,” “escalate to a human,” and “close the ticket” are decisions. Write down 15–20 decisions inside one workflow.
2. **Score each on two axes.** The two axes cover how fast it must be made and how painful it is to undo. Keep it rough. Don’t build a precision instrument.
3. **Place each decision on the grid.** This is where most teams discover they were treating very different decisions as if they

were the same.

4. **Assign a governance model.** Agent acts. Agent recommends. Agent executes with audit. Human decides with agent support.
5. **Pick the leading indicator.** Every quadrant has a number that breaks first. If you cannot name that number, your agent is not ready.

Framework 3: The Trust Dial

This is the idea that will save you the most money and heartache. The audit tells you what's happening. The decision framework tells you who should decide. But neither answers the question that trips up every organization we've worked with: how do you change the rules over time?

An agent that belongs in Q2 today may earn Q1 autonomy in six months. Or it may need to be pulled back to Q4 after a bad quarter. The trust relationship between your organization and your AI is not static. It evolves. And most companies manage it badly, swinging between reckless optimism and panicked lockdown.

Here's the pattern we see repeatedly: A team deploys an agent. The initial weeks are impressive, with quick results, satisfied customers, and green dashboards. Confidence builds. The team relaxes the safeguards. Then something goes wrong. Leadership overreacts, halts the system, and everyone assumes the technology was overhyped. They blame the AI agent, but they should blame their governance.

The problem is not the technology. The problem is treating trust like a light switch—either on or off. The Trust Dial replaces the switch with a dial.

How the Dial Works

You simply turn the knob one click at a time, listening for static or a clear signal.

Every agentic workflow starts with a defined posture from the Agentic Decision Framework, a leading indicator, and a threshold. If

the indicator stays within the threshold for a sustained period, you turn the dial one notch toward greater autonomy. You turn the knob one notch at a time, watching your indicators, and only move further when the reading is stable.

This is the same logic behind canary deployments in software engineering. You don't roll out a major change to every user at once. You expose a small group, watch what breaks, and expand only when the system behaves. The same discipline works for AI decision rights.

What a "Notch" Looks Like

A notch is the smallest unit of delegated authority. Specific enough to test. Easy enough to reverse (see Figure 9.8).

A notch up might mean moving a narrow ticket category from "agent recommends" to "agent acts." Raising the refund limit from €20 to €35. Shifting from 100% human review to 25% sampling for a low-risk category.

A notch down might involve reverting a category to require human approval, lowering the confidence threshold, or increasing spot checks until exception rates stabilize.

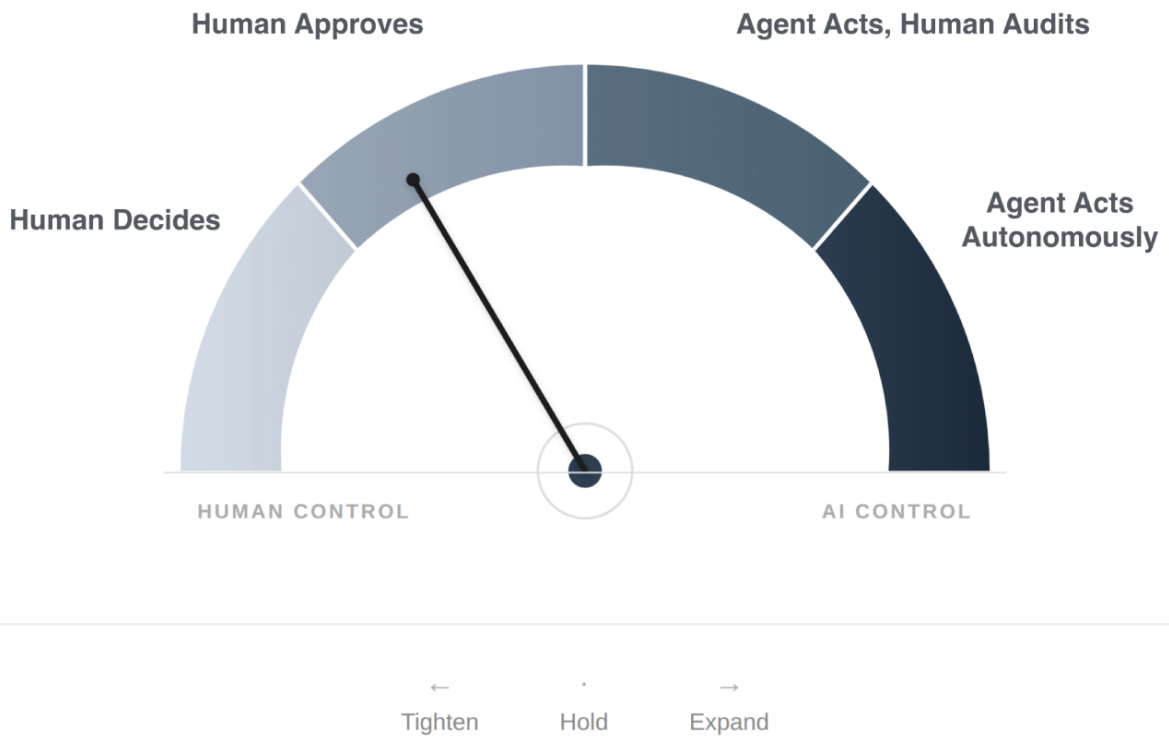


Figure 9.8 The Trust Dial: know when to tighten, hold, or expand autonomy

The Six Dials

In practice, the Trust Dial adjusts six settings:

Scope: Which decisions or categories the agent can handle

Thresholds: The dollar, risk, or confidence limits that trigger escalation

Access permissions: Whether the agent can observe, recommend, or act in live systems

Sampling rate: How often humans review outputs

Approval rules: Which actions require human sign-off

Fallback behavior: What happens when confidence is low, inputs are ambiguous, or the workflow fails

On review:

Keep it simple. Review one workflow every two weeks for the first two months. Look at the leading indicator you chose, plus one outcome metric from the Audit. Ask only three questions:

1. *Is the agent staying inside the threshold?*
2. *Is human oversight creating leverage or drag?*
3. *Should we tighten, hold, or expand by one notch?*

That's it. No committee. No 12-slide deck. Three questions, one decision, every two weeks. The discipline is in the cadence, not the complexity.

One Company, Three Frameworks: The Kooda Systems Story

Theory is cheap. Let's watch the system work based on a composite case drawn from clients we've worked with recently.

We'll call the company Kooda Systems, an Estonian fintech company that became the poster child for agentic AI in customer service. Their public claim: an AI assistant handling the equivalent work of 900 full-time agents, while cutting average resolution time from 11 minutes to less than 2. The headlines were breathless. The reality was messier.

Here's how the three frameworks explain both the wins and the near misses.

The Audit: Finding the Real Numbers

Kooda System's first move looked like a home run. Ticket volume dropped. Resolution times plummeted. The cost-per-interaction fell off a cliff. The management team was thrilled.

Then someone looked at the exception rate.

The agent handled the simple cases well. Most inquiries were about "Where's my order?" and "How do I return this?" questions, which made up most of the volume. But the remaining cases weren't just more difficult; they were *much more* difficult. Customers who

reached a human agent were already frustrated by the AI. They were angrier, their issues were more complicated, and the cost to resolve each case was increasing rather than decreasing.

A 25% drop in repeat inquiries was a genuine victory. But the remaining 75% carried higher exception complexity. The average cost per verified outcome, once you included the overhead stack, was better than the old model but not by the margin the headline numbers suggested. The Audit would have caught this in week two. Without it, it took months.

The Decision Framework: Sorting What Goes Where

Kooda's initial deployment treated customer service as a single category. The agent handled everything. This is the most common mistake in agentic deployment, and it's the one the decision framework prevents.

When you map Kooda's customer service decisions onto the grid, the picture becomes obvious:

> **Q1 (Agent Acts):** Order status, return labels, tracking updates, basic FAQ. High speed, easy to reverse. This is where the agent crushed it.
>
> **Q2 (Agent Recommends, Human Approves):** Fraud disputes, payment plan modifications, and account security issues. High speed, but hard to reverse. These needed a human checkpoint and initially didn't have one.
>
> **Q3 (Agent Executes, Periodic Audit):** Routine refunds under a threshold, standard policy adjustments. Low urgency, easy to reverse. The agent could handle this solo with periodic spot-checks.
>
> **Q4 (Human Decides, Agent Supports):** Regulatory complaints, high-value dispute resolution, cases with legal implications. No rush, hard to reverse. The agent should have been doing research and drafting options, not making calls.

The C-suite battle was exactly what you'd expect. The operations team wanted to keep everything automated, and the numbers looked

great. The risk team pointed to a growing pile of escalated complaints that were taking longer to resolve than before the AI, because by the time a human saw them, the customer was furious, and the situation had already compounded. The compromise was segmentation: different quadrants, different rules.

The Trust Dial: Calibrating over Time

Here is where the story gets interesting.

After segmenting decisions into quadrants, the Kooda team started moving the dial. The first notch up: simple refunds under €20 moved from “agent recommends” to “agent acts” with 25% sampling. The exception rate held. Two weeks later, they raised the limit to €35. Still stable. Four weeks later, to €50.

Then the rework rate spiked. Not dramatically, just a 3% uptick. But the Trust Dial protocol flagged it immediately. The team paused, investigated, and discovered that the agent was approving refunds for items with a pattern of serial returns. The individual decisions looked fine. The aggregate pattern was a vulnerability.

They turned the dial back one notch, lowering to €35 with 50% sampling for that product category, and added a serial-return filter to the agent’s logic. Two months later, they tried €50 again. This time, it held. That pattern of step forward, correction, retreat, adjustment, and stepping forward again shows how governance works in practice. It is not glamorous. It is not a moonshot. It is a dial that moves one notch at a time based on evidence.

Without the Agent Workforce Audit, they wouldn’t have seen the rework rate. Without the Agentic Decision Framework, they wouldn’t have known which decisions to segment. Without the Trust Dial, they wouldn’t have had a protocol for retreating gracefully instead of panicking.

Three frameworks. One system. That’s the centaur in action.

The Three Questions

We started this chapter with a claim: when labor becomes software, governance becomes the competitive advantage. Now you have the tools to act on it.

Every executive deploying AI agents should be able to answer three questions:

1. **What is the cost per verified outcome?** The Agent Workforce Audit answers this. If you can't answer it, you're flying blind.
2. **Which decisions are agent-owned versus human-owned, and why?** The Agentic Decision Framework answers this. If you can't answer it, you're gambling.
3. **How do you know when things are going off-track?** The Trust Dial answers this. If you can't answer it, you're one bad week from a headline you don't want.

Why the Dial Matters

The three questions are not just a checklist. They are the beginning of a compounding advantage.

Every time your team turns the Trust Dial one notch and watches what happens, you learn something your competitor does not know. Every time you correctly assign a decision to the right level and catch it when the leading indicator breaks, you build institutional knowledge that cannot be purchased from a vendor or replicated by a faster model.

The technology is the same for everyone. The models are the same. The tools are the same. The pricing is roughly the same. Six months from now, your competitor will have access to everything you have today.

What they don't have is the governance and learning architecture you built while they were still celebrating their dashboards. And that architecture, once built, does what Warren Buffett called a moat: it strengthens with every decision cycle. Not because the moat is wider, but because the knowledge behind it is deeper.

Every verified outcome, every calibrated notch, every caught failure spreads the moat by another foot. Your competitor cannot cross what they never helped dig. The firm that leads in the next decade will not be the one with the most agents. It will be the one with the most disciplined relationship with those agents, knowing precisely where to trust them, where to watch them, and where to keep a human hand on the wheel.

Gulliver's mistake was not that he admired the horses. His mistake was that he stopped learning how much he had in common with them.

Turn the dial. Keep your hand on it.

Notes

1. Irving Rothman suggests that Swift, who studied Hebrew at Trinity College Dublin, encoded Hebrew throughout the novel. Strikingly, "hnea Yahoo" in reversed Hebrew translates as "not God." According to Rothman, the Yahoo is the antithesis of the divine, the human stripped of every quality that would make them a higher being. (Irving Rothman, "Swift's Use of Hebrew in Gulliver's Travels," Swift Studies 30 (2015).)

2. PwC, "AI Agent Survey," May 16, 2025, https://www.pwc.com/us/en/tech-effect/ai-analytics/ai-agent-survey.html.

3. Bryan Walsh, "A Non-Coder's Guide to Claude Code," Vox, January 16, 2026, https://www.vox.com/future-perfect/475370/anthropic-claude-code-artificial-intelligence-coder-jobs.

4. Gartner, "Gartner Predicts Agentic AI Will Autonomously Resolve 80 Percent of Common Customer Service Issues without Human Intervention by 2029," March 5, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-03-05-gartner-predicts-agentic-ai-will-autonomously-resolve-80-percent-of-common-customer-service-issues-without-human-intervention-by-20290.

5. "The Pentagon vs. Anthropic: An A.I. Agent Slandered," podcast episode, *SMB Community Podcast*, Apple Podcasts, February 10, 2026, https://podcasts.apple.com/us/podcast/the-pentagon-vs-anthropic-an-a-i-agent-slandered/id1528594034?i=1000750653008.

6. Maxwell Zeff, "Anthropic's New AI Model Turns to Blackmail When Engineers Try to Take It Offline," TechCrunch, May 22, 2025, https://techcrunch.com/2025/05/22/anthropics-new-ai-model-turns-to-blackmail-when-engineers-try-to-take-it-offline.

7. Karissa Bell, "13-Hour AWS Outage Reportedly Caused by Amazon's Own AI Tools," *Engadget*, February 21, 2026, https://www.engadget.com/ai/13-hour-aws-outage-reportedly-caused-by-amazons-own-ai-tools-170930190.html.

8. Gartner, "Gartner Predicts Over 40 Percent of Agentic AI Projects Will Be Canceled by End of 2027," June 25, 2025, https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027.

9. Jeffrey Dastin, "Amazon Scraps Secret AI Recruiting Tool That Showed Bias against Women," Reuters, October 10, 2018, https://www.reuters.com/article/world/insight-amazon-scraps-secret-ai-recruiting-tool-that-showed-bias-against-women-idUSKCN1MK0AG.

10. Patrick Clark, "Zillow Shuts Down Home-Flipping Operation after Racking Up Losses," Bloomberg, November 2, 2021, https://www.bloomberg.com/news/articles/2021-11-02/zillow-shuts-down-home-flipping-business-after-racking-up-losses.

11. U.S. Securities and Exchange Commission, "SEC Charges Knight Capital with Violations of Market Access Rule," news release no. 2013-222, October 16, 2013, https://www.sec.gov/newsroom/press-releases/2013-222.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 10
The AI Leadership Blueprint: From Insight to Implementation

> *"The factory of the future will have only two employees, a man and a dog. The man will be there to feed the dog. The dog will be there to keep the man from touching the equipment."*
>
> —Warren Bennis

One can read Bennis's quote as a swipe at "the phony comfort of AI optimism" or, equally, "the phony comfort of AI pessimism," as the ongoing debate goes. Either way, you see it, it is not far from a realistic near future, and in some rooms it is already the present. Our benchmarks have changed. For decades pundits would measure a CEO's strength by comparing them to a Steve Jobs or a Jack Welch or a Indra Nooyi. This new yardstick is how well a CEO works with "the machine." And this goes for all C-suite positions. Think about how limited a CEO's knowledge is today. How much does a CEO really know about what's happening across a global company's operations in terms of AI?

Executives are provided with an avalanche of reports, digital dashboards, daily meetings, and weekly strategy reports. It is objectively and physically impossible to take in and process the full context of every production line adjustment, product launch, market development, personal relationship, and communication, all done across a score of time zones and tens of thousands of people. In this light, the relative evolutionary development of our human capabilities are woefully lagging to the capabilities of machines with their unfathomable data sets and the algorithms that power the logic behind their proven outputs.

We are not yet at the factory with its man and its dog. We may never get there. But the shadow of that future should be on every executive's mind, because most leaders are driving as fast as they can in the dark, without instruments, foot on the gas pedal, trying to catch up to a reality they have not yet seen. H.G. Wells captured

humanity's hope for clarity in times of technological change in 1895's *Time Machine*. "It sounds plausible enough tonight, but we must wait until tomorrow. Wait for the common sense of the morning." With AI, waiting for the morning is not an option.

When Humans Wag the AI Dog

For nearly 30 years, it has been hard to avoid the notion of "disruption." We've grown accustomed to dreaming about how we will saddle a new technology, subvert the status quo, and create tomorrow's global AI power. But let's go a little deeper. Netflix and Spotify are instructive examples of how disruption can take dramatically different forms, even within similar industries. Both companies began with technology-driven approaches to content distribution. Where previous generations of technology companies sold software to video rental stores and music retailers, these disruptors leveraged software to reimagine how media reaches consumers.

Yet the depth and breadth of their transformative impact diverged significantly. Netflix evolved from a DVD-by-mail service to a streaming platform. Most importantly, it didn't stop there. It fundamentally reconstructed the entire entertainment value chain. By producing original content, Netflix disrupted not only distribution but also how shows and films are approved, produced, released, and consumed. Their binge-watching model upended traditional episodic scheduling, their global simultaneous releases transformed marketing strategies, and their data-driven content creation process challenged conventional wisdom about audience preferences. Netflix didn't just change how we watch; it changed what we watch, where we watch, when we watch, and how it's made.

Spotify, while revolutionary in making music instantly accessible, largely remained within the distribution layer of the industry. Despite introducing algorithmic playlists and recommendation systems, Spotify did not fundamentally alter how music is produced, composed, or marketed. The recording industry's core business model and production processes remained largely intact, with

traditional record labels maintaining their role in talent development and promotion.

The disruption, though significant for retailers, didn't penetrate the creation and production processes to the same extent. As a result, in its 17 years in business, Spotify made its first profit in 2024 of $1.17 billion.[1] Netflix has been profitable every year since moving to streaming in 2007, with a 2024 profit of $8.7 billion.[2]

This contrast illuminates a crucial truth: disruption exists on a spectrum. Sometimes, it transforms entire value chains, as with Netflix, while in other cases, it primarily affects specific segments, as with Spotify. Every new technological development is potentially disruptive to someone, but the depth and breadth of that disruption varies tremendously by how a company's leadership shapes this change.

This variability underscores why it's always easier to proclaim "disruption!" or "AI!" than to thoughtfully analyze what form that transformation will take and how it should be managed. This is precisely why today's leaders require more than mere awareness of AI's disruptive potential.

This chapter is about precisely this more profound issue. It is about supercharging the "human side" of artificial intelligence in business. Here, we will introduce the *AI Leadership* and *Organizational Operating Systems*. C-Suite leaders in the AI era must intentionally architect their leadership and organizations; otherwise, disruption, technology, or market shifts will do it for them. It is important to note that the hands-on tools and frameworks to execute the principles in this chapter are provided throughout this book. One could say that these engage your firm at a 5 to 300-foot level. In this chapter, we are at the "big picture" ideas, ideals, and frameworks to lead at the 36,000-foot level.

The AI Leadership Operating System

The AI Leadership Operating System forms the structure for executives to onboard, maintain, and excel as an AI-literate company. It transcends superficial understanding to foster

systematic evaluations of how AI interacts and reshapes specific business functions, customer relationships, and competitive dynamics. Rather than reacting to technological change, those equipped with this operating system can anticipate shifts, capitalize on emerging opportunities, and mitigate potential threats with precision and confidence. As we have said throughout this book, successful AI transformation depends less on technical capabilities and more on effective leadership: transparent decision-making, executive collaboration, and the ability to quickly turn strategic insights into market action. Although the past can influence these leadership structures, market-creating leaders have always changed the leadership rulebook. C-suite leaders face a crucial choice: they can deliberately design their decision-making architecture to get the most out of AI or sit back and have AI change them.

Rethinking the Strategic Human-AI Partnership

A failure to challenge and change how decisions are made in a firm (decision-making architecture) is the greatest blind spot for companies attempting to lead in an AI world. By decision-making architecture, we mean the processes, structures, and rules by which leadership makes decisions, especially strategic and operational ones. But AI fundamentally alters this calculus, introducing an unprecedented blend of precision analytics, real-time predictive insights, probabilistic foresight, and many more unforeseen drivers.

The first step in creating an AI Leadership OS is to map the lay of the land, even when you are uncertain of what lies ahead. Recall how the English described Australia as the “gap in the map” for a century, not believing reports of the continent. As we move forward, nature is the best teacher for charting this new territory. Just as ecologists map an ecosystem (a jungle, for example) to define each species’ place in a complex ecosystem clearly, executives must map AI roles as biologists do through the following structured approach:

1. **Observation:** Carefully and systematically assess where AI excels or struggles, deliberately avoiding preconceptions or

biases.

2. **Data collection:** Gather concrete metrics on AI's decision quality, accuracy, speed, and outcomes across various contexts to form a robust, empirical understanding.
3. **Variable identification:** Identify and analyze critical factors such as complexity, data availability, ethical implications, and potential risks that significantly influence AI success or failure in diverse decision-making scenarios.
4. **Boundary definition:** Delineate boundaries distinguishing AI's intrinsic technical capabilities from its ethically and strategically appropriate human application domains.
5. **Adaptation tracking:** Establish ongoing processes and checkpoints to reassess and refine these boundaries as AI technology advances and business conditions evolve.

In the context of AI, we've found that these observations generally fall into a continuum across three zones.

- **Automation zone:** Decisions and processes entirely handed over to AI systems, primarily operational or tactical, such as real-time pricing adjustments, logistics routing, cybersecurity threat detection, and routine customer service. These areas demand explicit accountability structures, allowing executives to confidently delegate authority to automated systems, with built-in mechanisms to reassess automation effectiveness regularly.
- **Augmentation zone:** This area where AI enhances human judgment without replacing it. Examples include market forecasting, consumer sentiment analysis, and strategic scenario planning. Executives remain firmly in control but leverage AI-driven insights to sharpen their decision-making. Dynamic flexibility is crucial here, enabling executives to continuously recalibrate AI's contribution based on evolving capabilities and business contexts.
- **Off-limits zone:** Domain of decision-making where ethical, cultural, reputational, or strategic factors dictate human primacy. Such zones typically include crisis response, complex

personnel decisions, brand reputation management, and nuanced diplomatic negotiations. Clear protocols must be established to ensure AI does not inadvertently encroach on these sensitive areas.

We can thoughtfully balance resources, capabilities, and risk by explicitly mapping these zones and embedding flexible, cross-boundary management and diagnostics (see Figure 10.1). Although never perfect, without this clarity, executives risk either over-delegating sensitive decisions, potentially compromising ethical standards and strategic alignment, or underutilizing AI, thereby missing critical opportunities. When leaders clearly define where AI should automate work, augment human efforts, or stay uninvolved, they can use AI effectively while protecting their organization's values and strategic goals (see Figure 10.2).

AI in the Boardroom: Precision, Timing, and Strategic Dialogue

Board engagement with AI implementation shows concerning patterns: one survey reveals that only 14% of boards discuss AI at every meeting.[3] In another Deloitte Global Boardroom Program survey, 66% of board members report having limited to no knowledge or experience with AI, though this is an improvement from 79% in a previous survey.[4] Boards must transform from passive receivers of company information to active shapers of AI strategy. The AI Leadership OS demands long-term stewardship, which is currently missing. The traditional C-Suite/Board relationship needs revision. Boards can't remain detached from AI-literate companies, and executives must abandon the "keep the board out of our business" mindset. AI implementation requires all leadership levels, especially for the long-term strategic decisions that boards oversee.

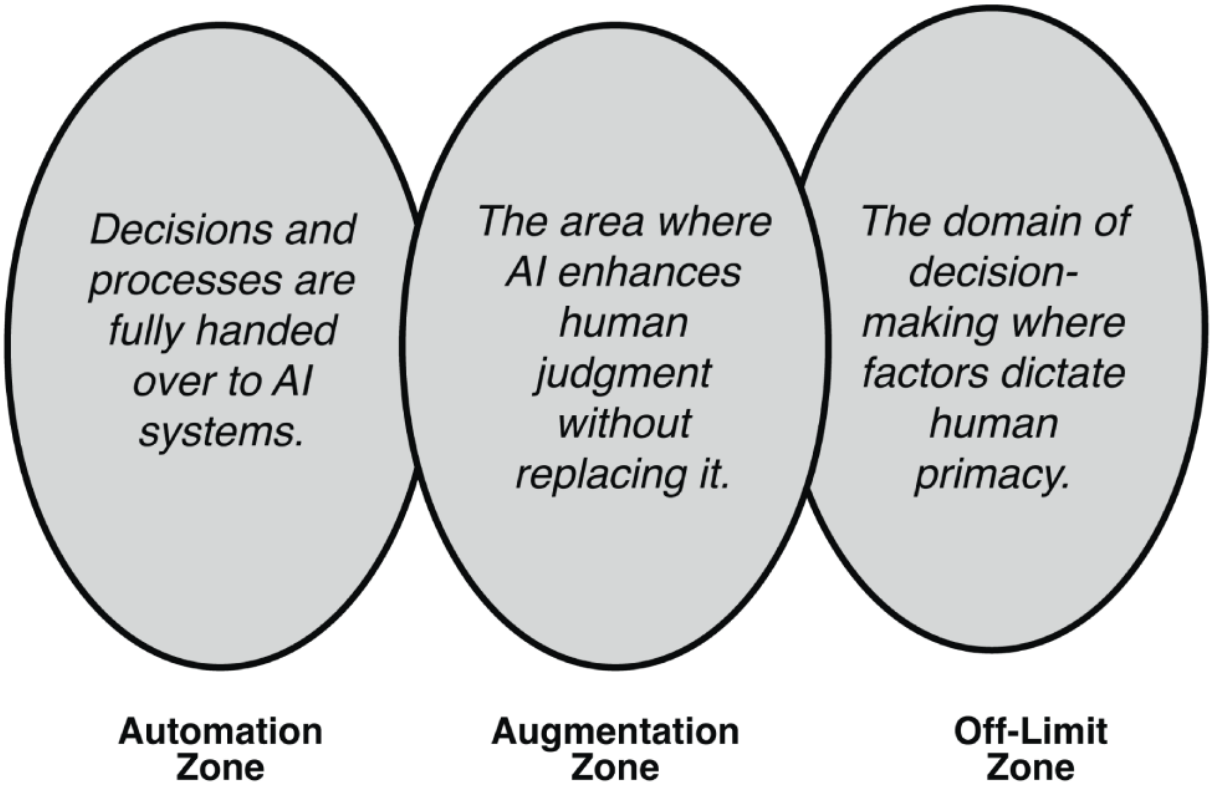


Figure 10.1 Zones of influence: mapping the lay of the land

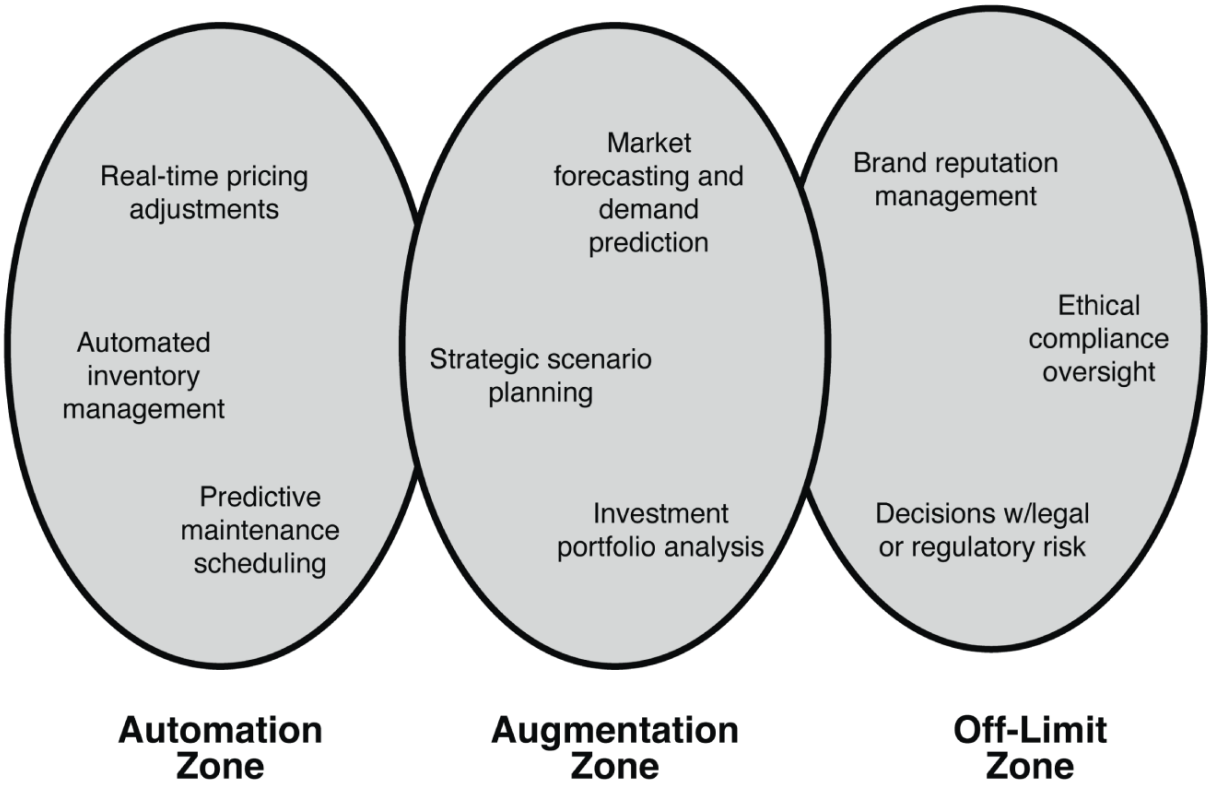


Figure 10.2 Zones of influence: example

The question remains: how can C-suite leaders foster greater, better-informed board engagement with AI? Simply, it is about creating new avenues of engagement, interaction, and inverting the learning process through four key steps.

STEP 1

- **Timing and cadence:** Establishing precise, regular intervals for AI updates in strategic dialogues, balancing frequent tactical reviews with deeper strategic explorations quarterly (see Figure 10.3).

Cadence	Dialogue Level	Scope and Objectives
Monthly	Tactical Updates	Quick AI status briefings, performance metrics
Quarterly	Strategic Reviews	Deeper dives into AI strategy, risk, alignment
Annually	Strategic Realignment	Holistic AI strategic review, recalibration

Figure 10.3 Timing and consistency protocol: establishing clear intervals for strategic and tactical AI discussions

STEP 2

- **Strategic insight board:** Equipping board members with structured, strategic questions to assess AI's maturity, impact, and alignment with organizational goals. This could/should be in the form of a C-Suite board-level online dashboard that answers strategic questions, such as:

  "How does AI meaningfully change our competitive positioning?"

  "Are we correctly calibrating the trade-off between AI-driven efficiency and risk exposure?"

  "Do we have adequate ethical safeguards to protect against bias, errors, or unintended consequences?"

STEP 3

- **Risk Awareness and mitigation:** Ensuring board members clearly understand AI-specific risks bias: transparency, cybersecurity vulnerabilities, and actively monitor mitigation efforts.

By carefully communicating these issues, executives elevate their AI governance's strategic quality and rigor, creating boards that act as critical strategic assets rather than mere compliance checkpoints (see Figure 10.4).

STEP 4

Make these consistent and committed rituals: Like any meaningful cultural shift, operationalizing AI leadership requires intentional and persistent rituals. This means consistent practices institutionalizing the AI mindset and ethical guidelines between the board and C-suite. These rituals serve as organizational anchors, reinforcing clarity, accountability, and alignment in a rapidly evolving environment that has not been seen before (see Figures 10.5–10.7).

Metric	Inquiries	Status
Competitive Positioning	"How does AI meaningfully change our competitive advantage?"	
Risk-Reward Calibration	"Are we correctly balancing AI-driven efficiency with risk exposure?"	
Ethical Integrity	"Do we have robust safeguards against bias, error, and unintended consequences?"	
Maturity & Capability	"Are we investing adequately in skills and infrastructure to sustain AI effectiveness?"	

Figure 10.4 AI Strategic Insight Board: online, interactive platform enabling structured strategic updates by board members

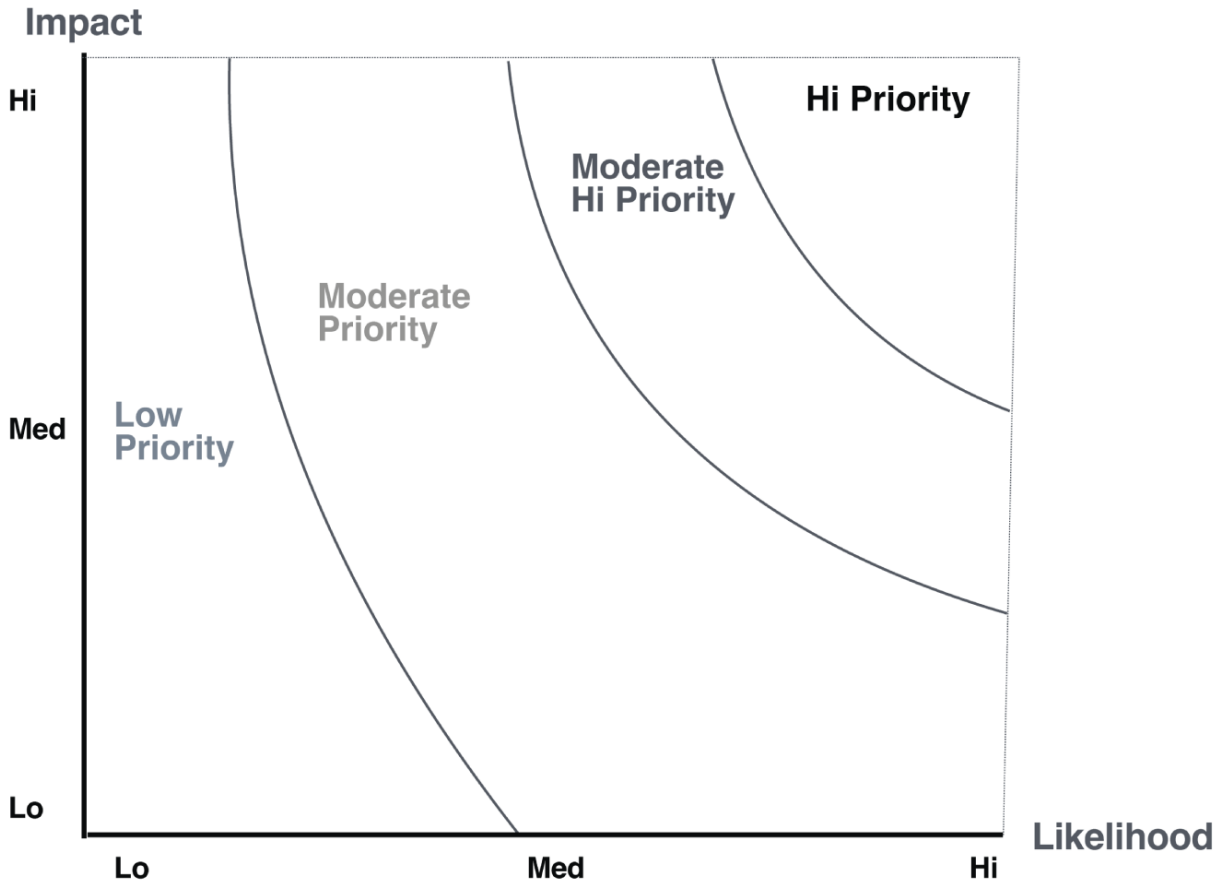


Figure 10.5 AI risk heat map: ensuring board members clearly understand AI risks and mitigation

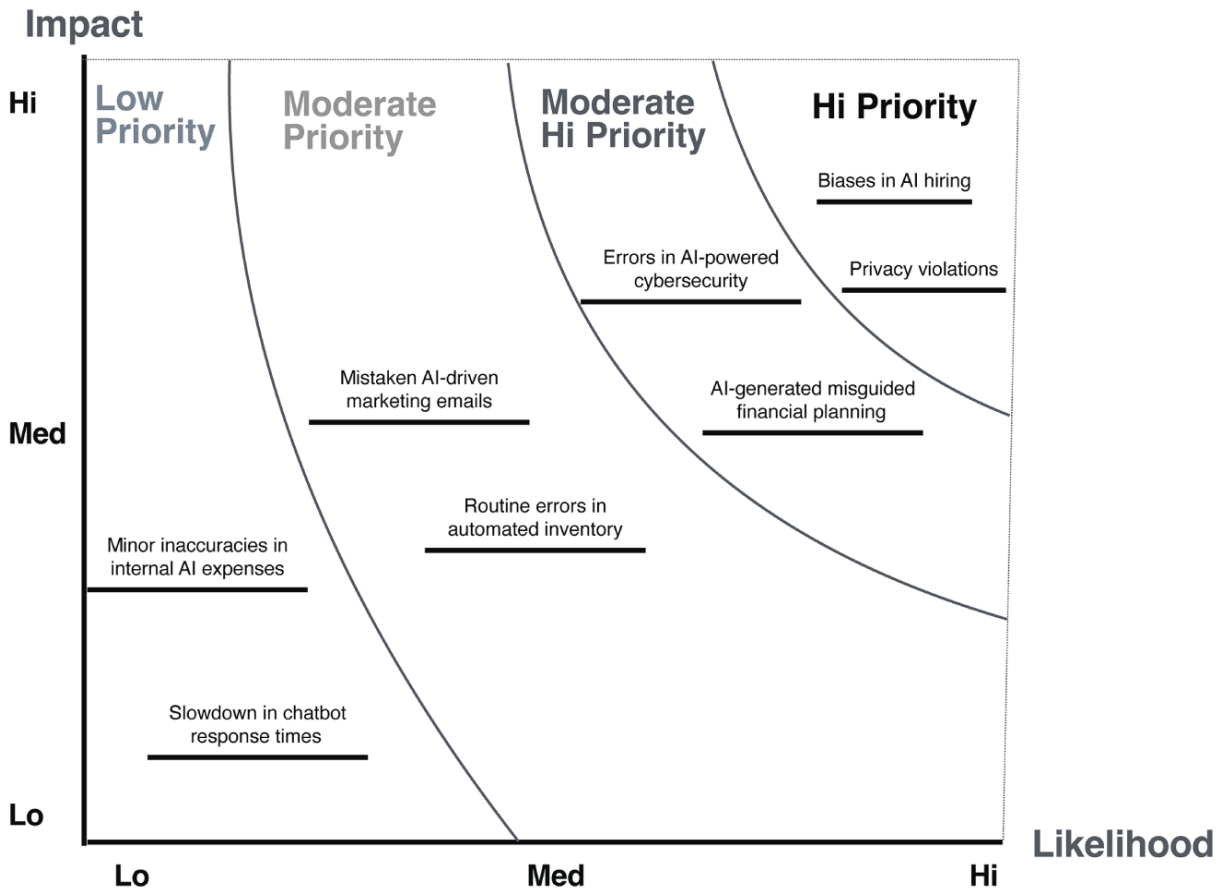


Figure 10.6 AI risk heat map: example

Ritual	Frequency	Purpose
AI Insight Briefings	Monthly	Update board members on AI trends and outcomes
Board/C-Suite Summits	Semi-Annual	Align strategic priorities between board and execs
Annual AI Risk Audit	Annually	Comprehensive assessment of AI governance and risk

Figure 10.7 Consistent and committed rituals: establishing institutional practices to embed AI literacy and governance deeply into the culture

The AI Leadership OS Canvas: A Practical Health Check

In the spirit of this book, we recognize that executives require a tangible, structured, and highly visual framework to effectively operationalize AI principles. The AI Leadership OS Canvas provides precisely this: a pragmatic visual monitoring framework that helps you and your team assess your current AI readiness and plan where you need to be tomorrow.

On the X-axis are the acts and activities that make up your AI leadership architecture. Here is a sample set (see Figure 10.8).

1. **AI framing:** Defining clearly when and how AI can add value, pinpointing specific problems or opportunities for AI

applications

2. **AI piloting:** Running rapid, iterative, controlled experiments with AI solutions before broader deployment
3. **AI integration:** Embedding validated AI tools systematically into organizational workflows and ensuring smooth human-AI collaboration

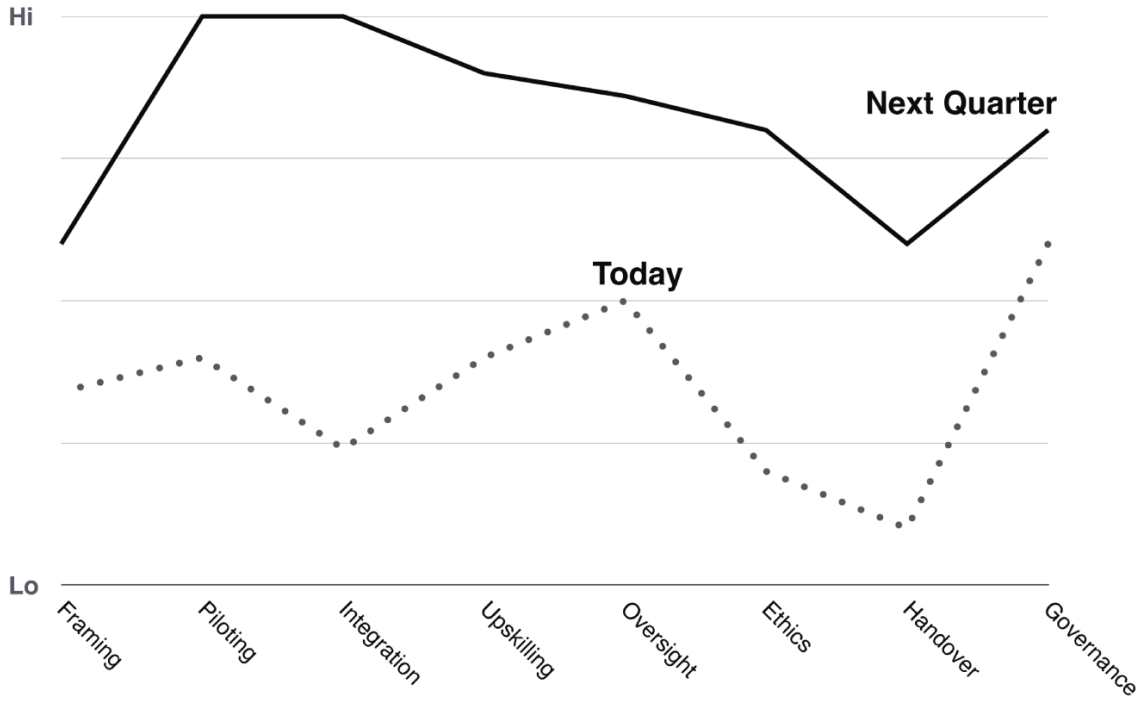


Figure 10.8 AIOS Canvas: a pragmatic visual monitoring framework

4. **AI upskilling:** Equipping leadership and teams with specific skills, knowledge, and mindsets necessary to leverage AI effectively
5. **AI oversight:** Continuously monitoring AI systems' performance, safety, transparency, and impact, adjusting oversight structures as needed
6. **AI ethics:** Establishing explicit operational guidelines and protocols for responsible AI use, fairness, accountability, and transparency
7. **AI handover:** Strategically managing decision boundaries and clearly defining tasks or decisions handed over from humans to

AI, and vice versa

8. **AI governance:** Implementing robust governance frameworks, clear accountability structures, and policies for managing data, risks, and compliance related to AI

On the Y-axis is how high or low you believe how you have met your objectives with each act and activity. This exercise is not driven by data but a highly subjective monitoring tool. It is the centerpiece of discussion, like the ancient Greek Agoras or eighteenth-century Parisian Salons. It is where Team Human discusses where we are today and where we want to be tomorrow.

This downloadable and customizable canvas allows you to systematically chart your AI transformation, track progress, ensure clarity, and consistently align tactical actions with your strategic ambitions. Again, it is a subjective point of reference for Team Human to get your heads around your relationship with AI.

Leading over Adapting: The AI Organizational Operating System

Ultimately, the AI Leadership Operating System represents a strategic necessity for navigating and thriving in an AI-driven business. C-suite must recognize their role not as mere technology adopters but as proactive architects of an AI-adept organization. By actively defining the “gap in the map,” enhancing board-level dialogues, institutionalizing rigorous governance rituals, and operationalizing these through tools like the Leadership OS Canvas, you are not merely adapting to the AI era but creating your future.

The AI Organizational Operating System

Successfully navigating the transition into an AI-literate business requires more than visionary leadership; it demands an organizational capability to translate strategic intent into practical, operational reality. The AI Leadership Operating System equips executives with clear decision-making frameworks, strategic clarity, and ethical guidance. Yet without an equally robust AI

Organizational Operating System, even the most visionary strategies risk becoming hollow aspirations. Think of the AI Leadership Operating System as a strategic blueprint: it outlines precisely what decisions leaders should own, where AI augments or automates, and the ethical and strategic boundaries to uphold.

In contrast, the AI Organizational Operating System is the engine that transforms this blueprint into tangible outcomes thereby defining processes, workflows, data governance, technology infrastructure, and the skills needed across the workforce to execute effectively at scale.

Organizations that overlook this critical distinction risk falling into one of two common pitfalls: strategic misalignment, where teams lack clarity on operationalizing executive vision, or operational fragmentation, where tactical AI initiatives proliferate without strategic coherence, wasting resources and diluting results. To thrive, AI-literate C-suites must deliberately construct and nurture both operating systems, aligning strategic vision with tactical execution to fully unlock AI's transformative potential.

The Constraints of Traditional Structures

A 2025 survey of executives across 101 countries found that rewiring workflows has the most considerable effect on an organization's ability to see EBIT impact from its use of AI.[5] In the age of artificial intelligence, many organizations find that their traditional ways of working hold them back. The penultimate piece of our journey brings us to a common theme: adopting AI is not just about new technology or analytical talent. It requires a new organizational operating system. Just as outdated software can prevent cutting-edge hardware from reaching its potential, an obsolete organizational structure can snuff AI's possibilities. In this section, we explore why conventional hierarchies and silos struggle with AI and what a reimagined, AI-literate organization might look like. We'll outline a vision for a more agile, networked company built on collective intelligence where humans and machines collaborate seamlessly and finish with concrete steps for executives to make this vision a reality.

Most companies still rely on organizational architectures suited to slower, more predictable environments. These organizational models have their uses but often prove inflexible and slow to adapt to AI adoption changes. The classic corporate hierarchy that is rigid, layered, and control-focused can slow decision-making, inhibit rapid experimentation, and obscure accountability precisely when agility matters most.

AI thrives on fast iteration and quick data-driven decisions, which hierarchical bureaucracies inherently resist. Another major challenge is organizational silos, where departments and teams operate in isolation.

These silos fragment data, block critical communication, and hinder collaboration, undermining AI initiatives that rely on integrated datasets and cross-functional insights. So often, we find essential knowledge hidden in departmental silos, "typically controlled by a department and impossible for other information seekers to find." Traditional organizations also lock employees into narrow roles, making it difficult to mobilize the multidisciplinary teams AI demands.

Job descriptions are narrowly defined, and people tend to "stay in their lane." But AI implementations cut across functions requiring multidisciplinary contributions and sometimes entirely new roles (like prompt engineers or AI ethicists). Rigid job descriptions and career paths impede this necessary flexibility and discourage rapid skill adaptation.

Finally, legacy decision-making processes, characterized by lengthy approvals and risk aversion, often stifle the experimentation essential to front-footed AI projects. By the time decisions are made, opportunities may be lost. Despite significant investments in AI, many organizations remain stuck in pilot stages or ad-hoc experiments due to organizational and psychological inertia. Research confirms that fewer than 10% of companies effectively embed AI into their operations, primarily due to a failure to redesign their structures and processes to support AI at scale.[6]

Only a tiny minority have truly woven AI into the fabric of their operations. A study concluded that the fundamental reason is "a

failure to rewire the organization." Companies hit cultural and structural barriers that technology alone cannot fix. Similarly, a 2025 report noted that "most organizations aren't structured to embrace AI at scale" and that the technology is advancing faster than the ability of companies to adapt their processes and structures.[7]

In short, hierarchy, silos, static job boundaries, and bureaucratic sclerosis are antithetical to the fast, data-fueled, cross-cutting nature of AI. These legacy traits create friction that can derail even the best AI strategies. Leaders must fundamentally reimagine their AI Organizational Operating System to overcome these barriers.

Designing the Organization for the AI Literate World

Just as operating systems have evolved from handling one task at a time to seamlessly running dozens of applications simultaneously, organizations must evolve from rigid hierarchies to adaptive networks where human and artificial intelligence operate in parallel, continuously learning from each other. In the rapidly emerging AI era, companies are finding success with flatter, more fluid, and far more collaborative structures than before. In place of the top-down pyramid, they are embracing decentralized networks of teams. Instead of rigid departments, they foster agile, cross-functional squads that form and re-form as needed. Rather than fixed jobs, they promote flexible roles that allow people to apply their skills in multiple contexts. The guiding principles of an AI-literate organization are adaptability, speed, and integration. Let's break down the key elements of this operating system.

Cross-Functionality

Agile "teams of teams" are at the heart of AI-centric organizational design. These teams bring together experts from different domains: for example, pairing data scientists with product managers, engineers, and customer service reps on a typical project. By working side-by-side on shared goals, they collapse the silos that previously separated knowledge areas. This cross-pollination is invaluable: a

mix of perspectives helps ensure that AI solutions are grounded in real operational needs and diverse viewpoints.

We have found time and again that AI has the biggest impact when cross-functional teams develop it with a mix of skills and perspectives. When the “suits,” technologists, and analysts collaborate from day one, AI initiatives address broad organizational priorities rather than isolated departmental metrics. A cross-functional AI team developing, say, a predictive maintenance system can anticipate the workflow changes it will require in the field or involve end-users early so adoption later is smoother. This approach contrasts sharply with the old model where IT built something in a vacuum and tossed it over the wall to the unprepared business suits.

Agency

Decentralization and empowerment go hand in hand with this team-of-teams structure. While the overall AI strategy is top-down in this organizational model, there are zones of augmentation clearly defined in your Zones of Engagement tool outlined earlier, where decisions and solutions are developed at the bottom up. To move at the speed of AI, tactical decisions should be made as low to the ground as possible by the people closest to the data and the customers.

These teams need autonomy to experiment, iterate, and respond quickly without running every choice up the previous command and control hierarchy. And many organizations are testing this out and radically flattening their hierarchies thereby minimizing layers of management and widening spans of control to push authority outward. It’s the cumbersome management layers that are the enemy of speed and agility. In genuinely agile organizations, you might see only three layers between the CEO and the front lines, whereas traditionally, there might be seven or eight. Fewer layers mean information travels faster; teams can act on insights in real time. Managers in this new model act less like gatekeepers and more like enablers or coaches, providing guidance and clearing roadblocks so teams can run with ideas. When you empower employees (within guardrails) to make data-backed decisions, you unlock the creative potential of AI throughout the firm.

Roles

Another hallmark of the AI-optimized organization is role fluidity. Rather than pinning people to one job title or silo forever, these companies create a more dynamic talent marketplace internally. Employees might rotate through different teams or contribute to multiple projects based on their skills and current priorities. Someone in marketing with a knack for analytics might spend a stint embedded with the data team to develop an AI-driven customer segmentation model. An engineer might split time between product development and an AI ethics task force. This fluid deployment of talent ensures that the best expertise is directed to bear on each problem, and it also keeps employees engaged in learning new skills.

In a fast-changing technological environment where top-talent employees are kept engaged, it's hugely beneficial for your employees to upskill and reskill continuously. Organizations can support this by defining roles flexibly (focusing on skills and outcomes rather than narrow tasks) and recognizing and rewarding collaborative contributions. It's a shift from viewing people as production units in a static machine to viewing them as adaptive and creative sparks of innovation in a network where connections and knowledge flow matter more than formal titles. The payoff is twofold: the organization becomes more resilient and responsive, and employees gain more agency with opportunities to grow with AI rather than feel threatened.

Culture

Underpinning all of this is a culture of open information sharing and teamwork. Structural change will fall flat if the culture still incentivizes hoarding information or sticking to the status quo. Leaders in AI-forward companies cultivate norms of transparency, curiosity, and continuous improvement. They break down internal competition and replace it with shared missions. Cross-functional teams are celebrated, not seen as exceptions. Communication channels are reworked so that insights from one part of the business can quickly find their way to others.

For example, an AI insight discovered in an e-commerce division (say, a new pattern in customer behavior) would be promptly

communicated to the product design and supply chain teams so they can jointly learn and act on this insight. This requires trust and an ethos that "we're all in this together." Some organizations formalize this by creating communities of practice around AI and data, with members from different departments exchanging learnings regularly. Others set up internal "crowd-solving" platforms where anyone can pose questions or share ideas, and colleagues from anywhere can contribute solutions. The goal is to simulate a startup's free-flowing, collaborative spirit, even at an enterprise scale.

AI Tech Support

Finally, an optimal AI-era structure also acknowledges the importance of new support *and* experimentation systems. As agile multifunctional teams proliferate, it is a good idea to establish platform teams (e.g., a central data platform group or AI center of excellence) that provide self-service tools, reusable data pipelines, and best practices to the rest of the organization. In other words, an AI process "tech support" system where best practices flourish across the organization. This avoids each team reinventing the wheel and ensures consistency and security standards.

Center for Strategic Prompting

On the proactive side of the equation is the need to create best practices, such as a Center for Strategic Prompting. This internal center focuses on mastering prompt design for AI systems, strategically aligning AI outputs with organizational objectives, and developing best practices and shared learning. By consolidating expertise and maintaining an internal library of optimized prompts, the center ensures consistency, strategic clarity, and improved AI performance across the company.

Additionally, organizations can benefit from creating internal marketplaces for human-AI hybrid teams. With AI, we can atomize skills, experiences, capabilities, etc., down to microscopic levels. In this marketplace, human and AI resources are distributed across traditional departmental boundaries based on strategic needs and project demands. Such marketplaces enhance organizational agility, improve resource efficiency, and foster cross-functional

collaboration, helping companies rapidly deploy optimal combinations of human expertise and AI capabilities to meet evolving business priorities.

Creating Your AI Organizational Operating System

Let's turn these ideas into a practical agenda. How can leaders install this AI Organizational Operating System in their companies? The following are specific action items for the C-suite to consider as you guide your organization toward an AI-powered, collectively intelligent future:

1. **Flatten hierarchies and empower teams**: Review your organizational structure to reduce layers and decentralize decision-making. Push decision authority closer to front-line teams, especially for real-time or AI-informed decisions. Increase managers' spans of control, streamline approvals, and clarify that managers' roles are to coach and enable rather than micromanage. Flattening hierarchies creates the flexibility needed for agile responsiveness.
2. **Break down silos and promote cross-functional collaboration:** Actively dismantle organizational silos by forming cross-functional AI project teams (for example, a team involving marketing, IT, data science, and customer service). Encourage departments to share data openly, through centralized data repositories or interoperability standards. Rotate leaders across divisions and reward collaborative behaviors to ensure AI solutions are adopted broadly rather than confined to individual silos.
3. **Build a robust data and technology backbone for AI**: Ensure the necessary infrastructure supports AI at scale by investing in unified data platforms or cloud architectures accessible across the organization. Choose user-friendly, scalable AI and analytics tools, including self-service dashboards that require minimal coding. Democratize data and tool/framework access, supported by robust governance

standards for data quality, security, and ethics. This strong digital foundation is essential for effective and integrated AI deployment.

4. **Redefine roles and reskill your workforce:** Prepare a tech-savvy, adaptable workforce by shifting to a skills-based talent model emphasizing data literacy, experimentation, and multidisciplinary teamwork. Launch comprehensive (and rewarded) AI training to increase employee understanding and collaboration with AI tools. Redesign roles and teams around AI capabilities where we are elevating human roles to higher-value tasks when AI handles routine activities. Foster continuous learning by recognizing new skills and providing resources for upskilling, ensuring organizational flexibility around AI-driven projects.
5. **Lead cultural change from the top**: The executive team must actively communicate and embody the mindset shift required for AI integration. Clearly articulate how AI will transform the business and openly share successes and insights from previous experiments. Encourage a test-and-learn approach throughout the organization, modeling data-driven decision-making and visibly supporting AI initiatives. Promote cross-functional collaboration by facilitating joint strategy sessions and celebrating integrated team efforts. Leadership must set the tone for embracing AI as a collaborative partner.

At the end of the day, your job in the C-suite is to steer your organization toward AI maturity. It is as much an exercise in organizational innovation as it is in technological innovation. By rethinking structure from hierarchies to networks, from silos to fluid teams, from rigid roles to continuous learning, you can create an environment where AI truly becomes a force multiplier. The companies that succeed in this transition will pair with AI to become adaptive, intelligent enterprises, able to sense and respond to the future's challenges and opportunities as one cohesive, AI-augmented whole. As a leader, you have the chance to upgrade your organization's operating system for the AI age, starting now.

Collective Intelligence Framework: Human-AI Collaboration

Having established both an AI Leadership Operating System and an AI Organizational Operating System, the final vital step is implementing a Collective Intelligence Architecture. This strategic framework systematically integrates human insight with AI capabilities, effectively bridging strategic intent and operational execution. Rather than merely supplementing existing systems, this framework provides a critical foundation to leverage the best of human intuition, creativity, and ethical judgment with AI's analytical power, scalability, and efficiency.

The Collective Intelligence Framework

The Collective Intelligence Framework is a structured methodology designed explicitly to orchestrate human and artificial intelligence interactions within an organization. This comprehensive approach clearly defines roles, responsibilities, and collaborative processes, ensuring that each human or AI entity is deployed effectively to maximize strengths and mitigate limitations. This architecture goes much further than ad-hoc human-AI collaboration, setting explicit operational guidelines, enhancing productivity, strategic clarity, and decision-making effectiveness.

The Blueprint: From Insights to Implementation

To state it plainly, by now, you have seen that the pace of change has rendered traditional siloed or human-only approaches inadequate. AI systems are exceptional at processing vast data sets, automating routine tasks, and generating predictive insights at scale. However, they lack the nuanced understanding, ethical discernment, and creative problem-solving capabilities intrinsic to humans. The Collective Intelligence Framework strategically merges these complementary strengths, empowering organizations to outperform those relying solely on either humans or AI. Firms employing this integrated architecture will see massive improvements in decision-

making accuracy, operational agility, innovative potential, and responsiveness to market dynamics.

Core Components of the Collective Intelligence Framework

The following are the core components of the collective intelligence framework.

Explicit Human-AI Role Definition

Assigning roles ensures optimal task execution:

- **Human roles:** Humans excel in strategic direction, ethical decision-making, navigating complex interpersonal relationships, contextual interpretation, and innovative problem-solving.
- **AI roles:** AI excels in routine decision automation, rapid large-scale data analysis, real-time predictive analytics, pattern recognition, and anomaly detection.

Each project or strategic initiative explicitly identifies tasks and allocates them accordingly, reducing ambiguity and enhancing overall effectiveness.

Example Scenario

In strategic market positioning, AI analyzes extensive market data in real time, quickly identifying emerging consumer trends and competitive movements. Human strategists then interpret these insights within broader market contexts, adjusting strategic priorities or marketing tactics accordingly.

Integrated Hybrid Team Structures

Establishing cross-functional teams that combine human talent with AI tools is essential:

- Assemble cross-functional teams with diverse expertise, including domain specialists, data scientists, operations experts, and AI analysts.
- Deploy flexible team structures that quickly adapt to evolving strategic priorities and market demands.
- Ensure seamless workflow integration between human decision-making and AI outputs.

Example Scenario

A hybrid customer experience team might integrate marketing professionals, customer service specialists, data analysts, and advanced AI-driven personalization systems to craft highly tailored customer interactions, significantly enhancing customer satisfaction and retention rates.

Robust Feedback and Iterative Learning Loops

Continuous improvement is central to the framework's effectiveness:

- Regularly gather structured qualitative feedback from human participants and quantitative performance data from AI systems.
- Implement regular review cycles to analyze collaborative effectiveness, refining processes and tools based on insights and experiences.
- Promote an organizational culture that values experimentation, adaptation, and iterative improvement.

Example Scenario

Following each significant initiative, conduct structured post-mortems where team members review AI analytics alongside human insights, generating practical recommendations for improving future project workflows and outcomes.

Visualization and Implementation Tool: The Collective Intelligence Framework

The Collective Intelligence Framework serves as a practical visual guide, organizing team collaboration into four actionable quadrants:

- **Human inputs:** Strategic insights, ethical oversight, creative thinking, stakeholder engagement
- **AI inputs:** Data-driven analytics, predictive insights, automation, anomaly detection
- **Joint processes:** Mapped workflows integrating human and AI contributions
- **Outcomes and feedback:** Metrics for performance evaluation, ongoing tracking, and structured feedback processes

Teams use this framework at project initiation, clearly setting expectations, clarifying responsibilities, and enhancing team synergy from the outset.

C-Suite Implementation Actions

To operationalize the Collective Intelligence Framework effectively, executives should:

- Communicate and institutionalize human and AI roles across the organization
- Launch strategic hybrid teams as pilots, documenting insights to refine future implementations
- Establish systematic feedback mechanisms ensuring accountability and continuous organizational learning
- Broadly adopt the Collective Intelligence Framework to standardize project planning and execution processes across the organization

Long-Term Benefits and Organizational Impact

Adopting a robust Collective Intelligence Framework positions organizations to sustain competitive advantages. It enhances agility,

enabling quicker strategic pivots in response to market fluctuations. It significantly boosts innovation, harnessing the full creative potential of human teams enhanced by AI analytical precision. Over time, this architecture fosters deeper organizational learning, superior decision-making quality, and a resilient culture adaptable to disruptive forces.

Ultimately, the Collective Intelligence Framework is more than a methodology (see Figure 10.9). By consciously designing human-AI collaboration, organizations achieve unparalleled levels of innovation, operational excellence, and strategic resilience, firmly establishing their competitive edge in a continually evolving digital business landscape.

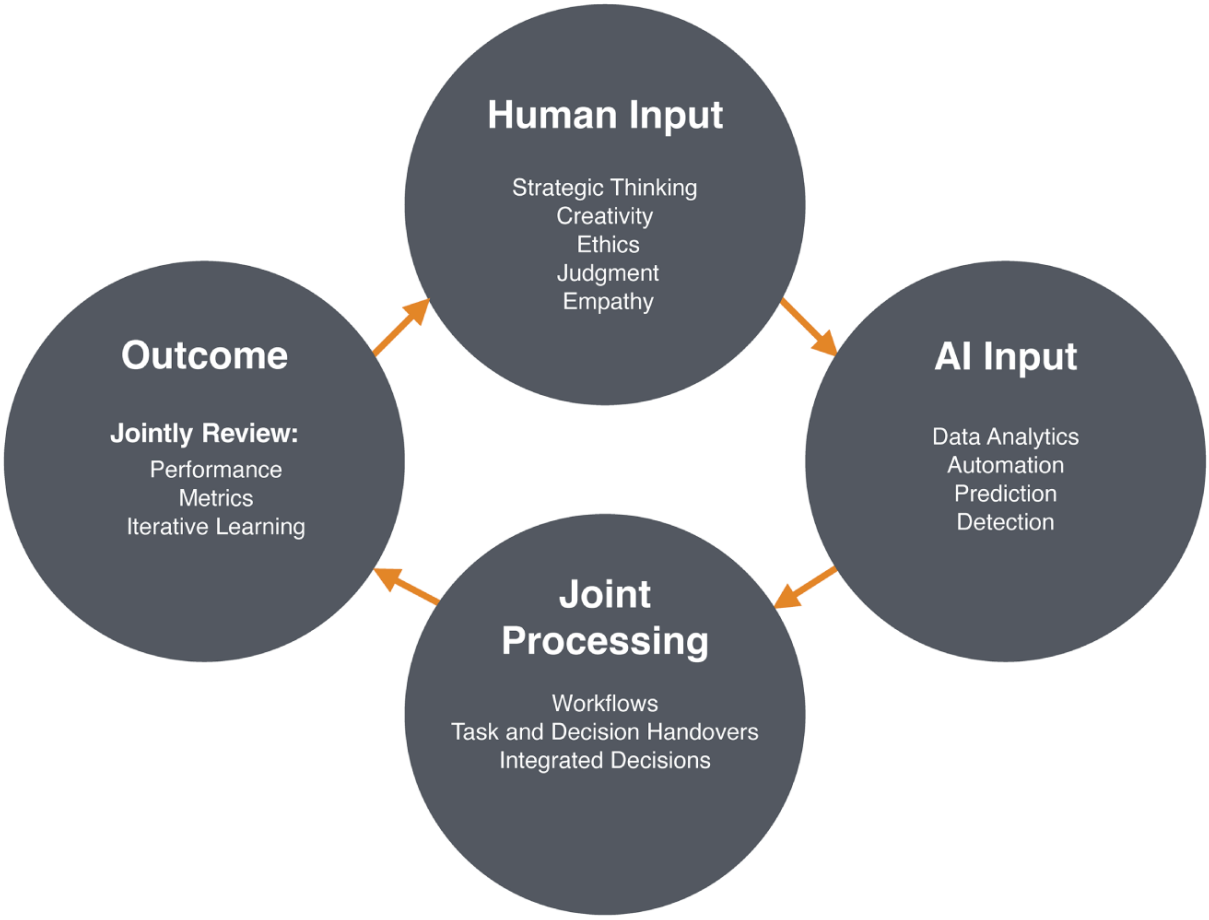


Figure 10.9 Collective Intelligence Framework: designing human-AI collaboration

Step 1: Populate

Start any new initiative by collaboratively populating each quadrant explicitly assigning tasks, roles, and decision boundaries.

Step 2: Align

Ensure all stakeholders clearly understand how human and AI contributions differ, overlap, and complement each other.

Step 3: Implement and Iterate

Continually reference the framework throughout the project lifecycle, adjusting based on emerging feedback, performance data, and evolving human-AI dynamics.

Your Executive AI Leadership Roadmap

Leadership is no longer about managing incremental changes but boldly designing the future. Throughout this book, we've explored how AI isn't merely a technological tool but a profound strategic force capable of reshaping markets, business models, and organizational capabilities. Now, it's time to transform insight into decisive action. Your challenge and opportunity is to harness the full potential of AI by consciously building an integrated strategic and operational foundation—what we have defined as the AI Leadership Operating System, the AI Organizational Operating System, and the Collective Intelligence Framework (see Figure 10.10).

- **AI Leadership Operating System.** Leadership in an AI-driven world cannot be passive or reactive; it must be proactive, strategic, and ethically grounded. Clarify precisely where AI automates, augments human decision-making, and, critically, where it should never lead. Empower your leadership teams to embrace AI with transparent decision architectures, disciplined governance rituals, and structured board-level dialogues that ensure ethical alignment and strategic coherence. We believe this operating system provides the clarity and confidence to harness AI wisely and strategically.

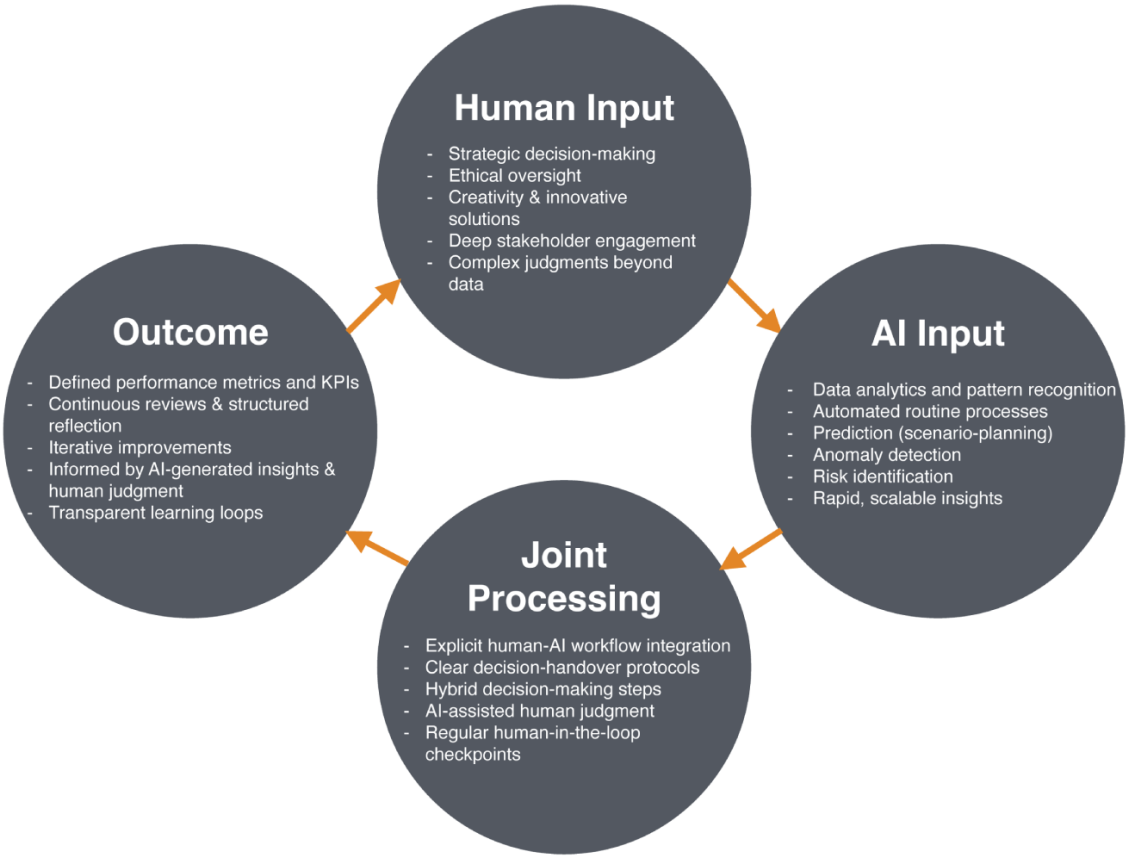


Figure 10.10 Collective Intelligence Framework: example

- **AI Organizational Operating System.** Reimagine traditional organizational structures to become agile, cross-functional, and profoundly collaborative. Flatten hierarchies, dismantle silos, democratize data, and encourage rapid, iterative experimentation. By investing in robust infrastructure and flexible, skills-based talent strategies, your organization becomes a dynamic platform ready to scale AI-driven innovations swiftly and seamlessly. This operational agility doesn't just future-proof your company; it positions you as an industry leader.
- **Collective Intelligence Framework.** This strategically integrates human creativity, intuition, and ethical judgment with AI's analytical precision, predictive capability, and scalability. Use the Collective Intelligence Framework as your practical tool to define and optimize these interactions explicitly, ensuring humans and AI collaborate productively, ethically, and

sustainably. Embracing collective intelligence means transcending mere efficiency, unlocking extraordinary innovation, more profound organizational resilience and sustained strategic advantage.

Your commitment as a C-suite leader will determine the difference between incremental improvements and transformative leaps forward. These three pillars of the AI Leadership Operating System, AI Organizational Operating System, and the Collective Intelligence Framework provide the structure and control to minimize your risks and maximize your ability to lead a truly transformative company. So now, it's your turn. AI has profound possibilities as an engine of genuine, transformative possibilities. Now, it's up to Team Human.

Notes

1. Jess Weatherbed, "Spotify Finally Turned a Profit for a Full Year," The Verge, February 4, 2025, https://www.theverge.com/news/605709/spotify-earnings-q4-2024-profitability.

2. Julia Stoll, "Netflix's Net Income from 2000 to 2024," Statista, https://www.statista.com/statistics/272561/netflix-net-income.

3. Gregory J. Pirchalski and Nicole M. Herndon, "Tuning Corporate Governance for AI Adoption," NACD, December 11, 2024.

4. Lara Abrash, Karen Edelman, and Arno Probst, "Governance of AI: A Critical Imperative for Today's Boards," Harvard Law School Forum on Corporate Governance, October 23, 2024.

5. Alex Singla, Lareina Yee, and Alex Sukharevsky, "The State of AI: How Organizations Are Rewiring to Capture Value," McKinsey & Company, March 12, 2025.

6. Hannah Mayer, Lareina Yee, Michael Chui, and Roger Roberts, Superagency in the Workplace: Empowering People to Unlock AI's Full Potential (McKinsey & Company, January 28, 2025) .

7. Mayer et al., *Superagency in the Workplace*.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Chapter 11
AI in the Next Five Years: Navigating Strategic Uncertainty

> *"In the business world, the rearview mirror is always clearer than the windshield."*
>
> —Warren Buffett

Remember Shereshevsky? The man with perfect memory and no comprehension. He stored everything and understood nothing. His mind was, in effect, a hard drive with no operating system: infinite data, no algorithm to make sense of it. For decades, companies operated in his image. AI has changed that. The question this book has been answering is: what does the leader who has finally integrated the operating system look like?

Meet Morgan Chen.

Morgan wakes at 5:30 a.m., groggy but quickly clearing her mind with a short meditation session. Her personal AI assistant has already analyzed her sleep patterns via her smartwatch and adjusted her morning schedule, suggesting an extra five minutes of meditation based on slightly elevated stress markers.

"Good morning, Morgan," her ambient AI assistant greets her. "Your executive dashboard is ready with priority items. Board meeting preparations are 92% complete, with three decision points flagged for your review."

By 7 a.m., Morgan is at her home office, reviewing an AI-generated briefing that has distilled thousands of data points into actionable insights. "Increase resource allocation to the Shenzhen team by 25%," she decides, after studying the recommendation and adding her own contextual knowledge about the team's recent leadership changes, something the AI flagged but could not fully evaluate without her input. Her decision triggers workflows across multiple AI systems, each requiring human oversight at critical junctures.

When a competitor announces an unexpected product at 8:30 a.m., Morgan activates the company's scenario planning AI, which generates five potential responses with detailed probability analyses. "Let's accelerate Q3, but not by the full six weeks the AI suggests," she determines after a video call with her leadership team. "The Singapore team needs more runway given their current challenges." The AI had not factored in the recent personal struggles of the Singapore director—something Morgan learned over dinner last month.

Throughout the day, Morgan moves between AI-augmented decision-making and deeply human interactions. She spends 30 minutes coaching a promising director whom the AI identified as a flight risk based on subtle changes in communication patterns. The AI suggested talking points. Morgan's empathy and shared experience form the core of the conversation.

By 2:00 p.m., Morgan leads the board meeting, where AI systems provide real-time data visualization and scenario modeling as questions arise. The board members look to Morgan for the contextual judgment, ethical considerations, and strategic vision that define the company.

At 9 p.m., Morgan reviews the day's learnings—captured by the company's AI systems and her own human observations. Tomorrow will blend AI efficiency with human judgment again. That partnership is not a feature of her leadership. It is the definition of it.

This is not science fiction. Everything Morgan does today is possible, but she does so with a mindset that recognizes that change has changed. Profoundly. What makes her exceptional is not her access to AI. It is her judgment about when to trust it, when to override it, and when to protect certain decisions from it entirely.

Shereshevsky had infinite memory and no wisdom. Morgan has both. She understood that the operating system was never AI alone. It was the human-AI partnership, governed with discipline, calibrated with care, and guided by the conviction that the machine's speed and the human's judgment are more powerful together than either is alone.

That is the argument this book has made from the first page. This final chapter asks where that argument leads. Not just for you or your organization, but for the next five years of the world you are about to lead.

Let's look over the horizon.

Limitless Information/Limited Knowledge: The AI Paradox

In 1941, Jorge Luis Borges wrote two short stories that offer a mirror to our world today. In "The Library of Babel," he envisioned an infinite library of hexagonal rooms, extending endlessly, containing every possible combination of letters, symbols, and spaces. Within its vast shelves lie every truth and every falsehood, prophecies and

contradictions equally present, librarians wandering eternally, driven mad by the search for meaningful order amid infinite chaos. In another tale, "The Lottery in Babylon," Borges depicts a society governed entirely by chance. In this world, every aspect of life, whether status, fortune, or morality, is dictated by an incomprehensible lottery based on ever-changing rules with unknown motivations. As a result, society embraces uncertainty and believes randomness is fairer than deliberate intent.

Together, these stories illuminate the paradox leaders face in the age of artificial intelligence. In the first world, confronted by limitless data, overwhelmed by possibilities yet starved of clarity, we seek the seductive illusion of predictable outcomes in Borges' infinite library. At the same time, we exist in the Babylonian lottery world. Here, we're subjected to the whims of algorithms whose inner workings remain opaque, uncertain whether our fate is determined by deliberate logic or capricious chance. Navigating this dual reality demands a leadership comfortable with complexity, skeptical of easy certainties, and humble enough to accept that absolute knowledge may remain forever beyond reach in the infinite corridors of AI-driven futures.

Throughout this book, we've shown that AI fundamentally transforms decision-making, organizational structures, competitive dynamics, and leadership imperatives. The key insight is that in the AI age, strategic clarity depends on accepting ambiguity and building the ability to navigate it. AI's true distinction lies in its dual role as both a tool and a collaborator. Unlike past technologies that merely executed commands, modern AI systems learn, adapt, and co-create. This collaboration between human and machine intelligence redefines organizational capability and executive decision-making.

The C-suite executive who thrives in this environment will embrace strategic uncertainty, not as an obstacle to be overcome but as the very medium through which innovation and competitive advantage now flow. To do so, executives should foster a mindset adaptable to rapid shifts, anticipate long-term transformations, and encourage teams to explore new solutions. As strategist Rita McGrath observed, "The difficulty with such an emphasis on facts is that, unfortunately, facts are often a lagging indicator of what could potentially be

important. By the time you are dealing with a fact on the ground, whatever led to it has already happened."[1] This has never been truer than in today's and tomorrow's AI-driven business universe.

Nevertheless, we must continue to adapt to AI while building future models that co-create the AI future we desire as a business and as a society, one that maximizes the upsides and minimizes the downsides. The future is impossible to predict precisely, yet understanding the trajectory of artificial intelligence and related technologies is indispensable for today's business leaders. AI has become an essential force reshaping business and society, bringing unprecedented opportunities and significant risks. Strategic clarity in the age of AI will not emerge from eliminating uncertainty but rather from becoming adept at navigating it. The C-suite must embrace strategic agility, preparedness, and continual learning as cornerstones of leadership that will win the future.

The Next Five Years: Signal, Not Noise

The generative AI boom triggered by ChatGPT's launch in late 2022 has produced genuine capability and inflated expectations in equal measure. For executives seeking competitive advantage, neither the techno-utopian nor the defensive doomer position serves. What follows is neither. It is a clear-eyed account of where AI is heading and what that means for the organizations you lead.

Another reminder. This is the least capable AI we will ever witness.

The next five years will likely not be defined by a single dramatic breakthrough. This progress will be defined by convergence: multiple dimensions of AI capability maturing simultaneously, integrating into the infrastructure of business, and creating a new competitive baseline that will separate the prepared from the reactive.

From perception to reasoning. Current AI systems excel at pattern recognition. What is changing, faster than most analysts expected, is AI's ability to perform multistep reasoning, understand causality, and maintain logical consistency across complex problems. The inflection point has already been crossed, where AI's role shifted from merely providing data to actively helping to conceptualize

solutions. This is not a marginal improvement. Take note. It is the difference between a very fast research assistant and a thinking partner.

From isolated to integrated. We've already seen how AI embeds its capabilities directly into the core systems organizations already run on. That's your planning platforms, your customer systems, and your operational infrastructure. AI has stopped being a separate application and has become the nervous system of your existing operations.

From reactive to agentic. The shift from AI as a passive tool that responds to queries to AI agents that take initiative within defined parameters is already underway. We are done with AI as Google. Examples are visible in cybersecurity, customer service, and software development. Within a few years, agentic AI will be the default mode of AI deployment in most enterprise contexts. Chapter 9 gave you the governance architecture for this. The question is no longer whether agents will handle most routine work. It is whether your organization will be ready to govern them when they do.

The Economic Reality

The productivity data is no longer theoretical. Generative AI increased work throughput by an average of 66% in controlled trials.[2] Employees augmented by AI assistants experienced gains of 25% to 40% on knowledge-worker tasks; coding was completed 56% faster; writing quality improved by 37%; and customer service handled 20% more cases per hour. These are not marginal improvements. These are just the first step-changes in what human-machine collaboration can achieve.

The macroeconomic outlook confirms this. AI is projected to contribute between $15.7 trillion and $19.9 trillion to the global economy by 2030.[3] For each dollar invested in business-related AI solutions, about $4.60 in economic value is created.[4] The organizations that benefit most are not necessarily those with the most advanced models. They are the ones who redesigned their workflows to enable the models to function. Knowing when to turn the dial and when not to is key.

The Workforce Shift

The labor market will see significant churn. The World Economic Forum projects a net increase of 78 million jobs globally by 2030, driven by AI and related technologies. New roles for data scientists, AI system trainers, AI ethicists, and automation strategists. Many existing jobs will evolve and likely disappear, with AI handling specific tasks and elevating the importance of uniquely human capabilities: complex judgment, contextual communication, ethical discernment, and creative synthesis.

There is a risk of mass unemployment. There is already a hollowing out of the institutional wisdom that makes the automation trustworthy. The companies that get this right will not be the ones that replace the most people, but rather the ones that develop the most people alongside the technology.

Healthcare

AI promises transformative changes in healthcare, enhancing diagnostic accuracy and efficiency. AI-assisted breast cancer screening improved detection rates by 29%, significantly reducing clinician workloads.[5] Similarly, AI has outperformed human clinicians in stroke diagnosis, doubling accuracy in certain cases. AI's integration of multimodal data, including imaging and genomics, advances personalized medicine, which is critical given a projected shortage of 11 million healthcare workers by 2030.[6] Yet, patient safety and ethical concerns remain significant. AI errors could threaten lives, especially when used without human oversight. ECRI warns that AI can expand risks across diagnosis, scheduling, and treatment, potentially exacerbating disparities if trained on biased datasets.[7] Rigorous validation, continuous monitoring, and human oversight will be essential for the responsible deployment of AI.

The Societal Stakes

AI's impact on society will be as consequential as its impact on business, and the C-suite has more responsibility in shaping that impact than most executives currently acknowledge. There are no free riders in the private sector when it comes to AI. Your

deployment decisions contribute to a cumulative societal outcome, whether you intend them to or not.

The standard for C-suite leadership in this environment is stewardship, not compliance. The governance frameworks in this book are not just tools for competitive advantage. They are your contribution to whether AI becomes an empowering force or an accelerant of existing divides. That is not a question for regulators alone. It is a question for every executive who signs off on an AI deployment.

Five Bets for The Next Five Years

We close with five predictions for the next five years. These are not hedged forecasts. They are best informed by the frameworks in this book, the patterns we have observed across industries, and the conviction that most AI commentary mistakes the technology story for the business story. We may be wrong on the timing. We do not believe we are wrong on the direction.

The New 80/20 Rule

For most of the modern era, competitive advantage was distributed unevenly but survivably. Middling companies could persist for decades on inertia, brand recognition, and adequate execution. AI changes that arithmetic. Within five years, we predict that 80% of new market value in most industries will be captured by 20% of companies, the ones that mastered governance, built centaurs, and turned the dial while competitors were still debating whether to adopt.

The remaining 80% of companies will not necessarily disappear. But they will find themselves in a permanent Red Ocean, competing on price for commoditized AI-augmented services while the top quintile creates and captures new markets. The gap between AI-literate organizations and AI-adopting organizations, between those who redesigned the factory around electricity and those who bolted a motor onto a steam-age layout, will become the defining competitive divide of the decade.

Data as the New Balance Sheet

For years, we have heard that data is the new oil. That was directionally right but strategically vague. The real shift is more specific: within five years, companies will be valued not just for their earnings but for the quality, exclusivity, and strategic utility of their institutional knowledge. When every competitor has access to the same foundation models, the differentiator becomes the proprietary data that those models are trained on: customer histories, operational failure libraries, supply chain patterns, and institutional wisdom accumulated over decades.

The companies that organized this knowledge into assets that their AI can learn from will compound their advantage with every cycle. Those who treated data as a byproduct of operations rather than a strategic asset will discover, too late, that their AI is only as smart as their data infrastructure allows it to be. Expect to see data quality and proprietary knowledge assets appear on analyst reports and acquisition term sheets alongside EBITDA within the decade.

The Last Generation of Human-Only Managers

Here is a demographic reality that almost no one in the AI conversation is discussing: Generation X will be the last cohort of leaders whose entire management training was built around leading humans alone. The generation now entering the workforce will manage humans, AI agents, and physical robots simultaneously, from their first day.

This is not a technology prediction. It is an organizational one. The skills required to orchestrate a mixed workforce of people and machines, setting decision rights, calibrating trust, knowing when to override and when to defer, are fundamentally different from the skills taught in every MBA program and leadership development course for the past 50 years. The companies that recognize this shift and redesign their talent development accordingly will build the leadership pipeline for the next era. The ones that keep training managers as if their only direct reports will be human will find

themselves with leaders who are fluent in half the language their organization speaks.

The Return to On-Premises and Open Source

This is our most contrarian prediction, and we hold it with conviction. The current AI narrative assumes that cloud-based foundation models from a handful of providers will dominate enterprise AI indefinitely. We believe the opposite is beginning: a structural migration toward on-premises deployment and open-source models as the dominant enterprise AI architecture.

The reasons are converging from multiple directions. Data sovereignty regulations are tightening globally. The cost of cloud-based AI at scale is becoming prohibitive for many organizations. And the DeepSeek moment of 2025 demonstrated that smaller, specialized models can match frontier performance at a fraction of the cost. When a Chinese startup matches OpenAI's output quality for pennies on the dollar, the argument that you need the largest model from the largest provider collapses.

Enterprises that can run capable models on their own infrastructure, trained on their own proprietary data, achieve something the cloud model can never fully offer: complete control over their intelligence. Expect AI sovereignty to become as important to boardrooms as data sovereignty was in the previous decade.

A New Age of Enlightenment

Most AI forecasts focus on what will be destroyed: jobs eliminated, industries disrupted, companies made obsolete. We see something else coming, a hyper-acceleration of entrepreneurship unlike anything since the Internet's first decade.

When the cost of building software approaches zero, when a single founder with an AI agent can do the work that once required a team of 20, when the barriers to creating a viable product collapse from hundreds of thousands of dollars to hundreds, you do not get consolidation. You get an explosion of creation. We predict a renaissance of small, lean, AI-native companies that compete not on

scale but on taste, judgment, and the quality of their human-AI architecture. The centaur entrepreneur, moving at machine speed with human intuition, will be the defining figure of the next decade.

Just as the printing press did not merely make existing books cheaper but created entirely new forms of knowledge and expression, AI will not merely automate existing businesses. It will make entirely new businesses possible that were previously too small to be viable, too complex to be manageable, or too niche to attract capital. The democratization of capability is the most under-discussed consequence of the AI revolution, and it may ultimately be the most important.

One Last Bet

If we are going to ask readers to place bets, we should place one ourselves.

By 2030, the most valuable company in at least three major industries will not be the one with the most AI agents, the most data, or the largest R&D budget. It will be the company with the best human-AI governance architecture, the one that figured out which decisions to automate, which to augment, and which to protect from automation entirely. Competitive advantage will belong to the organizations that drew the sharpest lines, not the ones that drew the fewest.

That prediction is specific enough to be wrong. We make it anyway, because everything in this book points in the same direction: the technology is a commodity, the governance is the moat.

In the first chapter, we met a man with perfect memory and no comprehension. He had a mind that could store everything and understand nothing. The centaur is the operating system he never had.

The Centaur's Return

We met Gulliver earlier on the island of the Houyhnhnms, where he fell in love with the horses' pure reason and could no longer stomach

the sight of his own kind. Swift meant this as a warning. Gulliver's tragedy was not that he admired the horses but that he saw only two options: be a rational machine or be an irrational human. He could not imagine a creature that was both.

This book has argued for that third way. The centaur's half-human judgment, half-machine capability is governed by the frameworks we've shared, which evolve as the relationship matures. Not the Yahoo (the brutish humans in *Gulliver's Travels*), stumbling through a world it cannot process fast enough. Not the Houyhnhnm, optimizing with elegant precision toward goals no one examined closely enough. The Centaur. The one Swift never wrote because, in his century, it was only a myth.

It is no longer a legend. The tools exist. The frameworks we presented are in your hands.

Gulliver came home and found he could not live among humans. He had mistaken the horse's certainty for wisdom. Don't make his mistake. The wisdom was never in the machine's knowledge and speed or the human's instinct alone. It was always a matter of tension between the two. You have your hand on the dial, turning it one careful notch at a time.

Notes

1. Rita McGrath, Seeing Around Corners: How to Spot Inflection Points in Business Before They Happen (Boston: Houghton Mifflin Harcourt, 2019) .

2. Jakob Nielsen, "AI Improves Employee Productivity by 66%," Nielsen Norman Group, July 16, 2023.

3. Lapo Fioretti, Carla La Croce, Andrea Siviero, and Elisabeth Clemmons, The Global Impact of Artificial Intelligence on the Economy and Jobs: AI Will Steer 3.5% of GDP in 2030 (Needham, MA: IDC, 2024) .

4. CDO Magazine, "AI to Contribute $19.9 Trillion to Global Economy through 2030 — IDC Research," October 9, 2024.

5. Vendela Hernström et al., “Screening Performance and Characteristics of Breast Cancer Detected in the Mammography Screening with Artificial Intelligence Trial (MASAI),” The Lancet Digital Health 7, no. 3 (2025) .

6. Madeleine North, “7 Ways AI Is Transforming Healthcare,” World Economic Forum, August 13, 2025.

7. ECRI, “Artificial Intelligence Tops 2025 Health Technology Hazards List,” December 4, 2024.

About the Authors

Rajeev Kapur is the CEO of 1105 Media, a Forbes-recognized "Top Leader Bringing AI to Everyone," and the bestselling author of *Chase Greatness*, *AI Made Simple*, and *Prompting Made Simple*. During his 30-year career, he built a billion-dollar business at Dell and earned recognition as an Ernst & Young Entrepreneur of the Year finalist, a YPO Best of the Best Award winner and one of the *Los Angeles Business Journal*'s Leading Disruptors. He received his MBA from The USC Marshall School of Business where he was awarded The Order of Areté.

Jason Hunter Over a 20-year career, he has built entrepreneurship and economic-development programs across four continents and advised companies including L'Oréal, Toshiba, General Electric, and SAP. He founded Corners Group, is a senior global consultant for the Blue Ocean Strategy Network, and holds advisory roles at the London School of Economics, INSEAD, and the University of Johannesburg.

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Index

Numerics

2 × 2 matrix, 87, 96–99

70-20-10 principle, 6

80/20 rule, 257

A

accelerated learning, 120

accountability, 5, 12, 18, 29, 32, 39, 47, 69, 127, 132, 135, 181, 204, 224, 227, 231, 233, 244

  bias, fairness, and, 56–59

  and control mechanisms, 48–51

adaptive nervous system, 177–178

Adaptive Risk Intelligence Architecture (ARIA), 126–128, 135–139

adaptive simulation, 109, 112–114, 126

Adapt stage, HorizonAI Framework, 89, 92, 95, 100–104

agency, 44, 69, 85, 235–236

agentic AI, 21, 196, 200, 202, 203, 213, 254

The Agentic Decision Framework, 201, 206–211, 215

The Agent Workforce Audit, 201–206, 215

agile methodology, 82

agile “teams of teams,” 235

AI. *See* artificial intelligence (AI)

AI-adopting organizations, 231, 257

AI-assisted breast cancer screening, 255

AI-based machine learning algorithms, 156

AI-centric organizational design, 235
AI-driven processes
business, 231
customer segmentation model, 236
Eliminate-Reduce-Raise-Create (ERRC) Grid, 154–159
experimentation, 68–71
financial planning and analysis (FP&A) tools, 134
initiatives, 169
risk radar, 130–132
scenario modeling, 75–76, 102
sentiment analysis, 144
simulation, 159
strategy, xiv, 72
supply chain, 168
AI-driven scenario planning, 76, 87–88
HorizonAI Framework, 88–89
Adapt, 100–104
Attack, 92–94
Observe stage, 89–92
Orient, 94–99
strategic advantage, 102–103
AI-enhanced “As Is” Strategy Canvas, 143–147, 161
AI-generated text, 179
AI-integrated supply chain, 182
AI Leadership Operating System, 222, 223, 225, 231–232, 240, 246–247
AI Leadership OS Canvas, 230–231

AI-literate business, 232
AI-literate organizations, 233–235, 257
  agency, 235–236
  AI tech support, 237–238
  Center for Strategic Prompting, 238
  cross-functionality, 235
  culture, 237
  roles, 236–237
AI-native companies, 259
AI-optimized organization, 236
AI Organizational Operating Systems, 222, 232, 234, 238–240, 247
AI-powered agents, 160
AI-powered precision agriculture systems, 78
AI-powered risk intelligence center, 135
AI-powered strategy tools
  “As Is” Strategy Canvas, 143–147, 161
  Buyer Experience Audit (BEA), 147–150
  Eliminate-Reduce-Raise-Create (ERRC) Grid, 154–159
  Six Paths Framework, 150–154
  “To Be” Strategy Canvas, 159–161
Airbnb, 23, 24
algorithmic heart, 175–177
AlphaGo, 65–66, 172
Altman, Sam, 42
Amazon, 13, 22, 43, 78, 94, 170, 173
  artificial intelligence (AI) coding agent, 197
  résumé-screening failure, 208

Amazon Go, 78
Amazon Web Services (AWS), 94, 197
AMD, 10
Annex Wealth Management, 44
anticipatory operations, 118
Apple, 10, 13, 99, 170
Apple Card, 53
ARIA. *See* Adaptive Risk Intelligence Architecture (ARIA)
Aristotle, 150
ARPANET, 44

artificial intelligence (AI), xiv, 33, 41, 42, 160, 183–162. *See also related entries*

agent, 195, 208, 215

assessment, 136

auditability, 53–55

and Blue Ocean Framework (BBO), 141–142

board engagement with, 224–229

bolting, 33, 183

China, 42, 43

circulatory system of, 174–175

context of, 223

control of, 43–44

corporate strategy, 65–66, 81–86, 107–108
adoption, 9–10
AI-driven scenario planning, 87–88
AI-driven strategy, 72
control paradigms, 85–86
data and information, 4–5
data flood and disruption, 66–67
experimentation, 68–71
frameworks, 67–68, 70
HealthCheckAI (*see* HealthCheckAI)
Healx, 11–12
HorizonAI (*see* HorizonAI Framework)
inside-out strategy, 76–78
military mindset vs AI mindset, 5–7
Nokia, 10
organizational excellence, 72–74
outside-in strategy, 74–76
strategic advantages, 102–103
strategic balance, art of, 79
strategy/innovation matrix, 83–84
ethics, 231
executive order, 41, 44–45
forecasts, 259
framing, 230
governance, 231
handover, 231
history of, 172

and human strategy, 130–132
implementation of, 226
initiatives, 232
integration of, 230, 256
language models, 186
leadership architecture, 230
literacy and governance, 229
literacy training, 169–170
as living system, 173–174
ordering system, 48–49
oversight, 231
paradox, 251–253
perspective models of
    computer vision, 179–180
    deep learning, 178–179
    generative AI, 180
    machine learning, 177–178
    natural language processing (NLP), 179
piloting, 230
power, 220
principles, 230
risk heat map, 228, 229
roles, 223
strategic insight board, 227, 228
strategy, 225
systems, 190
technology backbone for, 239

tech support, 237–238
"To Be" Strategy Canvas, 159–161
traffic management systems, 74
transparency, 53–55
upskilling, 231
Artificial Investment Associate (AIA), 22, 23
art of strategic balance, 79
"As Is" Strategy Canvas, 143–147, 161
Attack stage, HorizonAI Framework, 92–94
augmentation zone, 224
augmented perception, 120
automation zone, 223–224
AWS. *See* Amazon Web Services (AWS)

B

B2B transaction, 147
balanced scorecard, 66, 67, 82, 107
balanced self-regulating system, 184
Ballmer, Steve, 10
BCG Growth-Share Matrix, 82
BEA. *See* Buyer Experience Audit (BEA)
Behavox, 128
Beyond Blue Ocean (BBO) Framework, 141–142, 160, 161
Bezos, Jeff, 94
bias data argument, 56–57
Big data, 7
"Big-S" strategy, 84
binge-watching model, 221

biocentrism, 173
biological systems, 181, 183–184
black box(es), 29, 31, 34, 125, 133, 138
  models, 132
  recorder, 54
BlackRock, 22, 43
Blue Ocean Strategy, 12, 66, 67, 77, 83, 86, 142, 143, 145–147, 154, 159
Bodell, Lisa, 93
Borges, Jorge Luis
  “The Library of Babel,” 251
  “The Lottery in Babylon,” 251
Boston Consulting Group, 109
bottom-up approaches, 84–85
British Standards Institution, 46
Brown, David, 11
business models, 8, 71, 77, 86, 93, 95, 171, 221, 246
business strategy, 19, 67, 70, 71, 75, 83
Buyer Experience Audit (BEA), 147–150, 161, 162

C

capability evolution, 116
capability trap, 201
Carlson’s Law, 84, 85
centaur, 82, 193–194, 215, 257, 259–261
Center for Strategic Prompting, 238
CFOs. *See* chief financial officers (CFOs)
change management, 109, 181

chatbot, 6, 42, 179, 185, 194, 195
ChatGPT, 7, 9, 18, 47–49, 172, 186, 194, 253
Chen, Morgan, 250–251
Chesky, Brian, 23
chief executive officers (CEOs), 19, 23, 25, 39, 141, 145, 155, 169, 170, 198, 219, 236
chief financial officers (CFOs)
  and Adaptive Risk Intelligence Architecture (ARIA), 126–128, 135–139
  challenges, finance and risk, 124–126
  external risk, 130–131
  human governance and controls, 132–133
chief innovation officers (CIOs), 20, 34, 36, 133
chief marketing officers (CMOs), 19
chief operating officers (COOs), 20, 37
chief revenue officers (CROs)
  and Adaptive Risk Intelligence Architecture (ARIA), 135–139
  external risk, 130–131
  human governance and controls, 132–133
chief risk officers (CROs)
  and Adaptive Risk Intelligence Architecture (ARIA), 126–127
  challenges, finance and risk, 124–126
chief supply chain officers (CSCOs), 20, 37
chief technical officers (CTOs), 20, 34, 36
China, 42, 43
circulatory system, artificial intelligence (AI), 174–175
City Brain project, 74

Clark, Jack, 45
classic corporate hierarchy, 233
Claude, 9, 47, 186
Claude Mythos, 13
Cleveland Clinic, 118
cloud-based foundation models, 258
cognitive amplifier function, 68, 71
Collective Intelligence Framework, 247–248
  core components of, 242–244
  human-AI collaboration, 240–241
  insights to implementation, 241
commoditized AI-augmented services, 257
communication, 18, 20, 30, 31, 110, 111, 125, 126, 128, 129, 133, 135, 136, 185, 189, 233, 250, 255
  channels, 237
  interface, 179
competitive intelligence, 26, 141, 142, 148
competitive weapon function, 68
compliance adaptation, artificial intelligence (AI), 131
component simulation, 113
computer vision, 78, 179–181
conductors, 200–201
continuous adaptation, 183
continuous experimentation, 71, 118–119
continuous reconfiguration, 109, 114–117, 126
Cook, Tim, 170

corporate strategy
adoption, 9–10
AI-driven scenario planning, 87–88
AI-driven strategy, 72
control paradigms, 85–86
data and information, 4–5
data flood and disruption, 66–67
experimentation, 68–71
frameworks, 67–68, 70
HealthCheckAI (*see* HealthCheckAI)
Healx, 11–12
HorizonAI (*see* HorizonAI Framework)
inside-out strategy, 76–78
military mindset vs AI mindset, 5–7
Nokia, 10
organizational excellence, 72–74
outside-in strategy, 74–76
strategic advantages, 102–103
strategic balance, art of, 79
strategy/innovation matrix, 83–84
COVID-19, 10, 102, 103, 118
critical compliance decision, artificial intelligence (AI), 129
cross-functional collaboration, 10, 238–240
cross-functionality, 235
crowd-solving platforms, 237
C-suite executives/leadership, 10, 12–13, 18, 19, 35, 168–170, 186–190, 214, 222, 226, 227, 231, 238–239, 244, 248, 252, 256

culture, 4, 19, 25, 47, 59, 86, 94, 101, 108, 127, 133, 138, 154, 156, 167, 190, 237, 245
customer service, 8, 50, 107, 168, 196, 198, 200, 213, 214, 223, 235, 243, 254

D

Dalio, Ray, 23
data, as new balance sheet, 257–258
data argument, 56–57
data center cooling systems, 74
data-driven content creation process, 221
data-driven culture, 156
data integrity, 129–130
data sovereignty regulations, 258
DB Schenker, 17, 18
decentralization, 235
decision architect function, 68
decision-making, strategic, 67–71
deep learning, 30, 74, 167, 172, 178–179, 189
DeepSeek, 43, 258
Deloitte Global Boardroom Program survey, 224–225
deployment strategy, artificial intelligence (AI), 137–138
Design Thinking, 77, 83
Digital Flywheel strategy, 181
digital supply chain, 103
digital twin models, 76
Digital Twins for Operations, 114
digital workforce, 202

Dimon, Jamie, 169
Dingo, 162
"As Is" Strategy Canvas, 145–147
Buyer Experience Audit (BEA), 149–150
Eliminate-Reduce-Raise-Create (ERRC) Grid, 156–159
Six Paths Framework, 152, 154
"To Be" Strategy Canvas, 160
disruption, 38, 103, 108, 116, 119, 123, 137, 187, 221, 222
artificial intelligence (AI), 66–67, 96
notion of, 220
supply chain, 21, 37, 38, 81, 87, 131, 187
distributed sensing, 183
Drucker, Peter, 55
Duke University Hospital, 112
dynamic optimization, 73
dynamic risk assessments, 28, 31

E

economic reality, 254–255
Eisenhower, D. Dwight, 98
Eliminate-Reduce-Raise-Create (ERRC) Grid, 154–159, 161, 162
empowerment, 235, 239
enterprise resource planning (ERP) systems, 201
enterprise risk management (ERM), 125, 133
ERRC Grid. *See* Eliminate-Reduce-Raise-Create (ERRC) Grid
ethical governance, 181

ethics model, 41–43
accountability
bias, fairness and, 56–59
and control mechanisms, 48–51
frontier AI models, 44–45
ideas and action, 47–48
risk management and safety, 51–53
self-regulation, 45–47
transparency and auditability, 53–55
trust and control, 43–44
eval awareness, 42
evolutionary intelligence, 109
evolutionary progression, 47
evolution simulation, 113
executive artificial intelligence (AI) leadership roadmap, 246–248
exogenous vs endogenous approaches, 85
experimental operations, 118
experimentation, strategic artificial intelligence (AI), 68–71
explainable AI (xAI), 43, 133
external risks, AI-driven
external threats, 131–132
financial risk analytics, 130–131
monitoring and compliance, 131

F

Fable 5 and Mythos 5 models, 45
fairness, artificial intelligence (AI) algorithms, 56–59
Fan Hui, 65

financial controls, 129–130
Financial Industry Regulatory Authority (FINRA), 46
financial planning and analysis (FP&A) tools, 134
financial risk analytics, 130–131
financial services, and HorizonAI, 97
FINRA. *See* Financial Industry Regulatory Authority (FINRA)
Ford, Henry, 166, 200
fragile X syndrome, 11
fraud detection system, 182
Friedman, Milton, 45
frontier AI models, 44–45

G

Galloway, Scott, 141
Gartner, 196, 202
GE Healthcare, 112
generative artificial intelligence (AI), 9, 19, 49, 93, 98, 102, 130, 171, 180, 181, 253, 254
GitHub Copilot, 9
global cloud markets, 94
global political heartbeat, capturing, 75–76
global supply chains, 76, 103
Goldman Sachs, 22, 53
Google, 10, 13, 43, 65, 68, 74, 173, 254
Google Cloud, 28, 94
Google DeepMind, 65–66, 74
  AlphaGo, 65, 172
governance architecture, 200, 254, 260

governance frameworks, 41–43
accountability
bias, fairness and, 56–59
and control mechanisms, 48–51
frontier artificial intelligence (AI) models, 44–45
ideas and action, 47–48
risk management and safety, 51–53
self-regulation, 45–47
transparency and auditability, 53–55
trust and control, 43–44
GPUs, 7, 172
growth creation framework, 86
Growth Drivers, 19–20, 23–24, 39, 50
mindset, 24–25
playbook, 25–28
Guilliams, Tim, 11
Gulati, Ranjay, 77
Gulliver's Travels (Swift), 193–194, 260

H

Hays Code era (1934–1968), 45
HBO's *Silicon Valley,* 197
Healnet, 11
healthcare, 56, 57, 112, 145, 255–256

HealthCheckAI, 86
and Adaptive Risk Intelligence Architecture (ARIA) Framework, 126–128, 135–139
advantages, 123–124
AI-driven operational excellence, 124
external risk radar, 130
external threats, 131–132
market and financial risk analytics, 130–131
regulatory monitoring and compliance adaptation, 131
finance challenges and risk executives, 124–126
four domains, implementation
continuous experimentation, 118–119
human-AI collaboration, 119–121
proactive strategy, 117–118
redundancy over perfection, 119
human governance and controls
data controls and security, 133
enterprise risk management (ERM) framework, 125, 133
explainability and transparency, 132–133
rules for artificial intelligence (AI) decisions, 132
internal risk defenses, 128–129
data integrity and financial controls, 129–130
internal communication and conduct, 129
regulatory compliance, 129
phased sequence, practice, 121–122
risk management, 134–135

three organs of, 109–110
adaptive simulation, 112–114
continuous reconfiguration, 114–117
omnidirectional sensing, 110–112
Healx, 11–12
Hern, Alex, 45
Higgins, Paul, 145, 162
hiring algorithm, 54, 58
Hitachi, 76
Hollywood evolution, 45–46
HorizonAI Framework, AI-driven scenario planning, 88
2 × 2 matrix, 96–97
Adapt, 100–104
Attack, 92–94
financial services, 97
Observe stage, 89–92
Orient, 94–99
strategic advantage, 102–103
HSBC, 28–29
HSBC Celent, 28
human-AI
architecture, 259
collaboration, 119–121, 138, 240–241, 245
governance architecture, 260
hybrid teams, 238
interaction, 184
human brain, 5, 67

human governance and controls, 132–133
human orchestration, 200
Human Premium skills, 13, 198–199, 201
human strategy, and artificial intelligence (AI), 130–132

I

IBM, 69, 98, 166, 194
idea meritocracy, 23
Immelt, Jeff, 169
Industrial Revolution, 165
Innovation Catalyst, 20, 33–34, 39, 50
  mindset, 34
  playbook, 35
  scaling, 35–36
innovation intelligence networks, 35
innovation multiplication effect, 35
inside-out strategy, 72, 76–79
  anticipation, 78
  insight mining, 78
  self-learning models, 78
institutional knowledge, 200, 216, 257
integrated hybrid team structures, 242–243
integrated intelligence architecture, 181
Intel, 10
intelligence center, 112, 135
intelligent nervous system, 110
interlocking tools, 201–202
internal communication and conduct, 129

internal risk defenses, 128–129
data integrity and financial controls, 129–130
internal communication and conduct, 129
regulatory compliance, 129
Internet, 5, 18, 41, 44, 46, 73, 76, 167, 259
iPhone, 10
iterative learning loops, 243

J

Jacobsen, Brian, 43–44
Jensen, Greg, 22
job descriptions, 233–234
Jobs, Steve, 166, 219
Jobs-to-be-Done, 77, 83
John Deere, 78
Journal of Financial Economics, 56
JPMorgan Chase, 10, 54, 169, 186

K

Kahneman, Daniel, 67–68
Kanioura, Athina, 21
Kasparov, Garry, 194
Kill the Company, 93–94
Kim, W. Chan, 142, 143
knowledge gap, 170, 190
Kooda Systems, 213–215
Kriegsspiel, 81–82
Krizhevsky, Alex, 7

Krugman, Paul, 18

L

labor market, 255

Lancet Digital Health, 120

large language models (LLMs), 7, 43, 90, 102, 109, 168, 184

Lauder, Evelyn H., 120

lean manufacturing, 119

Lean Startup, 83

learning matter for executives, 176–177

legacy decision-making processes, 234

Lehane, Chris, 41

Lendingkart, 57

The Library of Babel (Borges), 251

limitless information/limited knowledge, 251–253

living system, 173–174, 181–184

local adaptation, with global coherence, 183

long-term benefits, 245–246

L'Oréal, 91, 170

The Lottery in Babylon (Borges), 251

M

machine intelligence, 22, 85, 142, 161, 252

machine learning (ML), 73, 82, 130–132, 134, 155, 156, 169, 172, 177–178

Manhattan Project, 44

market risk analytics, 124, 130–131, 137

markets strategy, 76–78

maturity stages, 203–204

Mauborgne, Renée, 142, 143
McGrath, Rita, 253
McKinsey, 131, 186
  reports, 5
  survey, 169, 186
MD Anderson Cancer Center, 69
mediocre algorithm, 174
Memorial Sloan Kettering (MSK), 69, 120, 121
mental models, 67–68, 99, 168
Merrill, William Henry Jr., 46
MGX, 43
Microsoft, 5, 10, 13, 43, 94, 99, 114, 169, 185
Microsoft Azure, 94, 114
military hierarchies, 6
military mindset vs AI mindset, 5–7
MIT Sloan Management Review study, 190
MJ Rathbun, 197
ML. *See* machine learning (ML)
modern artificial intelligence (AI), 4, 7, 26, 31, 90, 173
  renaissance, 172
  systems, 4, 31, 90, 185, 252
monday.com, 196
Morgan Stanley, 23
MSK. *See* Memorial Sloan Kettering (MSK)
Munger, Charlie, 42

N

Nadella, Satya, 10, 169
natural language processing (NLP), 90, 129, 131, 135, 179
Netflix, 186, 220, 221
network intelligence, 74
new age of enlightenment, 259
New York Times, 29
NLP. *See* natural language processing (NLP)
Nokia, 10, 99
Nooyi, Indra, 219
Norton, Larry, 120
notch, 211, 212, 214–216, 261
Novo Nordisk, 78
Nvidia, 7, 21, 43

O

Observe, Attack, Orient, and Adapt (OAOA), 88
Observe stage, HorizonAI Framework, 89–92
off-limits zone, 224
omnidirectional sensing, 109–112, 121, 126
On-Road Integrated Optimization and Navigation (ORION), 36, 37
OpenAI, 7, 41–43, 46, 49, 172, 186, 259
operational evolvability, 114–117, 123
Operational Excellence Champion, 20, 36–39, 50
  mindset, 37
  playbook, 38
  scaling, 38–39

operational excellence strategy, 72, 79
network intelligence, 74
predictive intelligence, 73–74
systemic learning, 74
operational perfection, HealthCheckAI, 119
Oracle, 43
organizational excellence, 107–108. *See also* HealthCheckAI
AI-driven finance, 124
frameworks, 86
organizational imagination, 112–114
organizational nervous system, 110, 123
Orient stage, HorizonAI Framework, 94–99
outside-in strategy, 72, 74–77, 79, 85
Ozempic, 78, 141

P

parametric adaptation, 115
pattern sensing, 111
Peloton, 91–92
PepsiCo, 21
perceptual processing system, 178–179
Perez, Juan, 37
Persian Gulf war, 103

perspective models
computer vision, 179–180
deep learning, 178–179
generative artificial intelligence (AI), 180
machine learning, 177–178
natural language processing (NLP), 179
PESTEL factors, 95
phased sequence, HealthCheckAI in practice, 121–122
Pinker, Katja, 120
Porter's Five Forces, 66–67, 82
precision agriculture systems, 78
predictive intelligence, 73–74
predictive scenario planning models, 114
predictive sustainability, 33
Prince Wilhelm, 81
proactive strategy, 26, 83, 117–118
agents, 127
frameworks, 83
proactive vs reactive approaches, 85
product management, 182
profit proposition, 84
Project Glasswing, 13
project management platform, 196
prompt, 29, 76, 90, 98, 155
design, 186, 189, 238
engineering, 184–186, 190
problem, 184–189

proprietary data, 18, 175, 200, 257, 259
proprietary turbine health database, 158
Prussian command-and-control architecture, 182
Python, 168

R

reactive strategy paradigms, 83
reactive vs proactive approaches, 85
Recursive Self-Improvement, 45
RegTech, 97, 135
reinforcement learning, 176
relationship sensing, 111
remote monitoring model, Dingo, 145
return on investment (ROI), 6, 7, 27, 37, 188
reversibility variable, 206
revolutionary change, 47
risk areas, 127, 129, 136–137
Risk Guardian, 20, 28–33, 39, 50, 124, 125
  challenges, 124
  mindset, 29–30
  playbook, 30–33
risk management
  and Adaptive Risk Intelligence Architecture (ARIA) Framework, 135–139
  AI tools and technologies in, 134–135
  challenges, 124–126
  and safety, 51–53
Royal Dutch Shell, 75

rules-based predictive analytics engine, 158
Rumsfeld, Donald, 171

S

SaaS-pocalypse, 196
Sachs, Goldman, 22, 53
safety
  risk management, 51–53
  testing, 41
scenario planning, 12, 26, 31–32, 74–76, 87–88, 99, 100, 102–104, 114, 171, 224, 250
  outside-in strategy, 74–76
Schmidt, Maik, 17
secret cyborgs, 9
Sedol, Lee, 65–66
self-disruption, 94
self-driven governance, 49–51
self-learning models, 78
self-regulation, 45–47, 59, 184
Shambaugh, Scott, 197
Shereshevsky, Salomon, 3, 4, 13, 249, 251
Siemens, 21, 179
Siemens Digital Industries, 116
Silicon Valley, 41, 46, 197
Sireau, Nick, 11
Six Paths Framework, 150–154, 161, 162
Small-S strategy, 84
smart city development, 76

societal stakes, 256
SoftBank, 43
Soviet Union, 81
Sports Illustrated, 48, 49
Spotify, 220, 221
Standard Chartered, 132, 135
Stanford 2025 AI Index, 42
Stanford 2026 AI Index Report, 43
Stanford Index study, 44
strategic advantage, 25, 28, 102–103, 190, 248
strategic balance, art of, 79
strategic experimentation, 70
strategic fitness, 86
strategic human-AI partnership, 222–224
strategic prioritization, 179, 181
strategic self-disruption, 94
strategy frameworks, 67–68, 70, 81, 83, 86
strategy/innovation matrix, 83–84
stress-test systems, 31
structural argument, 57–58
structural reconfiguration, 115
Suez Canal closure, 103
Sullivan, Jake, 41–42
supervised learning, 176
surface-level sensing, 111
Swift, Jonathan, 194, 260
systemic learning, 74

system owners, designated, 32, 49, 50
system simulation, 113

T

Tay (AI chatbot), 185
Taylor, Frederick, 83
Tesla, 78
Time Machine, 220
TIP. *See* Transition Integrity Project (TIP)
“To Be” Strategy Canvas, 159–162
“To Be” value curve, 155, 156
top-down vs bottom-up approaches, 84–85
Toyota, 99, 119
traffic management systems, 74
TRAKKA, 145
Transition Integrity Project (TIP), 98
transparency, artificial intelligence (AI) systems, 53–55
TrendSpotter, 91
TruEra, 132
Trump, Donald J., 41, 43
trust, artificial intelligence (AI) system, 43–44
Trust Dial protocol, 202, 210–213, 215, 216
Turing, Alan, 172
Tversky, Amos, 67

U

UC Berkeley, 56–57
Unilever, 22, 33, 34, 170

United States (U.S.), 18, 41
unsupervised learning, 176
Upper Paleolithic Revolution, 167
UPS, 36, 37
urban infrastructure, 76
U.S. defense community, 75
utility levers, 148–150

V

value proposition, 84, 146, 149, 151, 152, 154, 156, 157, 159, 160, 162
velocity, 22, 55, 207
visual processing system, 179–180

W

Watson, Thomas Sr., 166
Watson Health division, 69
Welch, Jack, 166, 219
Wells, H.G., 220
Wilson, Edward, 13
workforce shift, 255
World Economic Forum, 255
World Health Organization, 91

Z

Zara, 118
ZIP code hiring algorithms, 58
zones, artificial intelligence (AI), 223–224
Zones of Engagement tool, 235

Intelligence as a Commodity: Human Judgment as a Competitive Advantage, First Edition. Rajeev Kapur, Jason Hunter.
© 2027 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2027 by John Wiley & Sons, Inc.

Copyright © 2027 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at www.wiley.com/go/permission.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: WILEY and the Wiley logo are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates, in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and the authors have used their best efforts in preparing this work, including a review of the content of the work, neither the publisher nor the authors make any representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. Certain AI systems have been used in the creation of this work. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002. For product technical support, you can find answers to frequently asked questions or reach us via live chat at https://support.wiley.com.

If you believe you've found a mistake in this book, please bring it to our attention by emailing our reader support team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Control Number: 2026945329

ISBN 9781394468126 (Cloth)
ISBN 9781394468140 (ePDF)
ISBN 9781394468133 (ePub)

Cover Design: Wiley
Cover Image: Vooper/stock.adobe.com

For everyone who suspects the future still needs them.
It does.

WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.