

# **THE HUMAN AND MACHINE COMPANY**

LEADERSHIP, WORK, AND DECISION-MAKING IN  
THE AGE OF INTELLIGENT SYSTEMS

**RYAN L. KING**

Copyright © 2026 by Ryan L. King

All rights reserved.

No part of this book may be reproduced in any form or by any electronic or mechanical means, including information storage and retrieval systems, without written permission from the author, except for the use of brief quotations in a book review.

*For CIOs who have run enough pilots.*



## ABOUT THE AUTHOR



Ryan King is an executive strategist and founder of RLK Consulting, where she advises CIOs, CTOs, and senior IT leaders on strategy, performance, and organizational transformation. With nearly 15 years in consulting, including a decade across McKinsey & Company and Deloitte, she has spent her career working alongside technology executives responsible for enterprise-scale systems, large portfolios, and board-level accountability.

She specializes in advisory support for CIOs in transition, whether stepping into the role, stabilizing performance under scrutiny, navigating cost pressure, modernizing delivery practices, or preparing for succession. She works shoulder to shoulder with technology leaders as they redefine operating models, clarify decision rights, simplify portfo-

lios, and reposition IT as a front-office business capability rather than a back-office cost center.

Her work spans industries including insurance, banking, health-care, logistics, government, and consumer businesses. She founded RLK Consulting in 2025 on the belief that leaders should have access to top-tier strategic advice without the cost and overhead of a large consulting firm.

Ryan earned her BBA in Finance from the College of William & Mary. She lives in Virginia with her husband, a recovering consultant himself, and their two young boys.

# PREFACE

For nearly fifteen years, I have advised Chief Information Officers on how to run their organizations as real businesses.

Not as cost centers.

Not as back-office enablers.

Not as order takers.

As businesses.

Over the course of my career, I have worked alongside CIOs across nearly every sector of the economy. I have advised leaders inside Fortune 50 enterprises and small founder-led companies. I have supported federal agencies, state and local governments, nonprofit institutions, publicly traded corporations, privately held firms, and venture-backed organizations. I have worked in shipping and logistics, consumer retail, aerospace and defense, insurance, banking and financial services, and technology environments.

The industries differ. The regulatory environments differ. The governance models differ.

The leadership challenges are strikingly similar.

CIOs are responsible for operating complex technical ecosystems while simultaneously enabling revenue growth, protecting margin, ensuring resilience, and maintaining competitiveness. They must modernize infrastructure, manage cyber risk, optimize cost, deliver

innovation, and explain all of it coherently to executive peers and boards of directors.

My work sits at that intersection.

Because of where I sit, I could see what many leaders could not yet explain. Artificial intelligence was working, but only where management structures allowed it to work.

Across large enterprises, CIOs were launching pilot after pilot of artificial intelligence initiatives. Chatbots in customer service. Code assistants in engineering. Document summarization in legal. Forecasting enhancements in finance. Automation agents in operations.

Individually, these efforts worked. Productivity improved locally. Costs declined in pockets. Cycle times shortened within teams.

Yet at the enterprise level, very little changed.

Boards did not see meaningful margin expansion. Operating models remained largely intact. Decision speed continued to be constrained by legacy structures. CIOs could point to dozens of successful experiments but struggled to explain what they collectively meant for how the organization should be run.

At the same time, I observed small and mid-sized businesses achieving disproportionate productivity gains with minimal coordination overhead. A single professional could produce research and analysis that previously required a team. Mid-market firms compressed planning cycles because fewer approval layers stood between insight and action. The informational advantage historically associated with scale was beginning to erode.

What I was observing did not fit a technology adoption narrative.

It fit a management narrative.

Artificial intelligence was not failing inside large organizations. It was colliding with management structures designed for a different information environment. The tools were working. The organization surrounding them had not yet adapted.

This book is an attempt to describe that shift.

I am not writing from speculation. The perspective in these pages is drawn from years of advisory work, extensive conversations with CIOs across industries, and patterns visible across organizations facing the same problem simultaneously. The change underway is not hypo-

thetical and it is not distant. Leaders are already experiencing its effects, often without a framework for interpreting them.

Artificial intelligence is often discussed as a technology transformation. What I have come to believe is simpler and more consequential.

It is a management transformation.



# INTRODUCTION

Artificial intelligence is no longer a question.

Boards no longer debate whether it matters.

Executives no longer ask whether to experiment.

Employees no longer wait for permission.

The technology works.

Inside many organizations, individuals quietly report meaningful productivity gains. Work that once required hours now takes minutes. Documents are drafted faster. Research is synthesized almost instantly. Developers move through unfamiliar codebases with surprising fluency. Analysts generate structured thinking on demand. At the level of the individual contributor, the improvement is tangible and often dramatic.

At the level of the enterprise, however, the picture is less clear.

Cycle times have not collapsed. Financial results have not shifted materially. Competitive positions remain largely intact. Leadership teams struggle to point to organization-wide performance gains despite visible activity across departments. Companies feel busier, more experimental, more technologically sophisticated. They do not yet feel fundamentally different.

This contradiction is not anecdotal. Large-scale surveys report a consistent pattern: a majority of organizations now use artificial intelli-

gence in at least one business function, yet only a small fraction can point to measurable bottom-line impact.<sup>1</sup> Successful pilots proliferate, but few scale into sustained operational advantage. Employee adoption increases, while executive confidence in enterprise returns remains tentative.

Leaders recognize the tension immediately. The tools are powerful. The results are promising. And yet the organization itself does not seem to behave differently.

This book exists to explain why.

Most conversations about artificial intelligence begin with the technology. They focus on models, prompts, vendors, and speculation about future disruption. These questions matter. But inside organizations they are rarely the source of confusion.

The central challenge is not understanding what artificial intelligence is. It is understanding what artificial intelligence does to how a firm operates.

For more than a century, companies have been organized around the management of scarce information. Reliable knowledge was difficult to gather, interpret, and verify. Reports existed to transmit understanding upward. Meetings aligned interpretation across groups. Approval chains reduced uncertainty. Supervisory layers were not merely administrative structures; they were mechanisms for processing limited cognitive capacity. The architecture of the modern firm evolved in response to those constraints.

Artificial intelligence alters those assumptions.

Employees can now generate analysis, explanations, and structured reasoning in minutes that previously required days. Managers can interrogate data directly rather than waiting for summaries. Junior staff can explore questions that once required specialized expertise. The constraint is no longer access to information but how that information is coordinated, validated, and acted upon.

The immediate consequence is not industry collapse or overnight disruption. It is organizational tension.

---

1. McKinsey & Company. *The State of AI in 2023: Generative AI's Breakout Year*. McKinsey Global Survey on AI, 2023.

The structure of the firm still reflects a world in which knowledge was scarce. Daily work increasingly reflects a world in which knowledge is abundant. The mismatch between those two realities explains why so many leaders feel unsettled. The technology functions. The management system surrounding it lags.

Artificial intelligence is therefore not primarily a software deployment. It is a change in how knowledge work is coordinated.

Many organizations misread what they are observing. They see pilots without transformation, adoption without impact, productivity without performance. Some conclude the technology remains immature. Others assume meaningful gains will arrive gradually as models improve.

In reality, the organization has already adopted artificial intelligence. The organization itself has not yet changed.

This is not a book about tools, prompts, or vendors. It is a book about management.

Its argument is straightforward. Artificial intelligence changes how decisions move through organizations. Firms that recognize this early will gain advantage not by possessing better models, but by redesigning how authority, accountability, and information flow.

The chapters that follow develop three related ideas.

First, artificial intelligence initially alters behavior before it alters strategy. Individuals change how they work before organizations change how they function, and organizations change before competitive positioning visibly shifts.

Second, firms adopt artificial intelligence through a predictable progression. Individuals experiment. Organizations attempt to govern usage. Only later do some redesign workflows and decision authority. Many enterprises today appear stalled between governance and redesign, accumulating activity without structural impact.

Third, progress requires leadership intervention. Artificial intelligence forces reconsideration of governance, accountability, skills, organizational design, and performance measurement. The challenge is not experimentation. It is managerial redesign.

Artificial intelligence is the catalyst. Management is the subject.



# PART ONE

## THE MISUNDERSTANDING

For many executives, artificial intelligence appears as another technology trend.

It is discussed alongside cloud computing, cybersecurity, data platforms, and digital transformation initiatives. Organizations form committees, evaluate vendors, and debate tools. Leaders ask which use cases matter and which investments to prioritize.

This framing feels reasonable. It is also incomplete.

Artificial intelligence is not primarily a new category of software. It is a change in how work itself is performed. The shift does not begin with corporate strategy, industry structure, or competitive positioning. It begins with individual employees.

Across organizations, people are quietly changing how they produce analysis, draft communication, research problems, and prepare decisions. The change is uneven, informal, and often invisible to leadership. Yet it is already reshaping organizations from the inside outward.

The early evidence confuses managers because it does not appear where they expect. Productivity improves in pockets before it appears in enterprise metrics. Employees alter workflows before companies

authorize them. The organization begins changing before leadership formally decides to change.

As a result, many executives interpret artificial intelligence as overhyped or underperforming. In reality, they are observing the first phase of a different phenomenon: a behavioral shift that precedes an organizational shift.

Artificial intelligence does not initially transform companies through strategy.

It transforms them through behavior.

# CHAPTER 1

# AI IS NOT ABOUT TECHNOLOGY

WHY ARTIFICIAL INTELLIGENCE  
IS FUNDAMENTALLY A CHANGE  
IN MANAGEMENT RATHER  
THAN A CHANGE IN SOFTWARE

IN NEARLY EVERY company I work with, there comes a moment when a conversation about artificial intelligence reveals that people are not actually discussing the same subject.

The meeting typically begins as a technology discussion. A CIO explains model selection, vendors, security constraints, and architecture. The information is accurate and useful, but the reaction around the table is muted. The CEO listens but remains uneasy. The CFO asks questions that seem disconnected from the technical explanation. A board member eventually asks something that sounds simple but is difficult to answer: “What does this mean for how the company actually runs?”

That question exposes the real issue. Leaders sense that artificial intelligence matters in a deeper way than previous technologies, yet they cannot locate where the change will occur. They are accustomed to evaluating technology as infrastructure. Artificial intelligence does not behave like infrastructure.

For three decades, major corporate technology initiatives primarily altered processes. Enterprise resource planning systems standardized transactions. Customer relationship management systems centralized customer data. Cloud computing relocated computing resources. Even mobility changed access more than it changed organizational logic.

Each of these required investment and discipline, but they did not fundamentally redefine how managers supervised work.

Artificial intelligence does.

Generative systems now perform tasks historically tied to professional expertise: drafting written analysis, interpreting unstructured information, synthesizing research, and producing working software. These are not simply clerical functions. They are elements of cognitive labor. The introduction of a technology capable of participating in cognitive work alters how organizations coordinate activity, and coordination is the essence of management.

The reason leaders feel disoriented is that the impact appears simultaneously obvious and invisible. Employees are using these systems constantly, yet company-level performance indicators often appear unchanged.<sup>1</sup> This is not anecdotal. Microsoft's Work Trend Index research reported widespread daily generative AI use among knowledge workers across industries, including organizations without formal programs governing it.<sup>2</sup> At the same time, McKinsey's global surveys of AI adoption show that while a majority of organizations now report some AI use, only a minority report material financial impact from it.<sup>3</sup> The coexistence of widespread usage and limited measurable enterprise benefit appears contradictory only if one assumes productivity should appear immediately at the organizational level.

It does not.

The models are capable enough already. The constraint is managerial.

For most of modern corporate history, the structure of organizations assumed information moved slowly and was expensive to interpret. Work therefore proceeded through layers. Analysts gathered data. Managers reviewed it. Directors validated conclusions. Executives

---

1. Brynjolfsson, Erik. "The Productivity Paradox of Information Technology." *Communications of the ACM*, vol. 36, no. 12, 1993, pp. 66–77.

2. Microsoft. *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Microsoft, 2024.

3. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

made decisions. The structure served a rational purpose: it reduced error and distributed responsibility when information was scarce and verification was costly.<sup>4</sup>

Artificial intelligence changes the economics of information. Producing analysis, summarizing knowledge, and drafting recommendations no longer require large teams or long cycles. When the cost of generating and interpreting information falls, the justification for certain coordination structures weakens. Ronald Coase's classic theory of the firm argued that organizations grow to minimize the cost of coordinating work.<sup>5</sup> When coordination becomes cheaper, optimal structure changes. Artificial intelligence reduces certain coordination costs.

This explains why technical discussions about AI often feel incomplete to senior leaders. The true implications are not confined to software selection or infrastructure decisions. They concern supervision, accountability, and decision rights. If a system participates in producing analysis, who is responsible for the judgment? If employees can independently obtain expertise, how should oversight function? These are management questions.

Organizations are therefore encountering a new type of transformation. The challenge is not learning how to deploy a tool. It is learning how to manage a form of work in which human and machine cognition coexist. Companies attempting to treat artificial intelligence as another IT implementation frequently find themselves running pilots indefinitely. The technology functions correctly. The organization surrounding it remains unchanged.

Executives sometimes say they are waiting for the technology to mature. More often, the management model must mature.

Artificial intelligence will eventually reshape markets and competition. But before industries change, firms must learn how to operate differently internally.

The first disruption of AI is organizational, not competitive. It alters

---

4. Simon, Herbert A. *The New Science of Management Decision*. Prentice-Hall, 1960.

5. Coase, Ronald H. "The Nature of the Firm." *Economica*, vol. 4, no. 16, 1937, pp. 386-405.

how work is coordinated inside companies before it alters how companies compete outside them.

# CHAPTER 2

# ADOPTION STARTS IN SECRET

WHY PRODUCTIVITY IMPROVES  
AT THE INDIVIDUAL LEVEL  
LONG BEFORE ORGANIZATIONS  
RECOGNIZE CHANGE

THE MOST RELIABLE indicator that artificial intelligence has entered an organization is not a formal announcement. It is a behavioral shift.

In many companies, leadership believes AI adoption is limited because no official program exists. Yet employees are already relying on it daily. Engineers consult models to understand unfamiliar code. Analysts summarize documents. Product managers draft plans. Marketing staff develop campaign concepts. The change is happening quietly.

When I ask employees privately whether they use AI, most acknowledge they do. When I ask whether they formally reported doing so, most have not. This pattern is not unusual. Research into workplace adoption of generative AI shows that individuals frequently adopt useful tools before policies catch up.<sup>1</sup> Microsoft's surveys document large numbers of employees initiating use independently, and

---

1. Microsoft. *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Microsoft, 2024.

Deloitte’s enterprise AI studies note that organizational governance typically trails behavior.<sup>2,3</sup>

The result is a peculiar situation. Individual productivity improves while organizational metrics remain stable.

Economists have observed a similar phenomenon with prior general-purpose technologies. Erik Brynjolfsson and his colleagues described a “productivity paradox” in which technological improvements do not immediately translate into measured economic gains because complementary organizational changes lag.<sup>4</sup> A worker performing tasks faster does not automatically increase organizational throughput if surrounding processes remain unchanged.

Artificial intelligence is currently in that stage.

Knowledge work rarely produces easily measurable output. When a financial analyst reviews contracts in one hour instead of four, accounting systems do not record a productivity event. When a software developer resolves a problem quickly using generated explanation, the improvement is absorbed into normal operations. The organization benefits locally but not systemically.

The consequence is confusion among leaders. Vendors and employees describe meaningful improvement, yet executive dashboards show little change.<sup>5</sup> This does not indicate exaggeration. It indicates misaligned measurement.

Companies measure output, not effort avoided.<sup>6</sup>

Artificial intelligence reduces the time required to complete cognitive tasks, but organizations are still structured around workflows designed for longer cycles. If a process requires approvals, handoffs, and scheduled reviews, accelerating one step does not change the

---

2. Microsoft. *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Microsoft, 2024.

3. Deloitte. *State of Generative AI in the Enterprise 2024*. Deloitte Insights, 2024.

4. Brynjolfsson, Erik. “The Productivity Paradox of Information Technology.” *Communications of the ACM*, vol. 36, no. 12, 1993, pp. 66–77.

5. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*.

6. Brynjolfsson, Erik, and Lorin M. Hitt. “Beyond Computation: Information Technology, Organizational Transformation and Business Performance.” *Journal of Economic Perspectives*, vol. 14, no. 4, 2000, pp. 23–48.

duration of the overall process. Productivity gains remain trapped at the individual level.

This explains why early AI benefits are often invisible. They occur at the edges of organizations where work is performed rather than at the center where performance is measured.

Another factor reinforces this invisibility. Social norms lag technology. Professionals have historically associated effort with competence.<sup>7</sup> When work is produced quickly, employees sometimes fear that revealing their use of assistance will be interpreted as reducing effort rather than improving effectiveness. As a result, they refine outputs and present them as conventional work product.<sup>8</sup> The organization receives better work but not the knowledge of how it was produced.

This behavior creates what might be called a shadow workflow. The formal process still exists. Alongside it, a more efficient informal process operates quietly. Over time, some individuals become dramatically more effective than peers not because they are inherently more capable but because they have learned how to incorporate machine assistance into their thinking process.

*We've seen it before. In the early 1980s, many corporate finance departments banned personal spreadsheet software because it bypassed centralized data processing controls. Yet employees quietly adopted programs such as VisiCalc and Lotus 1-2-3 on personal machines because they dramatically reduced modeling time. Informal spreadsheet models proliferated before IT governance frameworks caught up. What began as unofficial use eventually redefined how financial planning and analysis functioned inside firms.<sup>9</sup>*

Organizations eventually notice these performance differences but often misattribute them to individual talent rather than systemic change.

The challenge for management is to recognize that the adoption has

---

7. Parasuraman, Raja, and Victor Riley. "Humans and Automation: Use, Misuse, Disuse, Abuse." *Human Factors*, vol. 39, no. 2, 1997, pp. 230–253.

8. Deloitte. *State of Generative AI in the Enterprise 2024*. Deloitte Insights, 2024.

9. Yates, JoAnne. *Structuring the Information Age: Life Insurance and Technology in the Twentieth Century*. Johns Hopkins University Press, 2005.

already occurred. The question is not whether employees will use artificial intelligence but whether the organization will learn from how they are using it. Firms that attempt to prohibit or tightly restrict use often drive it further underground, while firms that observe carefully can identify where real gains are emerging.

At this stage, leaders face a strategic decision disguised as a policy decision. They can treat AI use as a compliance problem, or they can treat it as a source of operational insight. The companies that eventually benefit most are typically those that formalize useful behavior after observing it rather than attempting to design adoption entirely from the top down.

The important realization is that productivity appears first in individuals because individuals can change behavior faster than organizations can change structure. Only when management redesigns workflows to incorporate these behavioral shifts does organizational productivity emerge.<sup>10</sup>

---

10. Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton, 2014.

# CHAPTER 3

# COMPANIES CHANGE BEFORE INDUSTRIES DO

HOW INTERNAL BEHAVIOR  
SHIFTS RESHAPE FIRMS  
BEFORE COMPETITIVE  
STRATEGY VISIBLY MOVES

WHEN PEOPLE THINK about technological change, they usually imagine competitive upheaval. They expect new companies to replace old ones, markets to reorganize, and industries to transform. Those events do occur, but they are rarely the first effects of a general-purpose technology.<sup>1</sup>

Before markets change, organizations change.<sup>2</sup>

Artificial intelligence will eventually influence competition, but its earliest and most significant impact is internal. It alters how information moves through a company and how decisions are made. Because companies are coordination systems rather than collections of tasks, changing coordination changes the firm itself.

Historical experience supports this pattern. Economists studying general-purpose technologies such as electricity and computing found that their largest productivity gains occurred only after firms reorganized processes around them.<sup>3</sup> The technology alone did not create

---

1. Bresnahan, Timothy F., and Manuel Trajtenberg. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics*, vol. 65, no. 1, 1995, pp. 83–108.

2. David, Paul A. "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox." *American Economic Review*, vol. 80, no. 2, 1990, pp. 355–361.

3. Brynjolfsson, Erik, and Lorin M. Hitt. "Beyond Computation: Information Technol-

value; the redesign of work did. Bresnahan and Trajtenberg described these technologies as enabling innovations that require complementary organizational adjustments before economic impact appears.<sup>4</sup>

Artificial intelligence functions similarly but at a cognitive level. It reduces the cost of generating and interpreting information. When information becomes easier to obtain and analyze, the logic of hierarchical supervision weakens. Some managerial layers exist primarily to aggregate and interpret information from below. If individuals can access analysis directly, the purpose of certain reviews changes.

Research on information technology adoption already observed this effect. Studies of firms adopting advanced information and communication technologies found that decision-making authority tended to move closer to the operational level as information processing improved.<sup>5</sup> Artificial intelligence accelerates that trend by assisting not only data transmission but also interpretation.

Inside companies, this manifests in subtle ways. Certain employees become central nodes of coordination because they can synthesize information quickly. Managers rely on them more frequently. Decisions occur faster but outside formal structures. Reporting relationships feel less predictive of influence than practical capability.

Executives often interpret these developments as talent concentration. In reality, they are structural signals. The organization is adjusting informally before it adjusts formally.

This explains why industry disruption has not yet materialized at scale. Firms are still learning how to operate with altered coordination costs. Once enough organizations redesign workflows and authority structures, competition will shift because operating speed and adaptability will differ dramatically between firms.

The progression therefore runs opposite to popular expectation.

---

ogy, Organizational Transformation and Business Performance." *Journal of Economic Perspectives*, vol. 14, no. 4, 2000, pp. 23–48.

4. Bresnahan, Timothy F., and Manuel Trajtenberg. "General Purpose Technologies 'Engines of Growth?'" *Journal of Econometrics*, vol. 65, no. 1, 1995, pp. 83–108.

5. Aghion, Philippe, and Jean Tirole. "Formal and Real Authority in Organizations." *Journal of Political Economy*, vol. 105, no. 1, 1997, pp. 1–29.

The technology first changes how companies function internally. Only later does it change how companies compete externally.

Leaders seeking evidence of AI progress should therefore observe management behavior rather than market outcomes. Shorter decision cycles, fewer review layers, and broader participation in analysis are more meaningful indicators than immediate revenue changes. They indicate that the firm is adapting to a new mode of work.

Artificial intelligence does not begin by transforming industries. It begins by transforming how companies organize thought. When enough firms reorganize successfully, industry change follows naturally.



# CHAPTER 4

## THE PILOT TRAP

WHY MOST AI INITIATIVES  
FAIL WHEN COMPANIES  
OPTIMIZE TASKS INSTEAD  
OF REDESIGNING WORK

ONCE LEADERS ACCEPT that artificial intelligence matters, their first instinct is usually practical. They want to begin doing something tangible. Organizations are accustomed to responding to new technology through structured initiatives, and so they construct one. A committee forms, or a working group, or a center of excellence. Meetings are scheduled. Departments are asked to submit opportunities for application.

Within weeks, a familiar catalog appears.

Customer service proposes summarizing support interactions. Marketing suggests generating draft content. Engineering wants assistance with code documentation and debugging. Finance proposes automated classification of contracts or invoices. Human resources considers resume screening. Each proposal is rational. Each appears valuable. Each is approved as a pilot.

The pilots usually succeed.<sup>1</sup>

The software performs as demonstrated. Customer calls are summarized accurately. Draft documents are produced quickly. Developers complete tasks faster. The project reports are positive. Yet several

---

1. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. "Generative AI at Work." *National Bureau of Economic Research Working Paper*, no. 31161, 2023.

months later, executives observe something perplexing. The organization itself has not meaningfully changed. Operating metrics remain largely stable. The promised transformation appears absent.<sup>2</sup>

This experience has now become common enough that it no longer surprises me when leaders describe it. They are not frustrated with the technology. They are confused by the outcome. The tools clearly work, yet the business does not feel fundamentally different.

The explanation lies in how companies understand work.

Organizations tend to interpret artificial intelligence as a way to perform existing tasks faster. This interpretation is intuitive because most prior business technologies behaved that way. Software automated bookkeeping, recordkeeping, scheduling, and transaction processing.<sup>3</sup> Productivity gains resulted from reducing the effort required to complete defined steps within an existing process. The process itself remained largely intact.

Artificial intelligence does not fit comfortably into that category. It does not merely execute predefined steps. It participates in reasoning, drafting, and interpretation. When a technology alters reasoning, the natural unit of improvement is no longer the task but the workflow.

A workflow is not a series of activities. It is a sequence of decisions. A customer support interaction, for example, is not simply answering a question. It involves identifying the issue, selecting a response, determining whether escalation is required, and documenting the outcome. Each step exists because information historically had to be gathered, interpreted, and validated by people at different levels.

When artificial intelligence is inserted into one step of this process, that step accelerates. The rest of the process does not automatically adjust. Reviews still occur. Approvals remain scheduled. Escalations remain defined by prior assumptions about how long information takes to produce and verify. The system's speed is therefore limited not by the accelerated task but by the unchanged structure around it.

Operations researchers have long observed this effect in complex

---

2. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*.

3. Davenport, Thomas H. *Process Innovation: Reengineering Work through Information Technology*. Harvard Business School Press, 1993.

systems. Improvements to individual components do not increase system performance if constraints elsewhere remain. In manufacturing, improving the speed of one machine does not increase throughput if another stage remains slower. Knowledge work behaves similarly, though the constraints are less visible.

*When advanced airport screening machines were introduced, early deployments did not immediately reduce passenger wait times. Security lanes moved faster at the scanning point, but bottlenecks persisted at document checks and staffing levels. Only when airport operators redesigned the entire passenger flow did overall throughput improve. Speeding one checkpoint did not accelerate the system.<sup>4</sup>*

In many companies, artificial intelligence shortens the time required to prepare analysis but not the time required to review it. Managers continue to require meetings, presentations, and multi-level approvals designed for slower information cycles. The organization experiences faster preparation and unchanged decision speed. From the perspective of leadership, little has changed.

This leads to what I have come to think of as the pilot plateau. Companies successfully implement use cases but fail to achieve transformation because they have improved activities without reconsidering coordination.

Research reflects this pattern. Studies of enterprise AI adoption consistently show that organizations can deploy specific applications yet struggle to scale impact across the enterprise.<sup>5</sup> McKinsey's analyses of AI transformations note that scaling requires changes in processes, governance, and roles rather than merely increasing the number of use cases. The challenge is not technical capability but organizational adaptation.

What organizations call an AI strategy is often a collection of pilots. What they actually need is a redesign of work.

---

4. Neufville, Richard de, and Amedeo R. Odoni. *Airport Systems: Planning, Design, and Management*. McGraw-Hill, 2003.

5. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

The distinction becomes clearer when considering decision-making. In traditional workflows, a report may be prepared by analysts, reviewed by a manager, revised, and then presented to leadership. The sequence exists because analysis is time-consuming and difficult to verify. Artificial intelligence dramatically shortens preparation. The manager's role should therefore shift toward evaluating conclusions rather than coordinating preparation. Yet many organizations preserve the same review structure, effectively adding faster analysis to an unchanged hierarchy.

The result is paradoxical. Individuals feel more productive, but decision cycles do not accelerate. The organization benefits locally but not systemically.

Over time, employees adapt informally. They begin using AI to anticipate questions from reviewers. They prepare multiple scenarios quickly. They shorten their own workflows even when the formal workflow remains intact. The organization develops an unofficial operating mode layered beneath the official one.

Leaders sometimes interpret the limited organizational impact of AI pilots as evidence that the technology has been oversold. More often, the technology has been misapplied. Artificial intelligence improves reasoning capacity. If reasoning capacity changes, workflows should be simplified. When workflows remain unchanged, the benefit is diluted.

The critical mistake companies make is assuming that the objective of adoption is to complete existing tasks more efficiently. The objective should be to reconsider why the tasks exist.

This is not easy because tasks are embedded in accountability structures. Approval layers often exist to manage risk rather than to process information. Reducing them requires trust in new forms of verification and oversight. Artificial intelligence alters what can be known quickly, but organizations must adjust how they decide based on that knowledge.

In practice, transformation occurs only when leaders stop asking where artificial intelligence can be inserted and begin asking which steps in a workflow are no longer necessary. At that moment, adoption moves beyond experimentation.

Until then, companies will continue to run successful pilots that leave the firm essentially unchanged.



# CHAPTER 5

## WHO BENEFITS FIRST

WHY SMALL FIRMS ADOPT  
FASTEST, MID-MARKET  
FIRMS SEE RETURNS  
EARLIEST, AND ENTERPRISES  
ULTIMATELY GAIN THE MOST

EXECUTIVES FREQUENTLY ASSUME that the largest organizations will benefit from artificial intelligence first. The reasoning appears obvious. Large companies have more capital, more data, and more technical talent. They have dedicated IT departments and established vendor relationships. If a technology promises productivity improvement, scale should amplify its value.

Yet the early pattern of AI adoption contradicts that expectation.

Small organizations often change behavior first. Mid-sized firms frequently capture measurable operational benefit soon after. Large enterprises tend to realize the most significant long-term gains, but only after a longer and more complex transition.

The order appears reversed because the constraint is not capability. It is coordination.<sup>1</sup>

*When personal computers appeared in the late 1970s, they were first adopted by small businesses rather than large corporations. Large firms relied on centralized mainframe departments and formal data processing procedures, which made individual experimentation difficult. Smaller firms and entrepre-*

---

1. Mintzberg, Henry. *Structure in Fives: Designing Effective Organizations*. Prentice Hall, 1983.

*neurs quickly used spreadsheets and word processing. Only later did large enterprises reorganize around distributed computing and realize larger gains.*<sup>2</sup>

A small business typically operates through a limited number of decision-makers. Information flows directly to owners or leaders. When an individual learns to incorporate artificial intelligence into daily work, the entire firm effectively changes at once. A single person may handle marketing, operations, and financial analysis with assistance that previously would have required staff or external providers. The organization does not need to redesign formal processes because it has few.

The impact is immediate. The owner writes proposals faster, responds to customers more quickly, and analyzes information independently. The business becomes more capable relative to its size. The change feels dramatic because there are minimal structural barriers to adopting new behavior.

Mid-market organizations, usually consisting of several hundred to a few thousand employees, occupy a different position. They possess defined departments and recurring processes but remain flexible enough to adjust them. When teams begin using artificial intelligence, leadership can modify workflows, redefine responsibilities, and adopt new practices within months rather than years.

This environment often produces the earliest measurable return on investment. Departments streamline reporting. Customer support improves response time. Product development cycles shorten because coordination between functions becomes easier. These improvements appear in operating performance because processes are structured enough to measure yet flexible enough to redesign.<sup>3</sup>

Large enterprises face a more complicated situation. They possess the resources and data to benefit significantly from artificial intelligence, but they also possess deeply embedded coordination structures. Responsibilities are distributed across business units. Approval

---

2. Ceruzzi, Paul E. *A History of Modern Computing*. 2nd ed., MIT Press, 2003.

3. Brynjolfsson, Erik, and Lorin M. Hitt. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives*, vol. 14, no. 4, 2000, pp. 23–48.

processes manage risk across legal, compliance, and operational domains. Systems integrate with many stakeholders. Any change to workflow affects multiple functions simultaneously.

As a result, behavioral adoption occurs widely within large organizations but formal adoption proceeds cautiously. Employees use artificial intelligence individually, while leadership establishes policies, governance frameworks, and oversight. The organization experiments extensively yet struggles to produce enterprise-level performance change. The complexity that enables scale also slows adaptation.

Over time, however, the potential advantage of scale reappears. Once a large organization redesigns workflows, the improvement affects thousands of employees simultaneously. Decision cycles accelerate across business units. Knowledge sharing improves. The magnitude of impact eventually exceeds that of smaller firms, but reaching that point requires structural change.

Research on organizational adoption of information technology offers context. Studies examining firms adopting advanced digital tools found that performance gains depended less on the presence of technology and more on complementary organizational changes such as training, new processes, and adjusted management practices.<sup>4</sup> Large organizations required more time to implement those complementary changes but ultimately realized larger gains when they did.

Artificial intelligence magnifies this pattern. The early advantage belongs to agility.<sup>5</sup> The enduring advantage belongs to coordinated scale.

This progression explains why executives observing small firms appear surprised by their capability. A small team may operate with the responsiveness and analytical capacity once associated with far larger companies. The technology temporarily narrows the gap between size and capability. However, once large organizations adapt their structures, their scale again becomes a powerful advantage.

The important implication is that no single category of firm perma-

---

4. Bresnahan, Timothy F., and Manuel Trajtenberg. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics*, vol. 65, no. 1, 1995, pp. 83–108.

5. McAfee, Andrew, and Erik Brynjolfsson. *The Second Machine Age*. W. W. Norton, 2014.

nently dominates the technology. Each benefits at a different stage. Small companies demonstrate what is possible. Mid-market firms operationalize it. Enterprises institutionalize it.

Leaders often misinterpret early observations. Seeing small firms adopt quickly, they assume scale is no longer relevant. Seeing large firms move slowly, they assume bureaucracy is inevitable. Both conclusions are incomplete. The sequence reflects the speed at which different organizations can modify coordination.

Artificial intelligence therefore changes competition not by eliminating scale but by altering when scale matters. Early in adoption, speed of behavioral change matters most. Later, the ability to reorganize complex operations determines advantage.

Understanding this sequence clarifies the broader argument of this book. Artificial intelligence is not merely a technology firms install. It is a shift in how organizations process knowledge and make decisions. Firms that adapt their structures earliest demonstrate its potential. Firms that eventually align structure with capability realize its full value.

The technology does not choose winners immediately. Organizational adaptation does.

# PART TWO

## THE AI ADOPTION CURVE

Once leaders recognize that artificial intelligence is altering work, a second confusion appears.

They expect adoption to follow the pattern of earlier enterprise technologies. A tool is selected, deployed across the organization, employees are trained, and benefits follow implementation.

Artificial intelligence does not follow that sequence.

Adoption occurs in stages, and those stages are driven less by technical rollout than by how decisions move through the firm.

First, individuals change how they perform their work.

Second, organizations attempt to manage the technology through pilots, policies, and oversight.

Third, a smaller number of firms redesign workflows and decision authority.

Many companies stall between the second and third stages. They deploy tools successfully but fail to achieve meaningful enterprise impact. They accumulate pilots yet struggle to produce operational change. The problem is not capability. The technology functions. The problem is alignment between how work is done and how the organization is structured.

Early adoption therefore appears informal and uneven. The pilot

phase expands but does not scale. Real value emerges only when workflows, roles, and management practices are adjusted to reflect the new economics of information.

Artificial intelligence does not create value simply because it is used.

It creates value when the firm reorganizes around it.

Understanding this progression clarifies why individual productivity improvements appear before enterprise results, why the pilot plateau occurs, and why transformation is ultimately organizational rather than technical.

# CHAPTER 6

## THE AI MATURITY CURVE

VALUE COMES FROM  
ORGANIZATIONAL  
CHANGE, NOT TOOL USAGE

BY THE TIME leaders begin asking how to scale artificial intelligence, they are usually already past the first stage of adoption without realizing it.

The question they ask sounds practical. They want to know which platform to standardize on, how to train employees, or how to measure return on investment. The assumption behind the question is that artificial intelligence adoption resembles prior technology rollouts. A new capability appears, the organization selects a tool, training follows, and results accumulate.

The difficulty is that the pattern does not unfold that way.

Artificial intelligence enters organizations unevenly. It does not begin with leadership approval. It begins with behavior.<sup>1</sup> Individuals experiment because the tool is accessible and immediately useful.<sup>2</sup> The organization notices only after those behaviors have already changed how work is performed. By the time executives create a formal plan, the company has already moved into the second phase of adoption, though it often believes it is still at the beginning.

---

1. Microsoft. *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Microsoft, 2024.

2. Noy, Shakked, and Whitney Zhang. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, vol. 381, no. 6654, 2023, pp. 187–192.

This is why discussions about artificial intelligence frequently feel misaligned. Different people in the same company are describing different realities.<sup>3</sup> Employees experience daily usefulness. Managers experience operational inconsistency. Executives experience strategic uncertainty.<sup>4</sup> All of them are correct because they are observing different stages of the same process simultaneously.

To understand this, it helps to think about artificial intelligence not as a project but as a progression.

Organizations adopt artificial intelligence through three stages. The stages are not determined by technical sophistication, and they are not determined by the quality of the models used. They are determined by how the firm changes its behavior and structure in response to the capability.

The first stage is behavioral adoption.

The second stage is managerial adoption.

The third stage is organizational adoption.

Each stage produces a different type of value, a different type of risk, and a different set of leadership challenges. Companies often confuse them because the same tools are present throughout all three stages. What changes is not the technology but the role the technology plays inside the company.

In the behavioral stage, individuals use artificial intelligence to improve their personal work. The company itself does not change. There is no formal policy, no coordinated workflow, and often little visibility. Employees adopt because it helps them complete tasks more easily. The organization neither directs nor meaningfully manages this behavior.

In the managerial stage, leaders recognize the capability and attempt to formalize it. Pilots are launched, governance committees form, vendors are evaluated, and use cases are documented. The company begins to manage artificial intelligence, but it still treats the technology as a set of tools applied to existing processes.

---

3. Microsoft and LinkedIn. *Work Trend Index Annual Report: AI Is Changing How We Work*, 2024.

4. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

In the organizational stage, something more substantial occurs. Workflows are redesigned. Roles shift. Decision rights change. Artificial intelligence becomes part of how work is structured rather than an addition to how tasks are performed. At this stage, measurable enterprise value appears.

The distinction between stages explains a pattern many executives struggle to interpret. A company can invest heavily in artificial intelligence and see limited impact while a smaller firm appears to benefit quickly.<sup>5</sup> The difference is not capability. It is stage.

A firm in the first stage may appear innovative because employees are highly productive individually. A firm in the second stage may appear organized because it has a visible program. A firm in the third stage is the one that actually changes performance because the organization itself adapts.

The confusion arises because leaders evaluate artificial intelligence through the lens of tools. They ask whether employees are using it, whether the technology is accurate, and whether vendors are reliable. Those questions matter, but they do not determine value. Value depends on whether the organization alters coordination.

Historically, general purpose technologies have required complementary changes before producing measurable gains. When electrification spread through manufacturing, factories initially replaced steam engines with electric motors but kept the same layout. Productivity improvements were limited. Only after factories redesigned the placement of machines and reorganized workflows did output increase significantly. Economic historians have documented that the largest productivity gains occurred not at the moment of technological introduction but after firms reorganized around the new capability.

*Businesses initially treated the telephone as a faster telegraph and routed calls through central operators. Over time, firms reorganized customer service, sales coordination, and management communication around direct voice contact. Departments that once relied on written correspondence shifted to immediate*

---

5. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven.*

*decision making. The economic impact emerged only after organizations changed how authority and communication were structured.*<sup>6</sup>

Artificial intelligence follows the same pattern. The technology allows work to be performed differently, but companies initially insert it into existing structures. They improve isolated tasks while preserving the system those tasks belong to. Leaders then conclude that the technology is promising but not transformative. The transformation arrives only when the organization changes.

The maturity curve therefore describes organizational readiness rather than technical maturity. A company does not progress from stage one to stage three because the models improve.<sup>7</sup> It progresses because leadership rethinks how work is coordinated.

The first stage creates hidden productivity.

The second stage creates visible experimentation.

The third stage creates enterprise performance change.

Recognizing the stages matters because each requires a different leadership response.

In the behavioral stage, the primary task of leadership is observation. The organization must learn how employees are actually using the technology. Attempts to control too quickly often suppress valuable experimentation. At this stage, artificial intelligence behaves more like a new form of literacy than a new system. People discover practical uses before formal instruction exists.

In the managerial stage, the task shifts to interpretation. Leaders must determine which patterns of usage represent genuine improvement and which are merely novelty. Governance becomes necessary, but governance without understanding tends to focus on compliance rather than value. Companies often spend this period creating policies while struggling to define outcomes.

In the organizational stage, the challenge becomes redesign. Processes are reconsidered, not merely accelerated. Some review steps

---

6. Fischer, Claude S. *America Calling: A Social History of the Telephone to 1940*. University of California Press, 1992.

7. MIT Sloan Management Review and Boston Consulting Group. *Achieving Individual and Organizational Value with AI*, 2023.

disappear because information can be validated differently. Some roles expand because individuals can operate across domains. Management becomes less about supervising activity and more about ensuring accountability for decisions informed by both human judgment and machine assistance.

A useful way to identify the current stage of a company is not to examine its technology stack but to examine its conversations.

In the first stage, conversations occur among employees about personal efficiency.

In the second stage, conversations occur among managers about use cases and vendors.

In the third stage, conversations occur among executives about structure and authority.

Many organizations remain in the second stage longer than they expect. They continue to add pilots and tools, assuming adoption will accumulate into transformation. It rarely does. Transformation requires a deliberate shift from asking where artificial intelligence can be inserted to asking which parts of the workflow are no longer necessary.

This distinction explains why companies sometimes feel stuck. They have invested, experimented, and trained employees, yet outcomes appear incremental. The problem is not effort. The problem is misdiagnosis. They believe they are moving along a technology curve when they are actually moving along an organizational one.<sup>8</sup>

The maturity curve is therefore less a roadmap than a recognition pattern. It does not prescribe a single sequence of actions. It helps leaders understand why their experiences feel contradictory. Employees report large productivity gains while executives see limited financial impact because the company is in the behavioral stage. Multiple pilots exist without integration because the company is in the managerial stage. Significant performance improvements appear only when workflows and decision making structures change because the company has entered the organizational stage.

The remainder of this book examines these stages in detail and

---

8. Stanford Institute for Human-Centered AI. *AI Index Report 2024*.

considers how they appear differently across types of firms. Small companies often move through the first stage rapidly because behavior spreads quickly. Mid sized firms often experience the second stage most intensely because they formalize practices while remaining flexible. Large enterprises experience the third stage most profoundly because coordinated redesign affects many parts of the organization at once.

The central point is simple but easy to overlook. Artificial intelligence does not create value when it is adopted. It creates value when the firm adapts.

Understanding that difference is the beginning of managing it.

## KEY TAKEAWAYS

- **Artificial intelligence adoption occurs in stages.** Organizations move from individual use, to managed pilots, to redesigned workflows. The technology remains similar across all three. What changes is how the company operates around it.
- **The stages reflect organizational behavior, not technical maturity.** Progress is not determined by model quality, vendor selection, or tool sophistication. It is determined by how decision-making and coordination change.
- **Different groups inside the same company experience different stages simultaneously.** Employees experience usefulness, managers experience inconsistency, and executives experience uncertainty because each is observing a different point on the adoption curve.
- **Most companies stall in the middle stage.** They deploy tools and run pilots successfully but fail to produce enterprise impact because workflows and authority structures remain unchanged.
- **AI creates value when the firm adapts, not when the tool is implemented.** Enterprise performance improves only after leaders redesign how work, roles, and decisions are organized.

# CHAPTER 7

## STAGE ONE: BEHAVIORAL ADOPTION

INDIVIDUALS CHANGE  
THEIR WORK BEFORE  
THE COMPANY NOTICES

THE FIRST STAGE of artificial intelligence adoption inside an organization does not look like a transformation. In fact, it rarely looks like anything at all.

There is no kickoff meeting. No executive sponsor. No project plan. No budget code.

Work simply starts getting done differently.<sup>1</sup>

An engineer pastes unfamiliar code into a model and asks for an explanation instead of scheduling a knowledge transfer meeting. A financial analyst summarizes a forty-page contract in minutes rather than hours. A product manager drafts a proposal, revises it, and iterates several alternatives before anyone else knows the work has begun.<sup>2</sup> A salesperson prepares outreach tailored to a client without waiting for marketing support.

Nothing about the organization's structure has changed. Job titles remain the same. Reporting relationships remain the same. The processes documented in operating manuals remain the same.

Yet the way work actually happens has already shifted.

---

1. Noy, Shakked, and Whitney Zhang. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, vol. 381, no. 6654, 2023, pp. 187–192.

2. Peng, Sida, et al. "The Impact of AI Pair Programming on Software Development Productivity." *arXiv preprint*, 2023.

*In 2023, Verizon discovered customer-support agents were already using generative AI tools to help draft responses to customers and interpret troubleshooting documentation before the company had formally implemented an AI support system. Management learned about the practice through internal observation rather than deployment planning. The productivity gains preceded the organizational decision.*<sup>3</sup>

This is the behavioral stage of adoption. Artificial intelligence enters the firm not as a system but as a habit.

Researchers have now begun documenting this phenomenon at scale. Microsoft's global workforce study found that a large majority of knowledge workers reported using generative AI in their daily tasks, and notably many began doing so on their own rather than through an employer rollout. The study also observed that employees frequently experimented first and informed their organizations later. The pattern was bottom-up rather than top-down.<sup>4</sup>

At the same time, leadership in many of those same organizations reported uncertainty about measurable business impact. McKinsey's global surveys similarly show broad employee usage across business functions while only a minority of firms report enterprise-level financial effect.<sup>5</sup>

Those two facts coexist because this stage is behavioral rather than organizational.

The company has not adopted artificial intelligence. The people inside it have.

## **WHY THE ORGANIZATION DOES NOT SEE IT**

Traditional corporate change is visible because it requires coordination.

---

3. Quach, Katyanna. "Verizon Finds Workers Quietly Using ChatGPT to Do Their Jobs." *The Register*, 2023.

4. Microsoft and LinkedIn. *Work Trend Index Annual Report: AI Is Changing How We Work*. Microsoft, 2024.

5. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

A new system must be installed. Employees must be trained. Procedures must be updated.

Artificial intelligence does not initially require coordination. It requires curiosity.

The cost of experimentation is extremely low. An employee can try a tool in minutes without asking permission. If the tool proves useful, the employee repeats the behavior. Over time, the employee adjusts how they approach problems. They write differently. They research differently. They prepare differently. The workflow inside the individual mind changes before the workflow inside the organization changes.

Economists studying general-purpose technologies have long noted that adoption often begins with individual improvement before organizational redesign. Erik Brynjolfsson and his colleagues describe a lag between technological capability and productivity statistics because organizations need time to reorganize around new methods of work.<sup>6</sup>

Artificial intelligence fits this pattern closely.

The firm measures outputs such as revenue, cost, and cycle time. Individuals experience cognitive effort. Artificial intelligence reduces effort first. The firm perceives improvement only when processes change.

During this stage, employees often hesitate to publicize their usage. Professional norms historically equate effort with competence. If a task once required a day of work and now requires an hour, workers sometimes worry that revealing assistance will be interpreted as cutting corners. Deloitte's enterprise AI adoption research notes that many organizations experience unofficial use of AI tools before policies are formalized, precisely because employees act before management frameworks exist.<sup>7</sup>

As a result, organizations experience a peculiar condition. Work improves quietly.

---

6. Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. "Generative AI at Work." National Bureau of Economic Research Working Paper No. 31161, 2023.

7. Deloitte. *State of Generative AI in the Enterprise: Now Decides Next*. Deloitte Insights, 2024.

## WHAT ACTUALLY CHANGES IN STAGE ONE

The change is not speed alone. It is cognitive workflow.

A knowledge worker normally moves through a sequence of thinking steps. They gather information, interpret it, outline an argument, and produce output. Artificial intelligence inserts a conversational partner into that process. The worker begins thinking interactively rather than sequentially. They test ideas earlier. They revise more frequently. They check assumptions continuously.

In effect, the employee now works with a form of augmented reasoning.

This matters because organizations are structured around assumptions about how long thinking takes. Meetings are scheduled to align understanding. Review cycles exist because preparation requires time. When preparation time collapses but the structure remains, the organization becomes misaligned with how work is actually being performed.

The consequences appear subtly at first. Certain individuals respond faster than expected. Some teams become unusually productive. Some employees appear able to operate across multiple domains. Managers interpret these as performance differences rather than structural change.

But they are observing the earliest signal that the firm's operating assumptions are becoming outdated.

## HOW STAGE ONE APPEARS IN DIFFERENT FIRMS

Although the behavior is similar everywhere, its visibility differs depending on organizational size.

### *Small Businesses*

In a small firm, behavioral adoption spreads almost immediately. The owner or a small leadership group performs many functions personally. When one person changes how they work, the entire firm effectively changes with them.

A consultant drafts proposals in half the time. A retailer analyzes inventory independently. A founder conducts market research without

hiring an external firm. The organization gains capability disproportionate to its headcount because cognitive tasks that previously required specialization can now be performed by a generalist using assistance.

The improvement is obvious because there is little hierarchy to buffer it. The business owner directly experiences the gain.

#### *Mid-Market Firms*

In a mid-sized organization, behavioral adoption is uneven. Some departments experiment rapidly while others do not. Performance variation becomes noticeable.

One team produces reports faster. Another continues operating normally. Managers observe inconsistent productivity and often attribute it to employee skill rather than differing workflows.

This stage can create internal tension. Employees who adopt early appear unusually effective, but the organization has not yet provided guidance on how work should be performed. Leadership senses improvement but cannot yet standardize it.

The firm is learning that change is occurring without yet knowing how to manage it.

#### *Enterprise Organizations*

In large enterprises, behavioral adoption is widespread but largely invisible. Thousands of employees experiment independently. Because outputs are aggregated across many layers, the change is difficult to isolate.

Security and compliance concerns also appear early in large organizations. Employees input internal documents into tools to obtain summaries or analysis. Leadership becomes aware not through productivity metrics but through risk discussions.

Research supports this pattern. Surveys consistently show that large enterprises report both significant experimentation and significant uncertainty simultaneously.<sup>8</sup> The organization experiences usage before it develops governance.

---

8. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

Here the firm's scale obscures the signal. Improvement exists, but no single manager can easily see it.

## THE FIRST REAL RISK

The risk in this stage is not technological failure. It is unmanaged learning.

Employees develop new workflows individually. Knowledge about effective usage remains local rather than shared. Some workers dramatically increase effectiveness while others do not. The organization accumulates performance differences without understanding their cause.

At the same time, policy often lags behavior. Leaders may respond by restricting use in order to manage perceived risk. However, banning behavior that provides clear benefit tends to drive it underground rather than eliminate it.

The more constructive response is observation. Leaders who ask how employees are actually completing work often learn more than those who begin by defining rules. Stage One is a discovery phase whether the organization acknowledges it or not.

## WHY STAGE ONE MATTERS

It is tempting to dismiss this stage because it produces few visible metrics. That would be a mistake.

Stage One is where the firm first encounters a new operating reality. Workers are no longer limited by the cognitive throughput assumed by the organization's structure. The company's processes still reflect an older model of work while its employees increasingly operate within a new one.

This mismatch explains why later stages become necessary. The managerial and organizational changes that follow are responses to a shift that has already occurred at the behavioral level.

Artificial intelligence adoption does not begin when leadership announces a program.

It begins when employees quietly change how they think.

## KEY TAKEAWAYS

- **The first phase of AI adoption is invisible.** There is no rollout or announcement. Employees independently change how they research, write, analyze, and prepare decisions long before leadership formally acts.
- **The organization has not adopted AI yet, but its people have.** Individual behavior shifts precede policy, governance, and strategy, which is why leadership often notices risk discussions before it notices productivity gains.
- **Early value appears as reduced cognitive effort, not improved business metrics.** Employees complete thinking tasks faster and more iteratively, but enterprise performance indicators lag because workflows and structures remain unchanged.
- **Uneven performance is a signal, not a coincidence.** Large differences in productivity between individuals or teams often reflect different working methods rather than differences in talent.
- **The leadership task in Stage One is discovery.** Attempting to control or prohibit usage too early drives behavior underground and prevents leaders from learning where the technology actually creates value.



# CHAPTER 8

## STAGE ONE RISKS

### UNMANAGED ADOPTION CREATES HIDDEN OPERATIONAL EXPOSURE

THE FIRST STAGE of artificial intelligence adoption feels productive. Employees work faster. Documents improve. Analysis appears more complete. From the perspective of the individual, the experience is almost entirely positive.

From the perspective of the organization, however, Stage One introduces a different kind of challenge.

The risk is not that artificial intelligence does not work. The risk is that it works without the firm understanding how it is being used.

*When USB flash drives became popular, employees quickly adopted them to move files between home and office computers and to share documents across teams. Organizations often discovered their widespread use only after sensitive data had been copied outside controlled systems. The technology functioned exactly as intended, but management had not yet adapted its security practices.<sup>1</sup>*

When behavioral adoption spreads quietly, organizations accumulate hidden operational exposure. The firm begins to rely on workflows

---

1. 1 Schneier, Bruce. *Secrets and Lies: Digital Security in a Networked World*. Wiley, 2004.

it has not designed, approved, or examined. The technology does not malfunction. The management system lags.

To understand the exposure, it is useful to examine what actually changes in Stage One.

Individuals are no longer producing work through exclusively human reasoning. They are collaborating with a model. That collaboration affects not only speed but method. Employees begin to:

- Input internal documents into external systems
- Rely on generated summaries rather than original sources
- Draft analysis based on synthesized patterns rather than raw data
- Share outputs that have not passed through traditional review steps

Each of these behaviors can be rational and efficient. Each also introduces risk when unmanaged.

Surveys of enterprise leaders reflect this tension. In Deloitte's research on generative AI adoption, executives frequently report concern about data leakage and compliance even as employees report daily use of AI tools in their workflows.<sup>2</sup> Similarly, Gartner has warned that many organizations underestimate how quickly generative AI use spreads beyond formal policy, particularly when employees perceive clear productivity benefits.<sup>3</sup>

The issue is not malicious intent. It is organizational invisibility.

## THE DATA EXPOSURE PROBLEM

The most obvious risk in Stage One concerns data. Employees experimenting independently may input confidential materials into tools that are not formally approved or monitored. Contracts, customer commu-

---

2. Deloitte. *State of Generative AI in the Enterprise*. Deloitte Insights, 2024.

3. Gartner, Inc. *Emerging Tech: Generative AI Governance and Risk Management*. Gartner Research, 2023.

nications, product documentation, and financial analysis can all become part of informal experimentation.

Even when vendors provide enterprise-grade assurances, usage often begins in consumer versions before enterprise agreements exist. Microsoft's workforce research has documented that many employees begin experimenting with publicly available AI systems before their organizations formalize internal deployments.<sup>4</sup>

In small firms, this exposure may be manageable simply because data volume is limited and oversight is direct. In large enterprises, the scale of experimentation increases the surface area of risk dramatically. Thousands of independent users represent thousands of potential policy variations.

The challenge is compounded by the fact that the productivity benefit is real. Employees are not acting carelessly. They are acting pragmatically. The organization therefore faces a dilemma: restrict behavior and lose efficiency, or allow behavior and absorb unmanaged exposure.

The solution is not immediate prohibition. It is visibility.

## THE QUALITY DRIFT PROBLEM

A second risk is subtler and more structural.

Artificial intelligence produces fluent output. That fluency can create an illusion of completeness. Employees may begin to rely on summaries or explanations without verifying underlying reasoning. Over time, the organization may experience what could be called quality drift.

Quality drift does not occur because the model is consistently incorrect. It occurs because verification habits weaken. When analysis appears polished and confident, managers may review less deeply. When documents are produced quickly, assumptions may be embedded that previously would have been surfaced through slower drafting processes.

---

4. Microsoft. *Work Trend Index Annual Report: AI at Work Is Here. Now Comes the Hard Part*. Microsoft, 2024.

Academic research on automation bias has long shown that humans tend to over-rely on automated systems when those systems produce confident outputs.<sup>5</sup> Generative AI introduces similar psychological dynamics into knowledge work.

In Stage One, each employee develops their own verification standards. Some are rigorous. Others are not. The firm therefore experiences variation in reliability across teams without realizing that the variation correlates with AI usage patterns.

The organization has not yet established norms for human oversight of machine-generated reasoning. That absence becomes more significant as usage increases.

## THE SHADOW DEPENDENCY PROBLEM

A third exposure emerges gradually.

As employees integrate artificial intelligence into their daily routines, they may begin structuring work around the assumption that the tool is available. Drafting, research, coding, and analysis become interactive rather than sequential. Over time, individuals build cognitive habits that depend on rapid machine feedback.

This is not inherently negative. In many cases it represents improved workflow. The risk arises when the dependency is informal and unsupported.

If a tool becomes temporarily unavailable, productivity may decline unexpectedly. If a vendor changes terms or pricing, the firm may discover that critical workflows depend on a system leadership never formally adopted. If models behave differently after updates, teams may experience confusion about output consistency.

In effect, the organization may become reliant on a system it has not officially integrated.

McKinsey's research on scaling AI highlights the importance of integrating systems into core workflows precisely because informal

---

5. Parasuraman, Raja, and Victor Riley. "Humans and Automation: Use, Misuse, Disuse, Abuse." *Human Factors*, vol. 39, no. 2, 1997, pp. 230–253.

experimentation, while valuable, does not create institutional stability.<sup>6</sup>

Stage One creates experimentation. Without progression to later stages, experimentation hardens into unexamined dependency.

## HOW RISK APPEARS ACROSS FIRM SIZES

The nature of Stage One exposure differs depending on organizational structure.

### *Small Businesses*

In small firms, risk is concentrated. The same individual often controls experimentation, oversight, and execution. Data exposure risk exists but is easier to monitor because fewer actors are involved. Quality drift may be less severe because work remains visible to the owner.

The more significant risk for small firms is over-reliance. A founder may begin structuring the business around personal augmented productivity. If growth occurs without formalizing workflows, scaling may prove difficult because knowledge remains embedded in individual-machine interactions rather than shared processes.

### *Mid-Market Firms*

In mid-sized companies, variation becomes the central risk. Some departments adopt aggressively. Others lag. Without coordinated visibility, performance differences emerge. Managers may reward outcomes without understanding process changes.

This creates uneven governance. One team may follow strict review practices for AI-generated content while another does not. Over time, inconsistency can undermine standardization and compliance expectations.

Mid-market firms often have enough structure to require policy but not enough oversight to monitor behavior comprehensively. The result is fragmentation.

### *Enterprise Organizations*

In large enterprises, exposure scales with adoption. Thousands of

---

6. McKinsey & Company. *Scaling AI Across the Enterprise*. McKinsey, 2024.

independent behavioral experiments occur simultaneously. Data governance concerns intensify. Legal and compliance functions become involved before leadership fully understands the scope of usage.

Enterprise firms frequently respond by accelerating formal governance efforts. However, governance frameworks designed for traditional software often struggle to keep pace with behavioral adoption.

Gartner has predicted that organizations failing to implement structured AI governance will face significant risk exposure, particularly related to data privacy and compliance.<sup>7</sup> The concern is not hypothetical. Regulatory scrutiny around data usage and algorithmic accountability is increasing globally.

In enterprises, unmanaged Stage One adoption can produce both compliance risk and cultural friction.

## THE LEADERSHIP RESPONSE

It is important to emphasize that Stage One risk does not imply that behavioral adoption is undesirable. On the contrary, this stage reveals where artificial intelligence creates genuine value. Suppressing it entirely would eliminate valuable experimentation.

The appropriate leadership response is not restriction but transition.

Leaders must first acknowledge that adoption has already begun. They must then create channels through which employees can surface how they are using AI effectively. Instead of asking whether AI should be allowed, leaders should ask how work is already being performed differently.

Organizations that progress successfully through Stage One typically take three steps, though not always consciously.

They observe behavior before prescribing policy.

They establish minimum guardrails rather than comprehensive restrictions.

They begin identifying patterns worth formalizing.

---

7. Gartner, Inc. *Emerging Tech: Generative AI Governance and Risk Management*. Gartner Research, 2023.

Research on digital transformation consistently shows that bottom-up innovation combined with top-down integration produces stronger results than either approach alone.<sup>8</sup>

Stage One is not a failure of governance. It is a precondition for informed governance.

## WHY THIS STAGE MATTERS FOR THE MATURITY CURVE

The risks of Stage One are not signs that artificial intelligence should be avoided. They are signals that the firm has entered a new coordination environment without yet redesigning its management systems.

*In 2023, public school districts across the United States discovered teachers were already using generative AI to draft lesson plans, write parent communications, summarize student assessments, and prepare classroom materials. The systems had improved productivity immediately, but institutional guidance, privacy controls, and evaluation standards had not yet been defined.*<sup>9</sup>

If left unmanaged, behavioral adoption produces exposure. If observed and integrated, it provides insight into where structural redesign is warranted.

The transition to the next stage occurs when leaders recognize that experimentation is widespread and begin formalizing it through structured initiatives. The firm moves from individuals changing how they work to managers attempting to guide how work should change.

The tension that emerges in Stage One is therefore productive. It reveals the distance between how employees operate and how the organization is structured.

Artificial intelligence does not introduce risk because it fails.

It introduces risk because it works faster than management adapts.

---

8. Fountaine, Tim, Brian McCarthy, and Tamim Saleh.

“Artificial Intelligence in Business Gets Real.” *MIT Sloan Management Review*, vol. 60, no. 4, Summer 2019, in collaboration with Boston Consulting Group.

9. Huang, Kalley. “Teachers Are Using A.I. to Grade Papers and Write Lesson Plans.” *The New York Times*, 14 Aug. 2023.

## KEY TAKEAWAYS

- **Stage One risk is not technical failure. It is invisible dependency.** Employees are incorporating AI into daily workflows before governance, policy, or visibility exist.
- **Data exposure increases before leadership awareness does.** Confidential information, draft analyses, and internal materials may already be moving through tools that the organization has not formally approved.
- **Quality variation becomes structural.** Without shared verification standards, teams develop inconsistent oversight practices, leading to uneven reliability across the firm.
- **Productivity gains can mask operational fragility.** Informal reliance on AI tools may create hidden dependencies that surface only when systems change, fail, or become restricted.
- **The correct response is visibility, not prohibition.** Leaders who observe and channel behavioral adoption into structured guardrails reduce risk while preserving the experimentation that reveals real value.

# CHAPTER 9

## STAGE TWO: MANAGERIAL ADOPTION

### COMPANIES DEPLOY AI BUT ONLY ACCELERATE TASKS

THERE IS A MOMENT, usually subtle, when leadership realizes artificial intelligence is no longer an individual curiosity.

At first it appears as scattered anecdotes. A department head mentions that their team is drafting reports faster. A manager observes that a junior employee is suddenly capable of handling work that previously required senior review. A compliance officer asks whether employees are entering company information into external systems. A vendor reaches out offering an enterprise contract rather than a demonstration.

Individually, these signals are small. Collectively, they force recognition. The organization concludes that artificial intelligence is happening whether it formally chooses it or not.

This recognition marks the beginning of the second stage.

The company decides to act deliberately.

An initiative forms. The firm appoints sponsors. Policies are drafted. Procurement becomes involved. A cross-functional group is created to evaluate tools and coordinate usage. What began as behavior becomes a program.

Leaders often interpret this moment as the beginning of adoption. In reality, it is the beginning of management.

In the first stage, individuals altered their personal workflows. In

the second stage, the firm attempts to govern those workflows. The organization does not yet change how it operates. It changes how it supervises the technology.

This distinction explains why the second stage feels busy but rarely feels transformative.

## THE ORGANIZATION'S FIRST INSTINCT: THE USE CASE LIST

Once leadership acknowledges artificial intelligence, the natural first question is practical.

Where should we use it?

The organization answers by collecting use cases. Business units are asked to identify opportunities. Proposals appear quickly. Customer service wants automated call notes. Marketing wants content generation. Finance wants document classification. Engineering wants code assistance. Human resources wants drafting support for internal communications.

The list grows rapidly because knowledge work contains many bounded activities. Artificial intelligence is particularly effective within a bounded activity. Summarization, drafting, categorization, and pattern identification fit the capabilities well.

The firm evaluates several options and launches pilots. Each pilot is assigned an owner, a timeline, and a set of expected benefits. The company has now moved from curiosity to experimentation.

*Morgan Stanley began testing a generative AI assistant for financial advisors using OpenAI technology. Advisors used the system to search internal research and draft client communications. Individual teams reported productivity improvements and faster access to information, leading the firm to expand the pilot across thousands of employees. The early value appeared within specific tasks rather than as an immediate firm-wide performance shift.<sup>1</sup>*

---

1. Son, Hugh. "Morgan Stanley Is Rolling Out a ChatGPT-Powered Assistant to Its Financial Advisors." *CNBC*, 19 Sept. 2023.

The results are usually encouraging.

Customer service representatives complete documentation more quickly. Marketing teams produce multiple campaign drafts in the time previously required for one. Developers navigate unfamiliar codebases with assistance. Analysts synthesize research efficiently. Employees report tangible productivity gains.

Empirical research supports this experience. In controlled studies, generative AI tools significantly increased task productivity for many knowledge workers, particularly in writing, research, and software development activities.<sup>2</sup>

From the perspective of the department running the pilot, the outcome is clear. The technology works.

From the perspective of the firm as a whole, the picture is less clear.

## WHY SUCCESS STILL FEELS INCOMPLETE

Executives at this stage often express a specific frustration. They have evidence of improvement yet cannot point to a corresponding change in overall company performance.<sup>3</sup> The organization is doing more projects related to artificial intelligence than ever before, yet operating metrics at the enterprise level move slowly.

The explanation lies in the difference between improving a task and changing a workflow.

A task is a component activity. A workflow is a sequence of coordinated decisions. Artificial intelligence, at this stage, primarily improves the former. It rarely alters the latter.

A team that drafts reports faster still presents them in the same meeting. A group that produces analysis more quickly still waits for the same approvals. A developer who writes code faster still follows the same release process. The company improves preparation but not coordination.

Management research has long shown that organizational perfor-

---

2. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. "Generative AI at Work." *Quarterly Journal of Economics*, 2024.

3. Fountaine, Tim, Brian McCarthy, and Tamim Saleh "Artificial Intelligence in Business Gets Real." *MIT Sloan Management Review*, Summer 2019.

mance depends less on the efficiency of individual steps than on the structure connecting them. When information flows through multiple review layers, accelerating one layer does not remove the delay created by the system itself. Artificial intelligence reduces effort required to create information, but it does not automatically change how organizations decide what to do with that information.

The firm has increased activity inside the existing structure. The structure itself remains intact.

This explains the paradox leaders experience. Artificial intelligence demonstrably improves work while organizational performance improves only modestly.

The company has adopted tools. It has not yet redesigned work.

## THE GOVERNANCE PHASE

As pilots multiply, another development occurs. Legal, security, and compliance functions become involved.

Their concerns are reasonable. Artificial intelligence introduces new questions about data confidentiality, intellectual property, and accountability. If employees are using models to analyze documents, where is the information stored? If a generated output contains an error, who is responsible? If a system suggests a decision, how should oversight occur?

Organizations therefore begin constructing governance frameworks. Approved vendor lists appear. Data handling policies are created. Employees receive training about appropriate usage.

Industry research consistently highlights governance as a central issue during this stage. Gartner has emphasized that organizations rapidly expanding generative AI usage must address accountability and risk management structures to avoid exposure related to data handling and regulatory compliance.<sup>4</sup> Deloitte's enterprise studies

---

4. Gartner, Inc. *Predicts 2024: Generative AI Governance and Risk Management Will Become a Board-Level Priority*. Gartner Research, 2024.

similarly note that formal policies often follow widespread employee experimentation rather than precede it.<sup>5</sup>

Governance is necessary. However, governance alone does not produce value. It produces stability.

Many organizations spend this stage becoming better at controlling artificial intelligence than at benefiting from it.

## HOW STAGE TWO APPEARS BY FIRM SIZE

Although the second stage occurs across organizations, its character differs depending on scale.

### *Small Businesses*

Small businesses tend to move through the managerial stage quickly. Decision authority is concentrated. A tool that proves useful becomes part of daily operations almost immediately. The distinction between pilot and process blurs. The firm formalizes adoption informally. Its advantage is speed. Its limitation is consistency. It may adopt multiple overlapping tools without evaluating long-term integration.

### *Mid-Market Firms*

Mid-market firms experience the stage most vividly. They possess enough structure to require coordination but remain flexible enough to experiment widely. Multiple departments run pilots simultaneously. Leadership compares outcomes and attempts to standardize successful practices. Many mid-sized organizations begin to see measurable return on investment because workflows are structured enough to observe efficiency gains yet not so rigid as to prevent change.

### *Enterprise Organizations*

Large enterprises often remain in this stage longest. Their complexity requires careful evaluation before organization-wide deployment. Governance processes expand. Vendor negotiations lengthen. Dozens of initiatives operate concurrently across business units. Each produces localized improvement, but the enterprise as a whole changes slowly. Surveys repeatedly show large organizations

---

5. Deloitte. *The State of Generative AI in the Enterprise: Now Decides Next*. Deloitte Insights, 2024.

investing heavily in artificial intelligence while still seeking scalable operational impact.<sup>6</sup>

The difference is not technological sophistication. It is coordination complexity.

## WHAT THE ORGANIZATION LEARNS

The managerial stage performs a critical function even though it can feel frustrating. It teaches the firm two things simultaneously.

First, it confirms that artificial intelligence meaningfully improves cognitive tasks.

Second, it reveals that improving tasks alone does not transform the enterprise.

The organization begins to understand that the challenge is no longer technological adoption. It is organizational adaptation.

The key question gradually shifts. Leaders stop asking which use cases to deploy and begin asking why certain steps exist at all. They notice that reviews occur because preparation once required time. They notice that approval chains exist because information once traveled slowly. They realize many coordination mechanisms were designed for a different information environment.

This realization marks the beginning of the next stage.

The company moves from managing artificial intelligence to reorganizing around it.

The tools remain the same.

The firm does not.

## KEY TAKEAWAYS

- **Stage Two begins when leadership intervenes.** AI shifts from individual experimentation to a managed program with sponsors, pilots, policies, and vendor evaluations.
- **Pilots improve tasks, not organizations.** Teams complete

---

6. McKinsey & Company. *The State of AI in 2024: Generative AI's Breakout Year*. McKinsey Global Survey on AI, 2024.

work faster and better, yet enterprise performance changes little because workflows and decision processes remain unchanged.

- **Governance stabilizes usage but does not create value.** Security, legal, and compliance involvement is necessary, but oversight alone cannot produce meaningful business impact.
- **The “AI frustration” executives feel is structural.** The company is improving preparation while keeping the same meetings, approvals, and coordination mechanisms.
- **The real lesson of Stage Two is diagnostic.** The organization learns that the barrier is no longer technology capability but how work is organized, which sets up the need for workflow redesign.



# CHAPTER 10

## STAGE TWO RISKS

### THE PILOT PLATEAU — BUSY, IMPRESSIVE, AND LARGELY WORTHLESS

THERE IS a stage in nearly every organization's artificial intelligence journey that produces a very specific emotional reaction among executives. It is not excitement, and it is not skepticism. It is unease.

The company has clearly invested effort. Teams are busy. Demonstrations are impressive. Multiple departments can point to successful implementations. Employees speak positively about the tools. The organization has created governance structures, assigned responsibility, and funded experimentation. From the outside, it appears that meaningful progress is underway.

Yet inside the leadership team a persistent question remains unanswered.

If all of this activity is real, why does the company still feel fundamentally the same?<sup>1</sup>

This moment is not unusual. It is a predictable phase in adoption that I refer to as the pilot plateau. It is the point at which a company has acknowledged the usefulness of artificial intelligence, has organized around it, and has deployed it across many functions, but has not yet changed the structure through which work flows. The organi-

---

1. Accenture. *Reinventing Enterprise Models in the Age of Generative AI*. Accenture Technology Vision 2024.

zation is managing artificial intelligence without yet operating differently because of it.

Understanding this stage matters because many companies misinterpret it. Some conclude the technology has been oversold. Others double down on experimentation, assuming more pilots will eventually produce transformation. In reality, the plateau exists not because the technology lacks capability but because the firm has not yet altered coordination.

Artificial intelligence improves tasks quickly. Organizations change slowly. The plateau is the distance between those two speeds.

## THE SCALING PROBLEM

The first difficulty leaders encounter is scaling impact rather than scaling usage.

Adding new pilots is easy. Integrating their results into the performance of the entire company is not.

In most organizations, a pilot begins within a functional boundary. Customer service deploys automated call summaries. Marketing introduces assisted content creation. Engineering adopts code generation assistance. Legal tests document review. Finance experiments with automated classification. Each project has a clear owner and a contained scope. Each produces measurable improvement within its domain.

The problem emerges when executives attempt to connect those improvements to enterprise performance.

Companies are systems of coordinated decisions. They do not operate as independent departments. A product launch requires engineering, marketing, operations, finance, and leadership approval. A pricing decision involves analytics, sales, and finance. A customer issue may touch support, operations, and legal review. The speed of the organization depends not on how quickly one department performs its tasks but on how quickly decisions move across departments.

Artificial intelligence, at this stage, accelerates preparation. It rarely accelerates coordination.

A marketing team can now generate campaign drafts in hours rather than days. But the campaign still moves through the same review meetings, approval committees, and scheduling cycles. A developer can write code faster, yet release management still follows the same validation gates. Analysts prepare research rapidly, yet executives still receive it at the next scheduled meeting.

The result is predictable. Local efficiency improves while system performance remains largely unchanged.

Empirical research reflects this pattern. Studies of enterprise AI adoption repeatedly show organizations achieving value within functions but struggling to translate those gains into enterprise-level financial performance. McKinsey's work on scaling artificial intelligence finds that only a minority of firms realize significant bottom-line impact even after deploying multiple use cases, largely because integration across workflows is limited.<sup>2</sup>

Leaders often expect transformation to emerge through accumulation. They believe ten successful pilots will produce ten times the impact of one. In practice, ten pilots often produce ten separate improvements that never converge into a different operating speed.

The company has scaled experimentation.

It has not scaled decision making.

After acquiring Kiva Systems, Amazon introduced thousands of warehouse robots to assist fulfillment operations. Individual tasks such as item retrieval became faster and more accurate. However, early deployments did not immediately produce proportional gains in overall fulfillment speed because surrounding processes and layouts still governed throughput. Larger improvements appeared only after Amazon redesigned workflows and facility organization around the new capability.<sup>3</sup>

---

2. McKinsey & Company. *Scaling AI Across the Enterprise*. McKinsey & Company, 2024.

3. Wurman, Peter R., et al. "Coordinating Hundreds of Cooperative, Autonomous Vehicles in Warehouses." *AI Magazine*, vol. 29, no. 1, 2008, pp. 9–20.

## THE MEASUREMENT PROBLEM

The next obstacle appears when organizations attempt to measure return on investment.

Artificial intelligence primarily reduces effort before it increases output. Modern accounting systems, however, measure output rather than avoided work. This mismatch creates persistent ambiguity.

Consider a team that completes analytical preparation in one-quarter of the previous time. If staffing levels remain unchanged and the process surrounding the analysis does not accelerate, financial records show little difference. Employees experience relief. Managers observe efficiency. Financial statements remain stable.

Economists have long documented this phenomenon. Erik Brynjolfsson described how general-purpose technologies often produce a lag between adoption and measurable productivity because complementary organizational changes must occur before benefits appear in performance statistics.<sup>4</sup> Artificial intelligence amplifies the issue because it targets cognitive effort rather than physical production.

Organizations therefore encounter a paradox. Workers report meaningful productivity gains, yet executives cannot demonstrate equivalent financial impact.<sup>5</sup> Some firms respond by imposing strict ROI requirements before expanding initiatives. Others continue funding projects in anticipation of future payoff. Neither approach resolves the core issue because the measurement problem reflects structure rather than investment.

Artificial intelligence shortens thinking cycles. Companies still operate on decision cycles.

Until the latter changes, accounting evidence remains incomplete.

## TOOL PROLIFERATION

During this period, experimentation continues. Each department seeks

---

4. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. "Generative AI at Work." *Quarterly Journal of Economics*, 2024.

5. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. "Generative AI at Work." *Quarterly Journal of Economics*, 2024.

solutions appropriate to its responsibilities. Vendors respond enthusiastically, offering specialized applications across nearly every business activity.

The organization accumulates tools rapidly.<sup>6</sup>

Marketing uses one system to generate content. Sales uses another for outreach preparation. Engineering uses a coding assistant. Legal adopts a contract analysis platform. Human resources experiments with policy drafting support. Data teams implement analytical copilots.

Each tool works. Together they introduce complexity.

Capabilities overlap. Data flows across multiple platforms. Employees navigate multiple interfaces. Integration challenges emerge. The firm begins managing an ecosystem of intelligence rather than a unified operating environment.

Industry analysts have begun describing this phenomenon as application proliferation. Gartner has warned that uncontrolled expansion of generative AI tools can increase governance and operational complexity faster than it produces enterprise value if not coordinated.<sup>7</sup>

The paradox becomes clear. Artificial intelligence promises simplicity in work execution, yet without workflow redesign it can increase complexity in system management.

This outcome does not arise because the tools are ineffective. It arises because each group optimizes locally rather than the organization optimizing collectively.

## VENDOR CONFUSION

At the same time, leaders confront a second challenge: understanding the technology landscape itself.

Artificial intelligence evolves faster than most enterprise software

---

6. Gartner, Inc. *Predicts 2025: Generative AI Will Require Platform Governance to Avoid Operational Fragmentation*. Gartner Research, 2024.

7. Gartner, Inc. *Predicts 2025: Generative AI Will Require Platform Governance to Avoid Operational Fragmentation*. Gartner Research, 2024.

categories.<sup>8</sup> Capabilities expand continuously. New providers emerge. Established providers reposition themselves. Demonstrations are persuasive and often accurate, but each emphasizes a different dimension of value.

Organizations therefore attempt to make a strategic choice prematurely. They hope selecting a single platform will clarify direction. Instead, the decision often becomes obsolete as capabilities shift.

Innovation research has long shown that during periods of rapid technological evolution, firms struggle not because solutions are absent but because categories are unstable.<sup>9</sup> Artificial intelligence currently exhibits this instability. The question “Which vendor should we choose?” appears strategic but often distracts from the deeper question “What role should this technology play in how we operate?”

Companies on the plateau invest significant effort evaluating providers while postponing structural decisions.

Technology selection feels like progress because it is concrete. Organizational redesign feels risky because it changes accountability.

## ORGANIZATIONAL INERTIA

The most powerful reason companies remain on the plateau is not technological uncertainty but institutional habit.

Organizations encode trust through process. Review steps, approval layers, and reporting structures exist not only to coordinate work but to distribute responsibility. A document reviewed by multiple managers shares accountability. A decision approved by committee distributes risk.

Artificial intelligence challenges these mechanisms indirectly. If analysis can be produced quickly and validated differently, some steps become less necessary. Removing them, however, alters who is responsible for outcomes.<sup>10</sup>

---

8. Sevilla, Jaime, et al. “Compute Trends Across Three Eras of Machine Learning.” *International Conference on Machine Learning (ICML)*, 2022.

9. Rogers, Everett M. *Diffusion of Innovations*. 5th ed., Free Press, 2003.

10. Fountaine, Tim, Brian McCarthy, and Tamim Saleh. “Artificial Intelligence in Business Gets Real.” *MIT Sloan Management Review*, Summer 2019.

Research on organizational behavior consistently shows that firms resist altering coordination mechanisms because those mechanisms define authority (Nelson and Winter, evolutionary theories of organizational change).<sup>11</sup> Leaders may recognize inefficiency yet hesitate to remove structures that manage risk and responsibility.

The pilot stage allows companies to avoid this decision. Pilots improve activity without redefining authority. They feel safe precisely because they do not require structural commitment.

Thus the plateau persists.

## WHY COMPANIES REMAIN

The pilot plateau is comfortable. The organization appears innovative without confronting disruptive internal change. Leadership can report progress to boards and stakeholders. Employees gain useful tools. Risk appears controlled.

What does not occur is transformation.

The firm is waiting for clearer evidence, yet the evidence it seeks cannot appear until it changes how it operates. The plateau ends not when technology improves but when leadership recognizes that the limiting factor is no longer the tool.

It is the firm.

The organizations that move beyond this stage do something difficult. They stop adding pilots and begin redesigning workflows. They remove approval steps designed for slower information environments. They redefine roles around decision ownership rather than process supervision. They accept that artificial intelligence changes not only what employees can do but how the organization must function.

At that moment the plateau ends.

Artificial intelligence ceases to be an initiative.

It becomes part of the operating model.

Transformation does not begin when the company experiments.

It begins when the company reorganizes.

---

11. Nelson, Richard R., and Sidney G. Winter. *An Evolutionary Theory of Economic Change*. Belknap Press of Harvard University Press, 1982.

## KEY TAKEAWAYS

- **The pilot plateau is a predictable phase.** Organizations can run many successful AI initiatives and still feel unchanged because work preparation has accelerated while decision coordination has not.
- **Scaling usage is not the same as scaling impact.** Departments improve locally, but enterprise performance depends on how quickly decisions move across the company.
- **ROI appears elusive because structure has not changed.** AI reduces effort first, and financial performance improves only after workflows and decision cycles are redesigned.
- **More tools and more pilots often increase complexity.** Without coordination, organizations accumulate overlapping systems, governance burdens, and vendor confusion rather than transformation.
- **The barrier is organizational inertia, not technical capability.** Companies remain stuck because changing workflows also changes authority, accountability, and risk ownership.

# CHAPTER 11

## STAGE THREE: ORGANIZATIONAL ADOPTION

### WORKFLOWS, ROLES, AND DECISION RIGHTS CHANGE

THERE IS a moment that arrives quietly inside companies that successfully progress beyond experimentation.

It is not the moment they install a platform. It is not the moment they train employees. It is not even the moment they deploy an impressive system.

It is the moment a leader asks a different question.

Instead of asking, “Where can we use artificial intelligence?”

They begin asking, “Why do we still do this step at all?”

That question marks the beginning of the third stage of adoption.

Up to this point the organization has treated artificial intelligence as a tool added to existing processes. In the third stage, the firm begins to reconsider the processes themselves. The company stops inserting AI into workflows and instead redesigns workflows around the presence of AI.

*In 2024, Klarna deployed an AI customer service assistant capable of handling support conversations. Rather than simply assisting agents, the company reorganized support workflows so many customer inquiries no longer required human handling at all. The change reduced resolution time dramatically. The value came not from adding a tool but from removing steps in the service process.<sup>1</sup>*

This is where measurable enterprise value finally appears.<sup>2</sup>

## **WHY THE FIRST TWO STAGES CANNOT PRODUCE FULL VALUE**

To understand the third stage, it helps to clarify why the first two stages, despite clear activity, rarely produce transformation.

In the behavioral stage, individuals improve how they perform tasks. Productivity exists but remains private. The firm's operating model does not change, so organizational performance barely moves.

In the managerial stage, companies coordinate and govern usage. Pilots scale across departments. Task-level efficiency increases. However, the structure connecting tasks remains intact. The firm supervises faster work but still manages the same sequence of approvals, reviews, and handoffs.

Both stages improve execution within the existing system. Neither changes the system.

Economists studying general-purpose technologies have repeatedly observed a delay between technological capability and productivity growth. When electricity replaced steam power, factories initially installed electric motors but retained layouts designed for central mechanical drives. Only after companies reorganized production lines around distributed power did productivity significantly improve.<sup>3</sup> The technology enabled change, but value appeared only after redesign.

---

1. Klarna. "Klarna's AI Assistant Handles the Equivalent Work of 700 Full-Time Agents." Klarna Press Release, Feb. 2024.

2. Ransbotham, Sam, et al. *Achieving Individual — and Organizational — Value with AI*. MIT Sloan Management Review and Boston Consulting Group, 2023.

3. David, Paul A. "The Dynamo and the Computer: An Historical Perspective on the

Artificial intelligence follows the same pattern.<sup>4</sup> As long as the organization preserves workflows designed for slow information processing, the gains remain localized.

The third stage begins when leadership accepts that the technology is not merely assisting work. It is redefining the assumptions underlying work.

## WHAT ACTUALLY CHANGES

In the third stage, companies alter coordination.

Historically, organizations built layered structures because producing reliable information required multiple steps.<sup>5</sup> An analyst gathered data. A manager interpreted it. A director validated it. A senior executive decided. Each level existed partly to process uncertainty.

Artificial intelligence reduces the effort required to generate, interpret, and cross-check information. As a result, the purpose of certain intermediate steps changes. Some reviews exist no longer because human judgment is unnecessary but because prior information constraints required them.

Companies in the third stage begin removing or redefining these steps.

Meetings that once existed to align understanding shrink or disappear because shared context can be established asynchronously.

Approval layers change because verification can occur differently.

Roles evolve because individuals can operate across functional boundaries with augmented knowledge access.

Research increasingly reflects this shift. Studies of generative AI adoption show that organizations capturing the most value are those

---

Modern Productivity Paradox." *The American Economic Review*, vol. 80, no. 2, 1990, pp. 355–361.

4. Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." *NBER Working Paper*, no. 24001, National Bureau of Economic Research, 2017.

5. Garicano, Luis. "Hierarchies and the Organization of Knowledge in Production." *Journal of Political Economy*, vol. 108, no. 5, 2000, pp. 874–904.

redesigning processes rather than simply deploying models.<sup>6</sup> The difference between adoption and impact lies in integration into core workflows.

The firm is no longer speeding up tasks. It is altering decision pathways.

## Decision Making Becomes the Center

At earlier stages, artificial intelligence helps produce information. In the third stage, it changes how decisions are made.

Managers no longer rely solely on pre-assembled reports created over long preparation cycles.<sup>7</sup> They interact dynamically with information. Instead of waiting for an analysis, they explore possibilities iteratively. Instead of receiving a single recommendation, they test multiple alternatives quickly.

This does not eliminate human judgment. It shifts where judgment is applied.

Previously, much managerial effort went toward verifying whether information was sufficient. Now more effort goes toward determining what actions to take based on abundant information.

The manager's role evolves from reviewer to orchestrator.

*Starbucks introduced AI tools to assist store managers with staffing and inventory decisions. Early deployments improved forecasting accuracy, yet operational performance improved meaningfully only after the company aligned performance metrics and decision rights around the new system. Managers shifted from manual scheduling toward supervising system recommendations. The transformation required redefining how store decisions were made.*<sup>8</sup>

---

6. McKinsey & Company. *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey Global Institute, 2023.

7. Dell'Acqua, Fabrizio, et al. "Navigating the Jagged Technological Frontier: Field Experimental Evidence of the Effects of AI on Knowledge Worker Productivity and Quality." *Harvard Business School Working Paper*, 2023.

8. Lucas, Amelia. "Starbucks Turns to AI to Improve Store Operations." *CNBC*, 2023.

Academic work on decision support systems anticipated this shift decades earlier. Systems that increase analytical access do not replace managers but change their cognitive workload, moving them from information gathering toward problem framing and evaluation.<sup>9</sup> Generative AI extends this pattern broadly across knowledge work.

Organizations that recognize this early begin redefining performance expectations. They evaluate managers less on the ability to supervise processes and more on the ability to make timely and accountable decisions.

## The Role Structure Changes

When decision making changes, role definitions follow.

In earlier stages, firms still assume narrow specialization because expertise required deep, domain-specific knowledge accumulation. Artificial intelligence allows individuals to operate competently across adjacent areas by reducing the barrier to accessing and interpreting information.

This does not eliminate expertise. It alters its expression.

Experts spend less time producing routine analysis and more time supervising higher-level reasoning. Generalists become more capable contributors. Small teams handle responsibilities that once required large groups.<sup>10</sup>

Research from productivity studies shows that AI tools often increase performance most for less experienced workers by narrowing knowledge gaps.<sup>11</sup> As a result, the organization's capability distribution changes. Competence becomes less concentrated.

Companies in the third stage recognize this and adjust staffing models. They rely on fewer handoffs between specialized roles. Work

---

9. Simon, Herbert A. *The New Science of Management Decision*. Rev. ed., Prentice Hall, 1977.

10. Noy, Shakked, and Whitney Zhang. "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence." *Science*, vol. 381, no. 6654, 2023, pp. 187–192.

11. Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. "Artificial Intelligence and the Modern Productivity Paradox: A Clash of Expectations and Statistics." *NBER Working Paper*, no. 24001, National Bureau of Economic Research, 2017.

groups become smaller but broader in scope. The structure of collaboration shifts from sequential to integrated.

## Accountability Becomes the Hard Problem

As workflows change, a new challenge emerges.

Who is responsible when work includes machine reasoning?

In earlier stages, organizations often avoided this question because artificial intelligence operated at the margins. In the third stage, it sits inside decision processes. Leaders must define accountability clearly.

Regulators and policymakers have begun emphasizing auditability and explainability precisely for this reason. As automated reasoning informs operational decisions, organizations must demonstrate how conclusions were reached and who authorized them.<sup>12</sup>

Successful firms respond by treating artificial intelligence as part of the decision system rather than an external tool. They design oversight processes ensuring that a human owner remains responsible for outcomes, while also documenting how machine assistance contributed.

The organization becomes neither fully automated nor unchanged. It becomes hybrid.

## HOW STAGE THREE APPEARS ACROSS FIRM SIZES

### *Small Businesses*

Small businesses experience this stage as leverage. A small team operates with capabilities previously associated with large organizations. A founder performs research, analysis, and planning without external support. The firm competes in markets that once required scale because cognitive capacity is augmented. The business model changes before the industry structure does.

### *Mid-Market Firms*

Mid-market firms experience this stage as efficiency. They reorga-

---

12. Organisation for Economic Co-operation and Development (OECD). *OECD Principles on Artificial Intelligence*. OECD, 2019, updated 2023.

nize workflows to remove coordination overhead. Decision cycles shorten. Fewer layers are required to operate reliably. These companies often see the earliest measurable financial improvement because they are structured enough to capture gains yet flexible enough to redesign processes.

#### *Enterprise Organizations*

Large enterprises experience this stage as restructuring. Their greatest benefits arrive later because coordination complexity is higher, but the magnitude of change is larger. Entire operational models shift. Certain reporting layers shrink while others expand. Governance frameworks become integrated into daily operations rather than existing as separate compliance exercises.

McKinsey's enterprise AI research consistently finds that organizations redesigning workflows and operating models around AI achieve significantly greater value than those treating it as isolated automation.<sup>13,14</sup>

The scale of impact correlates with scale of coordination.

## WHY THIS STAGE FEELS DIFFERENT

Unlike earlier stages, the third stage cannot remain a pilot.

Removing steps requires trust. Redefining roles requires leadership commitment. Changing accountability requires executive clarity. The organization must choose to operate differently rather than simply experiment.

At this point artificial intelligence stops being an innovation initiative. It becomes part of how the company functions.

Executives often describe a noticeable shift. The technology stops being a topic of discussion and becomes an assumption. Teams no longer debate whether to use AI. They design processes assuming it is present.

---

13. McKinsey & Company. *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey Global Institute, 2023.

14. Maslej, Nestor, et al. *AI Index Report 2024*. Stanford Institute for Human-Centered Artificial Intelligence, 2024.

This is the moment the firm has actually adopted artificial intelligence.

Not when it purchased software.

Not when it trained employees.

When it reorganized work.

## THE MEANING OF THE THIRD STAGE

The third stage is not the end of adoption. It is the beginning of a different type of organization.

The company now operates in an environment where knowledge is abundant rather than scarce. Coordination mechanisms designed for slow information processing give way to systems designed for rapid decision making. Management focuses less on supervising activity and more on ensuring accountability.

Artificial intelligence did not replace the firm. It changed what the firm is.

Earlier stages showed that the technology was useful.

This stage shows that the organization itself must evolve.

The value of artificial intelligence does not come from having better tools than competitors.

It comes from becoming a different kind of company.

## KEY TAKEAWAYS

- **Stage Three begins with a different question.** The shift occurs when leaders stop asking where to apply AI and start asking why existing steps exist at all.
- **Enterprise value appears only after workflow redesign.** Task acceleration in earlier stages improves preparation, but measurable performance gains emerge only when coordination, approval paths, and decision structures change.
- **Decision-making, not automation, becomes central.** Managers move from reviewing information to orchestrating

action in an environment where analysis is abundant and rapidly generated.

- **Role structures evolve.** AI reduces the need for narrow specialization and sequential handoffs, enabling smaller teams with broader scope and shifting expertise toward higher-level judgment.
- **Accountability becomes the defining leadership challenge.** As machine-assisted reasoning enters core processes, firms must clarify decision ownership and oversight within a hybrid human–AI system.



# CHAPTER 12

## STAGE THREE RISKS

### THE MANAGEMENT RISKS OF THE INTELLIGENT ENTERPRISE

THE THIRD STAGE of adoption is where organizations finally receive what they have been seeking.

Workflows change.

Decision cycles shorten.

Managers stop waiting for analysis and begin interacting with it.

Employees operate across domains with surprising effectiveness.

The firm becomes noticeably more capable.

For the first time, artificial intelligence affects not how a task is performed but how the company itself operates. Leaders often experience this moment as relief. After months or years of pilots and partial progress, something real has occurred. The organization is faster. Coordination improves. Certain problems that previously required significant effort become routine.

It is precisely here that the largest risks appear.

The earlier stages of adoption are risky in familiar ways. There is data exposure. There is inconsistent use. There are governance gaps. These are operational risks. They resemble the risks associated with any new enterprise system.

Stage Three introduces something different. The organization has redesigned how it makes decisions, but management practices were

designed for a slower informational environment. The firm now acts at a speed its control structures were never built to supervise.

Artificial intelligence has not been added to the organization.

It has altered the organization's operating physics.

The company becomes capable of moving faster than it understands.

## ACCELERATION WITHOUT UNDERSTANDING

Historically, managerial hierarchy served a practical purpose. It slowed action long enough to verify knowledge. Reports moved upward because gathering information required time. Reviews existed because analysis could contain errors. Approval layers were inefficient but stabilizing. They created a delay between interpretation and commitment.

Artificial intelligence removes much of that delay.

A manager can now examine market conditions, operational data, and external research interactively. Teams arrive at meetings with prepared analysis rather than waiting for it. Recommendations appear quickly and often persuasively. The organization experiences a reduction in uncertainty about what is happening.

What does not automatically improve is certainty about what should be done.

This distinction is critical. Artificial intelligence accelerates interpretation, not judgment.<sup>1</sup>

Psychological research has long documented that humans tend to trust automated outputs when they are presented confidently, a phenomenon known as automation bias.<sup>2</sup> Generative systems intensify this tendency because they do not merely output numbers. They produce reasoning language. The presence of explanation increases perceived reliability even when reasoning quality varies.

In earlier stages, an incorrect output affects a person's task. In the

---

1. Kahneman, Daniel, Olivier Sibony, and Cass R. Sunstein. *Noise: A Flaw in Human Judgment*. Little, Brown Spark, 2021.

2. Parasuraman, Raja, and Victor Riley. "Humans and Automation: Use, Misuse, Disuse, Abuse." *Human Factors*, vol. 39, no. 2, 1997, pp. 230–253.

third stage, it can affect the company's behavior. Decisions propagate operationally. A pricing change affects customers immediately. A forecasting error influences production planning. A policy interpretation alters compliance activity.

The firm has not become reckless.

It has become fast.

*Zillow expanded its home-buying program using an automated pricing model that rapidly generated property valuations and purchase offers. The system allowed the company to scale transactions quickly across markets. However, pricing errors accumulated faster than the organization could catch them, leading to large inventory losses and the eventual shutdown of the initiative. The model did not fail immediately. The speed of decisions amplified small misjudgments.*<sup>3</sup>

The organization now operates with tightly coupled decisions. Organizational theorist Charles Perrow described tightly coupled systems as ones in which failures spread quickly because processes are interdependent.<sup>4</sup> When decisions rely on rapidly generated analysis across many functions, small errors can propagate widely before humans recognize them.

The new risk is not more mistakes. It is faster mistakes with broader consequences.

Managers therefore face a new responsibility. They must supervise the pace of commitment, not merely the accuracy of information.

## THE PROBLEM OF DECISION OWNERSHIP

Artificial intelligence also alters a quieter but equally important feature of organizations: authority.

Traditional corporate hierarchies were partly justified by information asymmetry. Senior leaders had broader visibility. Lower levels had

---

3. Chen, Brian X., and Erin Griffith. "Zillow to Stop Buying Homes After Racking Up Losses." *The New York Times*, 2 Nov. 2021.

4. Perrow, Charles. *Normal Accidents: Living with High-Risk Technologies*. Princeton University Press, 1984.

partial visibility. Decision authority therefore concentrated upward. The structure aligned knowledge and permission.

Stage Three disrupts that alignment.

Employees across levels can now produce credible analysis. A product manager can model financial implications. A financial analyst can evaluate operational alternatives. A marketing leader can interrogate customer behavior directly. Real authority begins to spread because knowledge spreads.

Yet formal authority often does not change.

Economists Aghion and Tirole distinguish between real authority and formal authority in organizations.<sup>5</sup> Real authority derives from knowledge. Formal authority derives from position. Artificial intelligence expands real authority without automatically redefining formal authority.

The result is confusion.

Employees capable of deciding hesitate because permission is unclear. Managers review decisions they did not meaningfully contribute to understanding. Executives encounter escalating volumes of well-supported recommendations. Instead of faster decisions, the organization experiences extended deliberation.

More information does not automatically produce clarity. It often produces negotiation.

Without explicit redesign, artificial intelligence can create decision diffusion. Everyone understands enough to participate, yet no one clearly owns the outcome. When outcomes are ambiguous, organizations slow down to regain confidence.

The company does not lack intelligence.

It lacks defined responsibility.

Stage Three therefore requires something more difficult than tool adoption. It requires deliberate reallocation of decision rights.

---

5. Aghion, Philippe, and Jean Tirole. "Formal and Real Authority in Organizations." *Journal of Political Economy*, vol. 105, no. 1, 1997, pp. 1-29.

## ORGANIZATIONAL DEPENDENCY

As redesign continues, a second risk emerges gradually and often invisibly.

The organization begins to depend on machine-supported reasoning.

At first the dependence appears benign. Employees consult systems frequently because they are useful. Managers expect immediate analysis because it is available. Meetings are scheduled assuming live information access.

Eventually the firm stops designing workflows that function without these systems.

This is not technological dependence in the traditional sense. It is cognitive dependence. The company has incorporated artificial intelligence into how it thinks.

Research on enterprise adoption emphasizes that value emerges when AI becomes embedded in core workflows.<sup>6</sup> Embedding creates benefit. It also creates exposure. When reasoning infrastructure is unavailable, the organization struggles to operate at expected speed. When system behavior changes, teams must reinterpret familiar processes.

The company has effectively adopted a new utility: reasoning capability.

Yet many firms treat this capability differently from other utilities. They plan for power outages and network failures. They rarely plan for analytic interruption. Few organizations define fallback decision procedures. Few determine which decisions must pause if systems cannot assist.

The risk is not that artificial intelligence fails. The risk is that the organization no longer knows how to function at its prior cognitive speed.

The firm has gained capability but lost independence.

---

6. McKinsey & Company. *Scaling AI Across the Enterprise: A Guide to Capturing Value*. McKinsey Global Institute, 2024.

## ACCOUNTABILITY IN A HYBRID ORGANIZATION

A further challenge follows naturally. When machine reasoning participates in decision processes, responsibility becomes ambiguous unless explicitly defined.

In earlier stages, leaders could treat AI as advisory. In Stage Three, it becomes integrated into operational judgment. Organizations increasingly face questions about who authorized actions informed by automated analysis.

Regulatory frameworks emphasize human accountability for precisely this reason. The OECD AI Principles and the NIST AI Risk Management Framework both stress that automated systems cannot hold responsibility. Organizations and identifiable decision makers must.<sup>7</sup>

Without clarity, behavioral patterns emerge. Individuals begin attributing decisions to systems. Recommendations become explanations. Accountability weakens because no single person claims ownership.

This creates operational fragility. When failures occur, learning requires knowing who decided and why. Hybrid decision systems demand stronger, not weaker, managerial ownership.

Artificial intelligence increases the importance of management judgment. It does not dilute it.

## CULTURAL OVERCORRECTION

Many organizations respond to these risks predictably. They attempt to restore safety by restoring friction.

Approval layers reappear. Committees multiply. Review processes expand. Leaders attempt to regain confidence by slowing action. The company recreates the structures artificial intelligence made unnecessary.

This rarely works.

---

7. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.

Employees continue operating at higher cognitive speed while formal processes operate at lower administrative speed. The organization becomes internally inconsistent. Workarounds proliferate. Informal coordination replaces formal coordination.

High reliability organizations such as aviation and healthcare maintain safety not by slowing operations but by clarifying escalation, monitoring, and authority. The lesson applies here. Faster organizations require clearer governance, not heavier bureaucracy.<sup>8</sup>

The challenge is not to reintroduce delay.

The challenge is to introduce clarity.

## THE REAL STAGE THREE RISK

Earlier stages of adoption test technology. The third stage tests leadership.

The organization has altered its decision environment. Information is abundant. Analysis is immediate. Coordination accelerates. Yet management habits formed in a slower era remain.

The resulting mismatch produces the true risk of Stage Three.

The company is no longer limited by what it knows.

It is limited by how well it decides.

Artificial intelligence increases the organization's ability to act. If responsibility, authority, and oversight are not redesigned accordingly, capability outruns control. Firms do not fail because the models are inadequate. They fail because managerial systems lag the organization's new speed.

The firms that navigate this stage successfully do not simply adopt artificial intelligence. They redesign management around continuous decision making. They clarify ownership. They supervise outcomes rather than activity. They manage pace as carefully as they manage accuracy.

Artificial intelligence does not make companies fragile.

It reveals whether they were designed to decide.

---

8. Weick, Karl E., and Kathleen M. Sutcliffe. *Managing the Unexpected: Sustained Performance in a Complex World*. 3rd ed., Wiley, 2015.

## KEY TAKEAWAYS

- **Stage Three risk is systemic, not technical.** Once AI is embedded in workflows, errors affect coordinated decisions rather than isolated tasks.
- **Speed becomes a management problem.** Faster analysis increases the risk of rapid, organization-wide mistakes unless commitment processes are redesigned.
- **Authority must be clarified.** AI expands who has knowledge but not automatically who has permission, creating decision diffusion without explicit ownership.
- **Organizations develop hidden dependency.** Firms begin relying on machine-supported reasoning as operational infrastructure without planning fallback decision processes.
- **Accountability strengthens rather than weakens.** Hybrid human–AI decisions require identifiable human owners for operational learning and trust.
- **The central risk of success is managerial lag.** Companies fail not because the technology is immature but because management systems were built for a slower informational environment.

# PART THREE

## THE MANAGEMENT REWRITE

What is happening inside organizations is no longer speculative.

Individuals are changing how they work.

Organizations are experimenting.

Some firms are already redesigning how decisions move.

Observation alone is no longer sufficient. Leaders do not have the luxury of distance. They must choose how their organizations will operate.

The question is not whether artificial intelligence will alter the firm. It already has. The real question is whether leaders will shape that alteration deliberately or allow it to unfold unintentionally.

Artificial intelligence is not primarily a technical rollout. It is a governance challenge, a labor shift, a strategic inflection point, and ultimately an organizational design problem. It changes how work is performed, how authority is assigned, how performance is measured, and how competitive advantage is created.

The organizations that succeed will not simply deploy better tools. They will redefine roles, redesign workflows, and align management practices to a world where knowledge is abundant and reasoning can be partially automated.

The task is not adoption.  
The task is leadership.

# CHAPTER 13

# LEADERSHIP

## MANAGEMENT IN THE HUMAN AND MACHINE COMPANY

ARTIFICIAL INTELLIGENCE DOES NOT REMOVE management.

It removes the reason management was originally structured the way it is.

Modern organizations were designed for a world in which reliable information was scarce and slow. Managers coordinated people because knowledge had to be assembled piece by piece.<sup>1</sup> Reports moved upward so decisions could move downward. Meetings aligned understanding across functions. Review processes existed because verifying accuracy required time and expertise. Supervision was not merely oversight of employees; it was a mechanism for processing information that could not be gathered or interpreted quickly.

Artificial intelligence alters that environment.

When knowledge becomes abundant and immediate, coordination changes. The manager is no longer needed primarily to supervise activity. The manager is needed to direct judgment.<sup>2</sup>

For more than a century, the corporation has relied on a specific

---

1. Coase, Ronald H. "The Nature of the Firm." *Economica*, vol. 4, no. 16, 1937, pp. 386–405.

2. Davenport, Thomas H., and Nitin Mittal. *All-in on AI: How Smart Companies Win Big with Artificial Intelligence*. Harvard Business Review Press, 2023.

conception of managerial work: ensuring that work was completed correctly before action was taken. But when analysis can be produced rapidly and interpreted interactively, the problem management solves shifts. The manager is no longer the person who ensures work is completed. The manager becomes the person who ensures the organization decides.

*Netflix restructured much of its content decision process around data science and machine learning models that forecast audience engagement. Rather than layering additional approval committees, leadership clarified that executives were accountable for final greenlight decisions while data systems provided rapid scenario modeling. Managers were expected to interpret model output, challenge assumptions, and commit to direction rather than supervise analysis preparation. The structure shifted from processing information upward to orchestrating judgment around shared data.<sup>3</sup>*

This shift is subtle but profound. Outward activities remain familiar. Teams still meet. Reports still exist. Plans are still drafted. Yet the underlying purpose of these activities changes as the cost of producing and verifying knowledge declines. The human and machine company does not eliminate management. It relocates its center of gravity.

## **MANAGERS AS ORCHESTRATORS RATHER THAN SUPERVISORS**

Historically, managerial effort was devoted to supervision. Work moved through layers because each layer added interpretation. An analyst gathered information, a manager checked reasoning, a director reconciled competing inputs, and an executive authorized action. The sequence existed because reliable knowledge required coordination among specialists.

Artificial intelligence changes the economics of interpretation. Employees now arrive at discussions already equipped with synthe-

---

3. Gomez-Uribe, Carlos A., and Neil Hunt. "The Netflix Recommender System." *ACM Transactions on Management Information Systems*, vol. 6, no. 4, 2015.

sized analysis. A manager no longer waits for a briefing to understand a situation; the manager can examine the same materials directly and interactively.

As a result, the manager's primary responsibility shifts from supervising production to orchestrating decisions.

Orchestration means coordinating human judgment and machine reasoning into accountable action. The manager evaluates whether sufficient information exists, whether alternative interpretations have been considered, and whether the organization should act now rather than later. Timing becomes a leadership responsibility.

Herbert Simon anticipated this change decades ago.<sup>4</sup> As computational assistance improves, managerial work moves away from information processing toward problem framing and evaluation. The central difficulty is no longer obtaining analysis but determining which questions matter and when to commit to a course of action. Artificial intelligence extends this shift across the organization. It does not remove leaders. It increases the importance of their judgment.

Managers therefore spend less time monitoring activity and more time guiding attention. They ensure teams examine the right risks, consider the right alternatives, and recognize when analysis is converging toward a decision rather than expanding endlessly. The organization moves from coordinating effort to coordinating understanding.

## NEW PERFORMANCE MEASUREMENT

As managerial work changes, performance measurement must change as well.

Industrial organizations measured productivity through output: units produced, calls handled, transactions completed. Knowledge organizations preserved similar habits: reports written, projects completed, hours billed. These measures made sense when effort correlated with output.

Artificial intelligence weakens that relationship. A single employee

---

4. Simon, Herbert A. *The New Science of Management Decision*. Prentice Hall, 1960.

can now generate multiple high-quality analyses in the time previously required for one. Counting outputs becomes less meaningful because the cost of producing them has collapsed.<sup>5</sup>

Organizations therefore begin measuring different things.

*When Klarna deployed their AI assistants across customer support operations, the system began handling a large share of customer inquiries and drafting responses for agents to review. As preparation time collapsed, managers found that traditional performance measures (such as messages written or time spent per ticket) no longer reflected capability.*

*Instead, supervisors evaluated how employees interpreted system suggestions, corrected errors, and handled unusual cases the model could not resolve. Productivity no longer meant producing responses. It meant exercising judgment and knowing when intervention was necessary.<sup>6</sup>*

They observe decision cycle time rather than document production.

They track how quickly problems move from identification to resolution.

They examine whether teams revisit issues because context was insufficiently shared.

Productivity increasingly means clarity of action.<sup>7</sup> An effective organization understands what to do and acts responsibly without unnecessary delay.

Managers consequently evaluate reasoning rather than activity. They assess whether assumptions were examined, whether alternatives were considered, and whether decisions were taken at the right moment. Activity is visible. Judgment requires trust. Yet as preparation accelerates, activity ceases to reflect capability.

---

5. Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age*. W. W. Norton, 2014.

6. Wiggers, Kyle. "Klarna Says Its AI Assistant Does the Work of 700 Customer Service Agents." *TechCrunch*, 2024.

7. Weick, Karl E., and Kathleen M. Sutcliffe. *Managing the Unexpected: Sustained Performance in a Complex World*. 3rd ed., Wiley, 2015.

## NEW LEADERSHIP COMPETENCIES

These changes require new managerial skills.

First, leaders must understand machine reasoning. They do not need to be engineers, but they must recognize patterns. Artificial intelligence can synthesize information rapidly while presenting unwarranted confidence.<sup>8</sup> Managers must know when analysis requires verification and when it supports action.

Second, leaders must master problem framing. When information becomes abundant, the challenge shifts from finding answers to asking the right questions. Artificial intelligence can generate many plausible responses. Without clear framing, organizations risk analyzing indefinitely.

Third, leaders must integrate across boundaries. As knowledge barriers fall, work no longer remains strictly within functional silos. Managers coordinate collaboration between specialists and generalists and ensure accountability remains clear even as responsibilities expand.

Leadership effectiveness depends less on possessing information personally and more on guiding how information is used collectively.

## DECISION ACCOUNTABILITY

The most significant change concerns responsibility.

Artificial intelligence produces options. It does not assume consequences.

When machine reasoning participates in analysis, accountability must remain human.<sup>9</sup> A person must authorize action. The manager becomes the point at which possibility becomes commitment.

Leaders cannot defer judgment to systems. They must explain why a decision was taken, what evidence supported it, and what risks were accepted. Artificial intelligence increases information, but it does not

---

8. OpenAI. *GPT-4 Technical Report*. OpenAI, 2023.

9. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.

determine acceptable risk. That remains a human function grounded in context and responsibility.

In earlier organizations, slow analysis sometimes concealed indecision. Leaders could wait for more information. In the human and machine company, information appears quickly. Delay becomes a decision itself. The expectation shifts from careful preparation to timely judgment.

## WHAT MANAGEMENT BECOMES

Management no longer centers on supervising activity.

It centers on guiding judgment.

The firm's performance depends less on how much work employees perform and more on how well the organization interprets knowledge and acts upon it. Artificial intelligence does not diminish management. It concentrates it.

Preparation accelerates. Judgment becomes visible. Leadership shifts from overseeing effort to directing attention and responsibility.

The defining question for the modern firm is no longer whether it uses artificial intelligence.

It is whether its leaders can decide well in a world where understanding arrives faster than ever before.<sup>10</sup>

## KEY TAKEAWAYS

- **Management does not disappear.** Artificial intelligence removes the informational constraints that originally required layered supervision, shifting management's purpose from coordinating activity to guiding judgment.
- **The manager becomes an orchestrator.** Instead of verifying work and processing reports, leaders integrate human judgment and machine reasoning into accountable organizational decisions.

---

10. McKinsey & Company. *The State of AI in 2024: Adoption Expands, but Value Remains Uneven*. McKinsey Global Survey on AI, 2024.

- **Productivity measurement changes.** As cognitive effort falls, output counts matter less; organizations increasingly measure performance through decision cycle time, clarity, and the ability to act responsibly without delay.
- **Leadership accountability increases.** Machine assistance expands analysis but does not assume responsibility, making managers more directly accountable for interpreting evidence and committing the organization to action.
- **Effective leadership becomes a cognitive discipline.** The central managerial skills are framing the right questions, directing organizational attention, and deciding under conditions of abundant information rather than scarce information.



# CHAPTER 14

# GOVERNANCE

## AI ADOPTION IS A PERMISSIONS PROBLEM

WHEN ORGANIZATIONS BEGIN to confront artificial intelligence seriously, governance is usually the first executive instinct. Leaders worry about data exposure, model bias, intellectual property leakage, hallucinations, regulatory risk, and reputational damage.<sup>1</sup> Legal departments draft policy statements. IT creates approved tool lists. Security teams define access controls.<sup>2</sup> Risk committees request reporting.

All of this activity is rational. It is also incomplete.

The deeper governance challenge of artificial intelligence is not model validation, nor cybersecurity, nor regulatory interpretation. Those are components. The central governance issue is structural. Artificial intelligence changes who can know what, how quickly they can know it, and therefore who can reasonably decide.

AI governance, at its core, is a permissions redesign.

The industrial corporation was built on scarcity of information. Decisions traveled upward because information aggregated slowly. Reports were prepared, reviewed, escalated, and approved because

---

1. Gartner. *Top Strategic Technology Trends 2024: Trust, Risk, and Security Management (TRiSM)*. Gartner Research, 2024.

2. Deloitte. *State of Generative AI in the Enterprise: Now Decides Next*. Deloitte Insights, 2024.

few individuals had access to the full picture. Supervisory layers were not merely political constructs; they were mechanisms for synthesizing distributed knowledge.<sup>3</sup>

Artificial intelligence alters the cost of synthesis.

An employee who previously required days to assemble research can now generate structured analysis in minutes. A product manager can interrogate market data interactively. A finance analyst can simulate multiple scenarios quickly. An engineer can test architectural alternatives rapidly. Information that once justified managerial intermediation becomes directly accessible.

The firm must therefore answer a fundamental question: when information is abundant, who is authorized to act?

This is not a technical configuration decision. It is a governance choice about authority, responsibility, and risk tolerance.

## DECISION AUTHORITY IN AN ABUNDANT INFORMATION ENVIRONMENT

In traditional corporate design, decision rights are distributed according to a logic described extensively in organizational economics. Jensen and Meckling's theory of the firm emphasized that the allocation of decision rights and monitoring mechanisms defines how effectively an organization functions under conditions of incomplete information.<sup>4</sup> When information is costly and unevenly distributed, decision authority tends to concentrate at higher levels where broader visibility exists.

Artificial intelligence destabilizes that equilibrium.

When analysis can be produced and interrogated widely, lower-level managers often possess sufficient information to make sound decisions. Yet formal authority structures may not have changed. The result is a tension between capability and permission.

---

3. Coase, Ronald H. "The Nature of the Firm." *Economica*, vol. 4, no. 16, 1937, pp. 386–405.

4. Jensen, Michael C., and William H. Meckling. "Theory of the Firm: Managerial Behavior, Agency Costs and Ownership Structure." *Journal of Financial Economics*, vol. 3, no. 4, 1976, pp. 305–360.

*When Coca-Cola deployed generative AI for marketing, local teams suddenly gained the ability to produce campaign concepts and visuals without central creative groups. The issue was not model accuracy. It was brand authority. The company had to decide which marketing decisions still required corporate approval once many employees could independently generate professional-quality materials.<sup>5</sup>*

In some firms, this manifests as hesitation. Employees prepare sophisticated analyses using AI tools but defer decisions because formal approval pathways remain unchanged. In others, the opposite occurs. Employees act based on AI-generated insight without clear understanding of escalation thresholds, creating inconsistency and risk.

Neither outcome reflects technological failure. Both reflect unresolved governance.

Organizational research has long shown that decision rights must align with information distribution for performance to improve.<sup>6</sup> Artificial intelligence shifts information distribution dramatically. Unless authority is recalibrated accordingly, the firm experiences friction rather than acceleration.

This is why AI governance cannot be confined to tool selection. It requires explicit articulation of decision authority in the new information environment.

Who can commit capital based on AI-supported forecasts?

Who can alter pricing using algorithmic recommendations?

Who can approve customer actions informed by automated analysis?

What thresholds trigger human review?

Without clarity, artificial intelligence amplifies confusion.

---

5. Slefo, George P. "Coca-Cola Is Using Generative AI to Reinvent Marketing." *Ad Age*, 2023.

6. Aghion, Philippe, and Jean Tirole. "Formal and Real Authority in Organizations." *Journal of Political Economy*, vol. 105, no. 1, 1997, pp. 1–29.

## AUDITABILITY AND THE TRACEABILITY IMPERATIVE

As decision authority shifts, the second governance challenge emerges: traceability.

When artificial intelligence systems contribute to reasoning, organizations must maintain the ability to reconstruct how conclusions were reached. This requirement is not merely academic. Regulatory bodies globally are increasingly emphasizing transparency, explainability, and auditability in automated decision systems. The OECD's AI Principles highlight transparency and accountability as foundational governance requirements.<sup>7</sup> The European Union's AI Act codifies obligations for high-risk systems to document reasoning pathways and maintain human oversight.<sup>8</sup> In the United States, agencies including the Federal Reserve and the Securities and Exchange Commission have signaled increasing scrutiny of automated decision systems in financial contexts.

The concern is not that artificial intelligence is inherently opaque. It is that informal adoption can produce invisible reasoning chains.

In early-stage adoption, employees often consult AI systems conversationally. They may iterate through prompts, refine assumptions, and arrive at conclusions without preserving the intermediate reasoning steps. If a decision later proves problematic, the organization may struggle to explain how it was reached.

This creates governance exposure.

*As Walmart expanded AI-generated product descriptions and operational recommendations, managers discovered a new requirement: explanation. The systems worked, but leaders needed to know how a recommendation had been produced when questions arose. The company invested in logging and monitoring so decisions could be reconstructed afterward. Governance required traceability, not just accuracy.<sup>9</sup>*

---

7. Organisation for Economic Co-operation and Development. *OECD AI Principles*. OECD, 2019, updated 2023.

8. European Union. *Artificial Intelligence Act (Regulation (EU) 2024/1689)*. Official Journal of the European Union, 2024.

9. Walmart Global Tech. "Responsible AI at Walmart." Walmart Technology Blog, 2023.

Auditability must be designed deliberately. Systems must log inputs, assumptions, and outputs in ways that allow reconstruction. Decision documentation processes must incorporate machine assistance transparently rather than implicitly.

Importantly, auditability does not require slowing adoption. It requires embedding traceability into workflows from the outset. Firms that treat AI as an informal productivity tool often postpone this discipline, creating future risk.

The companies that mature into Stage Three are those that recognize early that artificial intelligence participation in decision processes must be visible, documentable, and attributable.

## ACCOUNTABILITY REMAINS HUMAN

One of the most persistent misconceptions in organizational conversations about AI is that machine participation diffuses responsibility. Leaders sometimes speak as if algorithmic recommendation creates a buffer. In reality, it does the opposite.

As machine-generated analysis becomes ubiquitous, the locus of accountability becomes clearer, not weaker.

The OECD AI Principles, along with the NIST AI Risk Management Framework, consistently emphasize human oversight and responsibility as non-negotiable elements of trustworthy AI deployment.<sup>10</sup> The rationale is straightforward. Algorithms do not bear legal liability. Organizations do.

In the human and machine company, artificial intelligence becomes a contributor to reasoning, not an autonomous decision-maker in the legal or ethical sense. A model may recommend credit approval, but a human authorizes it. A system may propose a hiring shortlist, but a manager validates it. A forecasting engine may project revenue, but an executive communicates guidance.

Governance must codify this explicitly.

---

10. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.

*UPS's AI systems began recommending operational changes such as routing and staffing adjustments. The challenge was not understanding the recommendations but deciding when a manager could implement them directly. The company defined escalation thresholds specifying which decisions required higher approval. Governance meant defining when a recommendation became a decision.<sup>11</sup>*

The danger lies not in overreliance on AI alone but in ambiguity about who ultimately decides. When responsibilities blur, organizations struggle to respond to failure. Clear accountability frameworks preserve trust both internally and externally.

Paradoxically, artificial intelligence increases the importance of managerial judgment. When analysis can be generated instantly, leaders cannot justify delay on grounds of insufficient information. They must decide. Accountability accelerates.

## **WHY AI GOVERNANCE CANNOT SIT SOLELY IN IT**

Given the technological nature of artificial intelligence, many organizations instinctively assign governance responsibility to IT or the Chief Information Officer. This approach captures infrastructure and security risks effectively. It fails, however, to address the broader organizational transformation.

Artificial intelligence influences marketing strategy, supply chain decisions, financial forecasting, legal review, product design, human resources processes, and customer engagement. It shapes the core of how value is created.

Research from MIT Sloan Management Review and Boston Consulting Group has repeatedly shown that digital transformations led exclusively by IT departments underperform those sponsored cross-functionally by executive leadership because they fail to alter organizational routines beyond technical deployment.<sup>12</sup>

---

11. O'Donnell, Jayne. "UPS Uses AI to Optimize Delivery and Operations." *USA Today*, 2023.

12. Ransbotham, Sam, et al. "Artificial Intelligence in Business Gets Real." *MIT Sloan Management Review* and Boston Consulting Group, 2018.

AI governance therefore must extend beyond platform control. It must address:

- Cross-functional decision rights
- Risk tolerance thresholds
- Escalation protocols
- Capital allocation criteria
- Documentation standards

These are executive responsibilities.

IT manages infrastructure.

Legal interprets regulatory requirements.

Risk oversees compliance exposure.

Finance monitors economic impact.

Executive leadership defines authority.

Artificial intelligence governance is not a compliance layer added to technology. It is a redesign of how decisions move through the firm.

## GOVERNANCE AS ENABLER, NOT CONSTRAINT

When governance conversations are framed solely around restriction, organizations respond defensively. Employees perceive artificial intelligence as risky, controlled, and constrained. Innovation slows. Informal shadow usage increases.

Effective governance performs the opposite function. It creates clarity that enables confident action.

When employees understand precisely which tools are approved, what data can be used, which decisions can be made autonomously, and which require escalation, they operate decisively within defined boundaries.

Organizational theory consistently demonstrates that clear boundaries improve performance more than vague autonomy.<sup>13</sup> Artificial

---

13. Simons, Robert. *Levers of Control: How Managers Use Innovative Control Systems to Drive Strategic Renewal*. Harvard Business School Press, 1995.

intelligence governance functions as a boundary system. It defines acceptable experimentation while preserving oversight.

Companies that articulate permissions clearly accelerate safely. Companies that avoid defining them oscillate between uncontrolled experimentation and heavy-handed restriction.

## THE STRATEGIC CONSEQUENCE

The firms that master the permissions problem first will enjoy disproportionate advantage.

They will move faster because employees understand when they can act. They will incur fewer regulatory surprises because traceability is embedded. They will avoid internal paralysis because accountability is explicit. They will scale adoption coherently rather than through disconnected pilots.

Artificial intelligence governance is therefore not a defensive posture. It is a structural advantage.

When information becomes abundant, coordination becomes the scarce resource. Governance defines coordination.

In earlier eras, scale created advantage because information traveled slowly and capital was concentrated. In the emerging era, coordination speed becomes decisive. Governance clarity directly influences that speed.

Artificial intelligence adoption is not ultimately about tools. It is about authority.

The companies that recognize this will not simply deploy technology more safely. They will redesign their decision architecture deliberately.

And in doing so, they will move beyond experimentation into transformation.

## KEY TAKEAWAYS

- **AI governance is not primarily a technical control system.**

The central governance problem is determining who is

authorized to make decisions when reliable analysis becomes widely accessible.

- **Artificial intelligence is a permissions redesign.** As information becomes cheaper to produce and interpret, traditional approval layers no longer align with where knowledge actually resides, creating tension between capability and authority.
- **Decision rights must follow information.** Organizations that fail to realign authority with distributed access to analysis experience hesitation, inconsistent action, and organizational friction rather than acceleration.
- **Traceability becomes essential to trust.** As machine-assisted reasoning enters operational decisions, firms must be able to reconstruct how conclusions were reached and who authorized action.
- **Accountability remains human.** AI systems contribute analysis but cannot assume responsibility; governance must explicitly define ownership of decisions rather than allowing responsibility to blur.
- **Governance is an enabler of speed, not a constraint.** Clear permissions, escalation thresholds, and documentation standards allow employees to act confidently while preserving oversight.
- **Competitive advantage shifts to coordination.** Firms that define authority clearly will move faster, scale adoption coherently, and realize enterprise value earlier than those treating governance as a compliance exercise.



## CHAPTER 15

# SKILLS AND TALENT

### JOBS DO NOT DISAPPEAR, WORK MOVES UP THE CHAIN

EVERY MAJOR TECHNOLOGICAL shift produces the same first reaction.

People do not ask what will become possible.

They ask what will disappear.

Artificial intelligence has followed this pattern almost perfectly. Public conversations, media coverage, and even executive dialogue frequently begin with a single question: how many jobs will AI eliminate?

The question is understandable. It is also historically incomplete.

Technologies rarely eliminate work in the way societies initially expect. They eliminate tasks. The distinction matters more than it appears. A job is a bundle of responsibilities, decisions, and accountabilities. A task is a component activity within that bundle. Automation tends to remove specific activities long before it removes entire roles.

Artificial intelligence is best understood as a task automation technology embedded within a cognitive workflow, not as a wholesale replacement for occupational roles.

*Inside Microsoft and several partner firms, engineers began using the coding assistant GitHub Copilot to help write and review software. Researchers found*

*developers completed substantially more code contributions and routine programming tasks faster, even though the job itself did not disappear. Instead of spending hours writing boilerplate functions, engineers spent more time debugging, architecture decisions, and reviewing alternatives. The role shifted from writing code line-by-line to deciding which code should exist.<sup>1</sup>*

Research from the Organisation for Economic Co-operation and Development consistently shows that most occupations consist of diverse task combinations, many of which remain resistant to full automation even when some activities within the role can be automated.<sup>2</sup> Similarly, work by economists David Autor and colleagues demonstrates that technological change historically restructures occupations by reallocating tasks rather than eliminating employment categories outright.<sup>3</sup>

Artificial intelligence therefore does not simply remove workers from organizations. It removes portions of their work.

## TASK REMOVAL VERSUS ROLE ELIMINATION

To understand the practical implications, consider how work is actually performed in most professional environments. An accountant does not only calculate financial results. The accountant gathers information, reconciles discrepancies, interprets regulations, communicates implications to leadership, and participates in planning decisions. A marketing professional does not only write copy. The role involves interpreting customer behavior, positioning products, aligning campaigns with strategy, and coordinating across functions.

Artificial intelligence disproportionately affects the preparatory components of these roles.

---

1. "Copilot Developer Productivity Findings." *InfoQ*, Sept. 2024, [www.infoq.com/news/2024/09/copilot-developer-productivity/](http://www.infoq.com/news/2024/09/copilot-developer-productivity/).

2. Arntz, Melanie, Terry Gregory, and Ulrich Zierahn. "The Risk of Automation for Jobs in OECD Countries: A Comparative Analysis." *OECD Social, Employment and Migration Working Papers*, no. 189, OECD Publishing, 2016.

3. Autor, David H. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *Journal of Economic Perspectives*, vol. 29, no. 3, 2015, pp. 3–30.

- It summarizes documents.
- It drafts communications.
- It organizes data.
- It generates first-pass analysis.

*Animation studios have begun using machine-learning systems to assist with rendering, scene setup, and early visual drafts. Artists now arrive at review sessions with multiple variations generated through AI-assisted workflows rather than a single prepared concept. Directors and production leads report that meetings no longer focus on whether work was completed. They focus on selecting among alternatives and determining creative direction. The managerial task has shifted from verifying progress to choosing among possibilities.<sup>4</sup>*

These activities are important but they are not the defining value of the role. They are the steps that enable the role to operate.

Economists refer to this pattern as task substitution rather than labor substitution. Autor's research on computerization demonstrated that technology often replaces routine cognitive tasks while increasing demand for non-routine analytical and interpersonal activities.<sup>5</sup> Artificial intelligence extends this dynamic to more sophisticated cognitive work, but the mechanism remains similar.

The immediate effect is not the disappearance of the professional. It is the disappearance of preparatory labor within the profession.

When that occurs, organizations face a decision. They may reduce headcount, or they may increase the scope of what each professional can accomplish. Historically, firms tend toward the latter, particularly in knowledge industries, because reduced preparation time enables broader participation in higher-value decisions.

This pattern occurred during the adoption of spreadsheet software. Early fears predicted dramatic reductions in accounting employment. Instead, the number of accounting and financial analysis roles

---

4. Failes, Ian. "How Machine Learning Is Entering Film Production Pipelines." *before & afters*, 2023.

5. Autor, David H., Frank Levy, and Richard J. Murnane. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Quarterly Journal of Economics*, vol. 118, no. 4, 2003, pp. 1279–1333.

expanded as organizations conducted more financial planning and analysis than previously feasible. The U.S. Bureau of Labor Statistics documented growth in financial analyst and accounting roles following widespread adoption of computational tools because analysis became cheaper and therefore more widely used.<sup>6</sup>

Artificial intelligence produces a similar expansion of analytical capacity.

Organizations do not stop making decisions when analysis becomes easier. They make more of them.

## THE UPWARD MOVEMENT OF WORK

When routine cognitive tasks shrink, work does not vanish. It moves upward in abstraction.<sup>7</sup>

The employee who previously spent hours assembling information now spends time interpreting it. The analyst who once prepared data now participates in decision discussions. The manager who previously requested reports now evaluates options.

This movement has a structural consequence. Organizations demand more judgment from a broader set of employees.

Research on technological change repeatedly shows that automation increases demand for workers who can integrate information and exercise discretion. Autor's labor market studies describe this as "skill-biased technological change," in which technologies complement workers capable of non-routine reasoning while reducing demand for routine activity.<sup>8</sup>

Artificial intelligence accelerates this effect because it handles not only numerical calculation but language-based preparation. Activities once reserved for specialists become accessible across the organization.

---

6. U.S. Bureau of Labor Statistics, *Occupational Outlook Handbook: Financial Analysts*; U.S. Bureau of Labor Statistics, *Occupational Outlook Handbook: Accountants and Auditors*.

7. Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. "GPTs Are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." OpenAI, 2023.

8. Autor, David H. "Polanyi's Paradox and the Shape of Employment Growth." *Federal Reserve Bank of Kansas City Economic Symposium Proceedings*, 2015.

More employees participate in analysis. Consequently, more employees must interpret implications.

This explains a pattern that initially confuses executives. Workers report both relief and anxiety simultaneously. Relief occurs because mechanical preparation declines. Anxiety arises because responsibility increases.

Artificial intelligence removes the excuse of insufficient information. Employees must now decide what the information means.

## NEW VALUABLE SKILLS

As tasks evolve, the skills that organizations value evolve as well.

The premium shifts away from information production toward information interpretation. Employees who can ask better questions become more valuable than those who merely generate answers.

Several capabilities become particularly important.

First is problem framing. Artificial intelligence can generate solutions quickly, but it does not reliably identify which problem matters most. Research in decision science shows that mis-framing problems leads to systematic decision errors.<sup>9</sup> Leaders and employees who define clear objectives therefore provide disproportionate value.

Second is contextual reasoning. Machine-generated outputs require evaluation within business context. A recommendation may be analytically sound yet strategically inappropriate. Understanding organizational priorities, regulatory environments, and customer relationships becomes more important than executing calculations.

Third is communication. As more employees gain analytical capability, coordination depends on shared understanding. Individuals must explain reasoning clearly and translate analysis into action. Studies in organizational behavior consistently show that communication quality correlates strongly with decision effectiveness in complex environments.<sup>10</sup>

---

9. Kahneman, Daniel. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011.

10. Edmondson, Amy C. *Teaming: How Organizations Learn, Innovate, and Compete in the Knowledge Economy*. Jossey-Bass, 2012.

Fourth is ethical judgment. Artificial intelligence introduces choices about acceptable risk, fairness, and transparency. Determining how to act responsibly cannot be delegated to algorithms. Human judgment remains central.

These capabilities are not technical skills. They are cognitive and managerial skills distributed more widely across the workforce.

## HUMAN JUDGMENT AS A PREMIUM CAPABILITY

As automation expands, the scarce resource becomes judgment.

Historically, organizations paid for effort and expertise because both were difficult to obtain. Artificial intelligence lowers the cost of expertise by making information broadly available. Effort becomes less differentiating. Judgment becomes the differentiator.

This aligns with earlier observations in management thought. Peter Drucker described knowledge workers as individuals whose primary capital is their ability to apply knowledge to decisions.<sup>11</sup> Artificial intelligence does not invalidate this insight. It intensifies it. When knowledge becomes instantly accessible, the ability to apply it appropriately defines value.

*Morgan Stanley deployed an AI to help its 16,000 financial advisors search internal research and answer client questions. Executives emphasized that the goal was not to replace advisors but to improve their decision conversations with clients.*

*The value of the role shifted away from locating answers and toward interpreting what the answers meant for a specific investor. Artificial intelligence reduced the labor of analysis but increased the importance of human responsibility for the recommendation.<sup>12</sup>*

Human judgment encompasses several elements machines do not possess in a business context: understanding of tradeoffs, appreciation

---

11. Drucker, Peter F. *The Effective Executive*. Harper & Row, 1967.

12. Morgan Stanley. "Morgan Stanley Wealth Management Announces Strategic Partnership with OpenAI." Morgan Stanley Press Release, 2023.

of stakeholder consequences, and responsibility for outcomes. A model can optimize for a defined objective. It cannot determine whether the objective itself is appropriate.

Consequently, the labor market evolves not toward fewer jobs but toward different expectations. Employees who remain focused on task execution alone become vulnerable. Employees who can interpret, synthesize, and decide become more valuable.

The premium shifts from execution to discernment.

## THE ORGANIZATIONAL CONSEQUENCE

Organizations that recognize this shift adjust how they hire, train, and evaluate employees.

Training moves from teaching procedural execution to developing reasoning capability. Career progression depends less on years of experience performing tasks and more on demonstrated ability to make sound decisions. Junior employees participate earlier in meaningful discussions because preparation work no longer consumes all of their time.

In this environment, artificial intelligence does not hollow out organizations. It flattens informational barriers. Employees engage in higher-level thinking sooner.

The economic effect is therefore nuanced. Some roles contract where work consisted primarily of routine preparation. At the same time, new responsibilities emerge across functions. Labor demand changes composition rather than disappearing.

The history of technological change supports this outcome. The World Economic Forum's employment studies indicate that automation both displaces and creates jobs, with the net effect depending on how organizations adapt skills and roles.<sup>13</sup> Artificial intelligence accelerates this dynamic but does not fundamentally reverse it.

---

13. World Economic Forum. *The Future of Jobs Report 2023*. World Economic Forum, 2023.

## WHAT THIS MEANS FOR LEADERS

For leaders, the implication is clear.

The workforce challenge is not primarily redundancy. It is readiness.

Employees must be prepared to operate at a higher level of abstraction. Organizations must teach interpretation, decision-making, and communication deliberately. Performance systems must reward judgment rather than mere activity.

Artificial intelligence does not remove the need for people. It increases the importance of capable people.

Work does not disappear.

It moves upward.

The firms that understand this earliest will not focus exclusively on replacing labor. They will focus on elevating it.

## KEY TAKEAWAYS

- **AI automates tasks, not occupations.** Most roles are bundles of activities; artificial intelligence removes preparatory work within jobs rather than eliminating the jobs themselves.
- **Work moves upward in abstraction.** As drafting, summarization, and data preparation shrink, employees spend more time interpreting information, participating in decisions, and coordinating action.
- **Judgment becomes the scarce skill.** When information and analysis are widely accessible, organizational value shifts from producing answers to framing problems, evaluating tradeoffs, and choosing actions.
- **Skill demand changes, not labor demand alone.** Capabilities such as problem framing, contextual reasoning, communication, and ethical decision-making become more important than procedural execution.
- **Organizations must redesign talent development.** Training, evaluation, and career progression should emphasize reasoning and decision-making rather than task completion.

- **The leadership challenge is readiness, not redundancy.** Firms that focus only on workforce reduction will miss value; firms that elevate employee responsibility will capture it.
- **AI elevates work rather than eliminating it.** Competitive advantage will come from organizations that expand employee decision capability, not those that merely automate activity.



# CHAPTER 16

# STRATEGY

## COMPETITIVE ADVANTAGE SHIFTS FROM SCALE TO COORDINATION

FOR MOST OF modern business history, scale has been the dominant strategic advantage.

Large firms could purchase inputs more cheaply, spread fixed costs across broader customer bases, invest in specialized expertise, and operate complex distribution networks that smaller competitors could not match. Economists described this advantage through economies of scale and scope.<sup>1</sup> The logic was straightforward: when capital, information, and expertise were scarce and expensive, organizations that could concentrate them would outperform those that could not.

Artificial intelligence does not eliminate scale. It changes what scale means.

The scarce resource is no longer primarily access to information or analytical capability. It is the ability to coordinate decisions quickly and coherently across an organization.

In the human and machine company, coordination speed becomes a strategic asset.

---

1. Chandler, Alfred D., Jr. *Scale and Scope: The Dynamics of Industrial Capitalism*. Harvard University Press, 1990.

## FROM INFORMATION ADVANTAGE TO DECISION ADVANTAGE

To understand why coordination now matters more than traditional scale, it is useful to consider what large organizations historically did better than small ones. They gathered information more effectively.

They employed research departments.

They maintained specialized analysts.

They funded forecasting and planning functions.

They hired consultants to synthesize external knowledge.

Information moved upward, was interpreted centrally, and decisions flowed downward. Smaller firms, lacking these capabilities, relied on local knowledge and intuition.

Artificial intelligence reduces the cost of analysis dramatically. A small firm can now access structured market research, financial modeling, competitive analysis, and operational optimization tools that previously required teams of specialists. The informational advantage of size therefore weakens.

Research on general-purpose technologies helps explain this shift. Brynjolfsson and McAfee argue that digital technologies reduce the marginal cost of information processing, enabling smaller firms to access capabilities once reserved for large enterprises.<sup>2</sup> Artificial intelligence accelerates this phenomenon because it operates directly on language and reasoning rather than merely on numerical data.

As access equalizes, strategic differentiation moves elsewhere.

If multiple firms can analyze markets effectively, advantage depends on which firm can act coherently on that analysis. This is a coordination problem, not an analytical one.

Coordination involves aligning product decisions, operational execution, marketing messaging, customer interaction, and resource allocation rapidly. It requires not merely understanding what to do but doing it without internal friction.

Organizational scholars have long argued that firm performance

---

2. Brynjolfsson, Erik, and Andrew McAfee. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton, 2014.

depends heavily on coordination efficiency. Alfred Chandler’s historical analysis of the modern corporation showed that managerial coordination capabilities determined competitive outcomes as firms grew more complex.<sup>3</sup> Artificial intelligence shifts this principle into a new context. Instead of coordinating production processes, firms now coordinate decision processes.

The faster an organization can move from understanding to action, the stronger its competitive position.

## COORDINATION SPEED AS COMPETITIVE EDGE

Artificial intelligence shortens analysis cycles. What once required weeks can now occur within hours. But shortening analysis does not automatically shorten execution. Many organizations discover that internal approval processes, communication structures, and functional boundaries still operate at older speeds.<sup>4</sup>

This creates a new competitive dynamic.

Firms capable of converting insight into coordinated action quickly outperform those that generate insight but act slowly. The bottleneck moves from knowledge to alignment.

*Shopify integrated generative AI tools directly into product, support, and merchant operations through its Shopify Magic features. Instead of separate research, marketing, and support workflows, teams could generate product descriptions, customer responses, and campaign plans immediately. The company’s goal was to allow decisions to happen in real time inside workflows rather than waiting for prepared reports. AI did not give Shopify more information than competitors. It allowed the company to coordinate action faster.<sup>5</sup>*

Research in organizational behavior supports this pattern. Studies

---

3. Chandler, Alfred D., Jr. *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press, 1977.

4. McKinsey & Company. *The State of AI in 2024: Generative AI’s Breakout Year*. McKinsey Global Institute, 2024.

5. Lütke, Tobias. “Reflections on AI and Shopify.” Internal memo publicly reported by *The Verge*, 2023.

of high-performing organizations consistently emphasize decision cycle time and cross-functional alignment as predictors of performance in complex environments.<sup>6</sup> Artificial intelligence amplifies this importance because information ceases to be the limiting factor.

Consider a pricing decision. An AI-assisted analysis may identify an opportunity immediately. However, implementing that change requires coordination among finance, sales, marketing, customer support, and operational systems. If each function operates sequentially, the opportunity may pass before action occurs. If the organization coordinates simultaneously, it captures value.

The competitive advantage therefore lies not in discovering opportunities first but in acting on them first.

Coordination speed becomes a measurable strategic capability.

## **WHY SMALLER FIRMS TEMPORARILY COMPETE ABOVE THEIR WEIGHT**

This shift initially favors smaller organizations.

Smaller firms possess structural characteristics that support rapid coordination. They have fewer hierarchical layers, fewer formal approval processes, and closer communication between decision makers. When artificial intelligence reduces analytical barriers, these firms suddenly gain capabilities previously limited to larger enterprises while retaining agility.

Entrepreneurial research has long shown that smaller organizations adapt faster to environmental change due to reduced bureaucratic complexity.<sup>7</sup> Artificial intelligence magnifies this advantage because it equips small firms with analytical resources they historically lacked.

A consulting firm of one or two professionals can now perform research, financial modeling, drafting, and scenario analysis at a level approaching that of much larger teams. A small retailer can optimize pricing and marketing using predictive tools. A software startup can

---

6. Edmondson, Amy C. *Teaming: How Organizations Learn, Innovate, and Compete in the Knowledge Economy*. Jossey-Bass, 2012.

7. Mintzberg, Henry. *Structure in Fives: Designing Effective Organizations*. Prentice-Hall, 1983.

produce technical documentation, test code, and analyze customer behavior rapidly.<sup>8</sup>

This does not mean small firms become identical to large firms. It means they temporarily overcome a specific historical disadvantage: informational capacity.

For a period, they compete above their traditional scale.

*Startup legal platform Harvey built an AI system to draft and analyze legal documents using large language models. Within a year, major law firms including Allen & Overy deployed it across thousands of lawyers to assist research and contract drafting. What mattered was not the firm's computing power but how quickly attorneys could coordinate around shared analysis. Associates could produce work in hours that previously required multi-day research chains. Partners reported that client response times accelerated because decisions could be made during a single working session rather than across multiple review cycles.<sup>9,10</sup>*

Market observations increasingly support this pattern. New companies can enter industries that once required substantial support infrastructure because artificial intelligence reduces the cost of planning, coordination, and communication. The barrier to informed participation declines.

However, this advantage is transitional rather than permanent.

## WHY ENTERPRISES EVENTUALLY REGAIN ADVANTAGE

Although smaller firms gain early competitive strength, large enterprises possess structural resources that become valuable again once coordination problems are addressed.

Large firms control capital, customer relationships, distribution networks, regulatory expertise, and operational infrastructure. These

---

8. McKinsey & Company. *The Economic Potential of Generative AI: The Next Productivity Frontier*. McKinsey Global Institute, 2023.

9. Roose, Kevin. "A.I. Is Coming for Lawyers, Again." *The New York Times*, 10 Apr. 2023.

10. Allen & Overy. "Allen & Overy Adopts Harvey AI Platform Firmwide." Press Release, 2023.

assets do not disappear because analysis becomes cheaper. They become more powerful once organizations learn to coordinate them effectively.

Historically, large firms struggled with bureaucratic inertia. Information traveled slowly, and decision cycles were long. Artificial intelligence does not automatically fix this problem, but it creates pressure to do so. Firms that redesign workflows, clarify decision rights, and adopt coordinated operating models can leverage both analytical capability and resource scale.

Research on organizational adaptation shows that established firms that successfully integrate new technologies often outperform entrants because they combine innovation with infrastructure.<sup>11</sup> Artificial intelligence becomes another instance of this principle.

Once enterprises align their internal processes with faster decision cycles, they regain strategic advantage through coordinated deployment of resources. A large company that can interpret data quickly and act across thousands of employees simultaneously possesses a capacity smaller firms cannot replicate.

Thus the competitive sequence becomes clear.

First, small firms benefit because information democratizes.

Then, large firms adapt coordination.

Finally, scale returns as advantage, but in a new form.

The enterprise does not win by possessing better information. It wins by mobilizing its organization effectively.

## STRATEGIC IMPLICATIONS FOR LEADERS

This shift has several implications for strategy.

Leaders should not view artificial intelligence primarily as a cost reduction tool. Its strategic impact lies in enabling faster and more coherent action. The central question is not “How many tasks can we automate?” but “How quickly can we align our organization around insight?”

---

11. Teece, David J. “Dynamic Capabilities and Strategic Management.” *Strategic Management Journal*, vol. 18, no. 7, 1997, pp. 509–533.

Second, firms must examine internal processes as carefully as external markets. Competitive advantage increasingly depends on reducing internal friction rather than merely identifying opportunities.

Third, executives must recognize the temporal dimension of competition. Smaller competitors may outperform briefly due to agility. Large firms must respond by redesigning coordination rather than relying solely on scale.

The firms that succeed will treat artificial intelligence as a coordination technology rather than merely an analytical one.

## THE NEW NATURE OF SCALE

In the emerging environment, scale no longer means simply being large. It means being able to coordinate a large number of actors effectively.

Scale becomes organizational, not merely financial.

A company with thousands of employees can act like a small firm if its decision processes are aligned and responsive. A company with many resources can still move slowly if coordination remains fragmented.

Artificial intelligence does not determine which firm wins. Organizational design does.

Competitive advantage therefore shifts from accumulation of resources to orchestration of resources. Firms that coordinate faster, learn faster, and act faster will outperform those that merely know more.

Scale still matters.

Coordination now determines whether scale helps or hinders.

## KEY TAKEAWAYS

- **Scale advantage is changing, not disappearing.** Artificial intelligence reduces the informational advantage large firms historically held, shifting competition away from who knows more toward who can act faster.

- **Coordination speed becomes the strategic differentiator.** Competitive performance increasingly depends on how quickly an organization can align decisions across functions, not how quickly it can produce analysis.
- **Insight is no longer scarce; alignment is.** Many firms can now generate similar market intelligence, but firms that convert understanding into coordinated action first capture the value.
- **Small firms gain early advantage through agility.** With AI expanding analytical capacity, smaller organizations temporarily compete above their traditional weight because they can coordinate decisions rapidly.
- **Large enterprises regain advantage after redesign.** Once big organizations clarify decision rights and streamline internal processes, they combine speed with capital, distribution, and customer scale.
- **AI is a coordination technology, not just an analytics technology.** The central strategic question becomes how quickly the organization can mobilize around insight rather than how accurately it can produce it.
- **The new meaning of scale is orchestration.** Winning firms will not simply possess more resources; they will coordinate people, systems, and decisions more coherently than competitors.

# CHAPTER 17

## THE NEW FIRM

### DESIGNING ORGANIZATIONS FOR HUMANS AND AGENTS

EVERY ERA of business is defined not only by the technologies firms use but by the shape of the firm itself.

The industrial era produced the hierarchical corporation. Work was divided into specialized functions. Information moved upward through layers of supervision. Decisions moved downward through chains of command. Frederick Taylor's scientific management formalized the separation between planning and execution. Managers analyzed and designed work. Labor carried it out.<sup>1</sup> Efficiency depended on control of knowledge.

The information age digitized communication and introduced matrix structures, project teams, and global coordination, but the underlying architecture remained recognizable. Functions persisted. Reporting lines governed authority. Human beings remained the primary processors of information. Software stored and transmitted data, but interpretation and reasoning were overwhelmingly human responsibilities.

Artificial intelligence alters that premise.

For the first time in the history of the firm, non-human systems

---

1. Taylor, Frederick Winslow. *The Principles of Scientific Management*. Harper & Brothers, 1911.

participate meaningfully in reasoning. Not as passive databases. Not as static automation scripts. But as agents that interpret, recommend, prioritize, and in some cases initiate action.

This is not an incremental tool upgrade. It is a structural inflection.

The firm must now be designed as a coordinated system of humans and machine reasoning agents.

That requires rethinking not only workflow, but the architecture of authority itself.

## **ORGANIZATIONAL STRUCTURE IMPLICATIONS**

An organizational chart appears to be a map of authority. In practice, it is a map of information specialization.

Marketing exists because customer insight requires dedicated interpretation. Finance exists because capital analysis requires specialized judgment. Operations exists because process reliability demands domain knowledge. Each function historically processed a different stream of information. Hierarchy coordinated those streams into coherent action.

Artificial intelligence softens the boundary between those streams.

When a product manager can generate financial scenario modeling interactively, the informational asymmetry between product and finance narrows. When an operations leader can analyze customer sentiment data without waiting for a centralized analytics team, the dependency structure changes. When a sales leader can receive machine-supported pricing guidance in real time, negotiation authority shifts closer to the edge.

This does not eliminate functional specialization. It alters its purpose.

Functions increasingly exist not to produce analysis, but to steward decision standards.

Finance shifts from generating forecasts to defining acceptable capital thresholds and scenario governance. Marketing shifts from crafting messaging line by line to governing brand positioning rules and customer segmentation logic. Operations shifts from compiling

performance data to defining acceptable process tolerances and escalation protocols.

The CarMax example illustrates this shift concretely. When generative AI began drafting vehicle descriptions and summarizing inspection reports, human specialists did not disappear. Their work changed. They focused on defining disclosure standards, validating edge cases, and setting pricing accuracy thresholds.<sup>2</sup> The department moved from producing text to governing publication authority.

The broader implication is structural.

Departments become stewards of decision rules rather than factories of information output.

Historically, Alfred Chandler described the rise of the managerial corporation as a response to coordination complexity.<sup>3</sup> Artificial intelligence introduces a new form of complexity: coordination between human and non-human reasoning actors. The organizational problem shifts from managing human specialization to managing hybrid intelligence systems.

Reporting lines alone will not describe how work occurs.

Workflows will increasingly cross formal structures continuously. A customer exception decision may integrate automated risk scoring, human negotiation judgment, compliance rule checking, and system-based pricing constraints simultaneously. Authority must be distributed across these actors deliberately.

Organizational charts will continue to define accountability. They will no longer define activity.<sup>4</sup>

## THE NEW OPERATING MODEL

An operating model defines how an organization converts signals into decisions and decisions into coordinated action. It is not simply a process map. It is the temporal architecture of the firm.

---

2. Nellis, Stephen. "CarMax Uses Generative AI to Write Car Listings." *Reuters*, 2023.

3. Chandler, Alfred D., Jr. *The Visible Hand: The Managerial Revolution in American Business*. Harvard University Press, 1977.

4. Yoo, Youngjin, et al. "Organizing for Innovation in the Digitized World." *Organization Science*, vol. 23, no. 5, 2012, pp. 1398–1408.

For most of the twentieth century, the operating model of large organizations was built around informational scarcity. Analysis was expensive. Data was fragmented. Coordination required physical proximity or formal reporting cycles. As a result, work occurred in batches. Information was gathered over time, synthesized into reports, reviewed in scheduled meetings, and translated into decisions that were then communicated downward for execution.

Even after digitization, this structure persisted. Enterprise systems accelerated data storage and communication, but they did not fundamentally alter the sequencing of work. Monthly reporting cycles remained. Quarterly planning cadences endured. Annual budgeting processes dominated capital allocation. Digital tools made the machinery more efficient, but the rhythm remained periodic.

Artificial intelligence disrupts that temporal structure.

When analysis can be generated interactively rather than assembled manually, the rationale for periodic batching weakens. When scenario modeling can be recomputed instantly as assumptions change, the logic of static planning cycles becomes questionable. When anomaly detection systems can surface irregularities continuously, waiting for retrospective reporting becomes inefficient.

Research on high-reliability organizations demonstrates that performance improves when feedback loops shorten and decision authority is exercised closer to the point of operational activity.<sup>5</sup> Lean manufacturing principles similarly emphasize the reduction of batch sizes and the movement toward continuous flow.<sup>6</sup> Artificial intelligence introduces the possibility of continuous flow in knowledge work.

This shift is not merely about speed. It is about structural redesign.

In a continuous operating model, monitoring becomes persistent. Systems track relevant variables in real time. Signals are evaluated against predefined thresholds. When conditions exceed those thresholds, escalation occurs automatically. Humans intervene not to compile information, but to interpret significance and exercise judgment.

---

5. Weick, Karl E., and Kathleen M. Sutcliffe. *Managing the Unexpected: Sustained Performance in a Complex World*. 3rd ed., Wiley, 2015.

6. Womack, James P., and Daniel T. Jones. *Lean Thinking: Banish Waste and Create Wealth in Your Corporation*. Simon & Schuster, 1996.

Consider the difference between retrospective risk reporting and real-time risk governance. In a traditional model, a risk committee reviews quarterly summaries of exposure metrics. In a continuous model, AI-driven monitoring systems flag unusual patterns immediately. Risk officers evaluate context and decide whether to intervene. Authority is triggered by conditions rather than calendar events.

This same pattern applies to revenue management, supply chain optimization, workforce allocation, and capital deployment. The operating model moves from event-scheduled decision making to condition-triggered decision making.

However, the transition is not automatic.

If organizations retain meeting-centered approval structures while introducing real-time analytics, the informational gain is neutralized. Data accumulates more quickly, but decisions still wait for formal review sessions. In such cases, artificial intelligence improves preparation but not outcome.

The new operating model therefore requires explicit specification of decision rights.

Which decisions are fully automated within guardrails?

Which are machine-assisted but require human validation?

Which require collective executive deliberation regardless of informational availability?

Research on hybrid intelligence systems emphasizes that effective human-AI collaboration depends on clearly defined interaction protocols and authority boundaries.<sup>7</sup> Without those definitions, organizations oscillate between over-automation and unnecessary escalation.

The operating model of the AI-enabled firm becomes a coordination architecture among four elements: continuous monitoring, threshold-based escalation, human judgment, and accountable authority. Each must be designed intentionally.

Artificial intelligence does not eliminate hierarchy. It changes its timing.

---

7. Dellermann, Dominik, et al. "The Future of Human-AI Collaboration: A Taxonomy of Design Knowledge for Hybrid Intelligence Systems." *Business & Information Systems Engineering*, vol. 61, 2019, pp. 363–376.

The organization shifts from a cadence of information gathering followed by decision to a state of persistent situational awareness punctuated by authority intervention.

That is a structural change, not a technological one.

## THE FUTURE MANAGEMENT STACK

Historically, the management stack of the firm consisted of layered supervision built atop process and system infrastructure.

At the base were transactional systems that recorded activity. Above them were reporting systems that summarized information. Above those sat managers who interpreted summaries and made decisions. The managerial layer existed because interpretation required human cognition and because access to underlying data was limited.

Artificial intelligence introduces a new infrastructural layer between data and management: algorithmic reasoning.

Systems no longer merely store and transmit information. They evaluate patterns, simulate consequences, prioritize options, and in some contexts initiate action. This insertion alters the logic of supervision.

The management stack now includes at least four interdependent layers.

First, data infrastructure, which governs how information is captured, structured, and accessed.

Second, reasoning systems, which interpret that information and generate recommendations.

Third, workflow orchestration, which determines how recommendations translate into action pathways.

Fourth, human accountability, which assigns responsibility for outcomes.

The introduction of reasoning systems creates a new managerial obligation: the governance of system behavior.

Managers must determine acceptable error tolerances, escalation thresholds, override authority, and validation procedures. This is not purely a technical design exercise. It is an organizational one.

Research on enterprise systems has long demonstrated that tech-

nology delivers performance gains only when organizational processes adapt in parallel.<sup>8</sup> Artificial intelligence intensifies this dependency because it operates directly on decision-making logic.

If reasoning systems are introduced without redesigning escalation rules, confusion results. Employees may not know whether machine-generated recommendations are advisory or binding. If autonomy levels are unclear, risk exposure increases. If accountability for machine-supported outcomes is ambiguous, trust erodes.

The future management stack therefore requires clarity at each layer.

Data governance must ensure lineage, integrity, and accessibility.

Reasoning systems must be explainable within acceptable tolerance bands.<sup>9</sup>

Workflow orchestration must encode authority structures explicitly.

Human accountability must be documented and visible.

Over time, managerial attention shifts upward in the stack. Less time is spent reconciling data discrepancies. More time is spent calibrating system behavior and refining authority boundaries.

Supervision evolves from oversight of human compliance to oversight of socio-technical interaction.

This aligns with sociotechnical systems theory, which emphasizes that performance emerges from the joint optimization of social and technical subsystems.<sup>10</sup> Artificial intelligence deepens this interdependence. Managers must design not only human teams but hybrid teams composed of humans and agents.

The management stack does not flatten. It transforms.

Authority becomes less about controlling information and more about defining permissible action within dynamic systems.

---

8. Brynjolfsson, Erik, and Lorin M. Hitt. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives*, vol. 14, no. 4, 2000, pp. 23–48.

9. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.

10. Pasmore, William A., et al. "Sociotechnical Systems: A North American Reflection on Empirical Studies of the Seventies." *Human Relations*, vol. 35, no. 12, 1982, pp. 1179–1204.

## WHAT A COMPANY WILL LOOK LIKE IN TEN YEARS

Projecting forward is not speculation about sentient machines. It is extrapolation from observable trends in information processing.

Ten years from now, the most visible difference inside firms will not be robots walking hallways. It will be the compression of time between signal and decision.

Planning cycles will shorten materially. Annual strategic reviews will remain for governance purposes, but operational reallocation will occur continuously. Finance teams will rely on rolling forecasts supported by live scenario simulation rather than static models assembled quarterly. Capital committees will evaluate investment shifts based on continuously updated projections rather than periodic slide decks.

Meetings will change character. Preparation will no longer require days of assembling and formatting data. Participants will arrive with interactive analysis already available. Discussion will focus on trade-offs, risk appetite, and authority rather than information validation.

The composition of roles will shift gradually but significantly.

Entry-level roles that once focused on compiling reports will evolve toward interpreting machine-supported insights. Junior employees will be expected to exercise judgment earlier because cognitive scaffolding will be widely available.

Mid-level managers will spend less time supervising task completion and more time calibrating decision thresholds. Their authority will derive from their responsibility for outcomes within defined domains rather than from informational gatekeeping.

Senior executives will increasingly evaluate system performance rather than individual output. Dashboards will reflect decision latency, escalation frequency, and threshold calibration effectiveness.

Operational monitoring will become continuous. Artificial intelligence agents will track supply chain disruptions, customer behavior anomalies, credit exposure shifts, and workforce utilization patterns in real time. Routine adjustments will occur automatically within predefined limits. Escalation will be triggered by deviation from acceptable bounds rather than by scheduled review.

Research on digital platform ecosystems suggests that competitive advantage increasingly derives from the orchestration of interconnected actors rather than isolated capabilities.<sup>11</sup> In a human-agent firm, orchestration extends to algorithmic actors as well.

The physical experience of work will feel more responsive and less bureaucratic.

However, this future is contingent on intentional design.

If authority boundaries remain ambiguous, speed will produce confusion rather than clarity. If governance frameworks lag, regulatory incidents will erode trust. If cultural adaptation fails, employees will resist hybrid workflows.

The defining characteristic of the successful firm will not be the raw intelligence of its systems. It will be the coherence of its authority structure.

During Moderna's use of machine learning in drug discovery, computational models were integrated directly into research workflows, guiding experiment prioritization continuously rather than sequentially.<sup>12</sup> Scientists described their role as supervising and refining model-generated hypotheses rather than manually generating each experimental path. The cadence of research changed. Authority shifted closer to real-time evaluation.

Extend that pattern across finance, operations, customer service, product design, and compliance, and the structural consequences become visible.

The firm of the next decade will be characterized by persistent situational awareness, shorter decision loops, clearer authority thresholds, and hybrid teams in which human judgment and machine reasoning are deliberately integrated.

The organization will not disappear.

It will become more explicit.

Artificial intelligence does not remove the need for management.

---

11. Iansiti, Marco, and Karim R. Lakhani. *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Review Press, 2020.

12. Loftus, Peter. "Moderna Bets Big on Artificial Intelligence to Speed Drug Discovery." *The Wall Street Journal*, 27 July 2023.

It forces management to become more intentional about how decisions are made.

## DESIGNING THE INTENTIONAL FIRM

Artificial intelligence does not dictate the future shape of the organization. It exposes it.

For more than a century, firms were organized around the scarcity of information. Knowledge was costly to gather and slow to distribute. Hierarchies existed not only to control labor but to process uncertainty. Supervisory layers aggregated data upward. Reports reduced ambiguity. Planning cycles existed because analysis required time.

Artificial intelligence alters that constraint.

When analysis becomes abundant and interactive, the bottleneck shifts. Information is no longer scarce. Judgment becomes the scarce resource. Authority design becomes the performance lever.

The firm of the coming decade will not succeed because it possesses the most advanced models. It will succeed because it has redesigned how decisions move.

That redesign requires deliberate choices.

Leaders must specify which decisions can occur autonomously and within what thresholds. They must define when machine-supported outputs are binding and when they are advisory. They must redesign planning cadence to reflect compressed analytical cycles. They must clarify accountability in hybrid workflows where human and machine reasoning are intertwined.

This is not a technological project.

It is an architectural one.

The central question is no longer “What can artificial intelligence do?” It is “How should authority be structured when intelligence is abundant?”

Firms that treat artificial intelligence as a tool will optimize locally. Firms that treat it as a structural condition will redesign globally.

The difference will not be incremental.

It will be institutional.

# CHAPTER 18

# THE CIO'S

# RESPONSIBILITY

## REDESIGNING AUTHORITY IN THE AGE OF ABUNDANT INTELLIGENCE

IF YOU ARE THE CIO, artificial intelligence is not simply a capability upgrade. It is a structural stress test of your operating model.

The board wants assurance that the company is not falling behind.

The CEO wants competitive leverage.

The CFO wants economic clarity.

Risk leaders want containment.

Business units want acceleration.

Your teams are already experimenting.

Where your organization sits in the maturity curve determines the texture of these pressures. A mid-sized firm in Stage One faces ambiguity and experimentation risk. A global enterprise in Stage Two faces pilot proliferation and structural inertia. A firm entering Stage Three faces authority recalibration and managerial redesign.

Before acting, you must locate yourself honestly.

## A CIO DIAGNOSTIC

You can identify your stage by examining authority rather than activity.

- Has any recurring, cross-functional decision moved materially faster because machine-supported analysis reduced uncertainty?
- Have any approval layers been formally removed in the past twelve months due to improved information quality?
- Do your capital allocation cycles reflect faster analytical preparation, or do they still assume weeks of manual synthesis?
- Do you measure decision latency anywhere in your enterprise?
- Are AI investments framed as workflow redesign or as software procurement?

If intelligence has increased but authority has not shifted, you are likely in Stage One or Stage Two. If authority structures are changing, even in isolated domains, you may be entering Stage Three.

The stage determines your responsibility.

In Stage One, your task is legitimacy and visibility.

In Stage Two, your task is structural courage.

In Stage Three, your task is systemic recalibration.

But beyond stages, the CIO role itself changes across several core dimensions.

## **THE CIO ACROSS CORE DIMENSIONS**

Artificial intelligence does not simply alter workflows. It alters how you perform your job across the boardroom, the finance office, the risk function, the talent model, the architecture stack, the vendor ecosystem, and the CEO relationship.

Each of these dimensions evolves across the maturity stages.

### **The CIO and the Board**

The relationship between the CIO and the board has historically been episodic.

In many organizations, the CIO appeared before the board

primarily in moments of risk. Cybersecurity incidents. Major system outages. Large modernization investments. Occasionally, digital transformation updates. The tone was often defensive or operational. Assurance was the goal.

Artificial intelligence changes both the frequency and the substance of that interaction.

Boards now see artificial intelligence as strategically material. Competitive narratives, market valuations, and media attention have elevated it from technical curiosity to board-level priority. Directors are asking not only whether the company is secure, but whether it is structurally positioned for a world of abundant intelligence.

This shift alters the CIO's presence in the boardroom.

In Stage One organizations, the board's primary posture is inquisitive but externally influenced. Directors read about AI-driven disruption. They observe competitors announcing partnerships and model deployments. They hear earnings calls where peers reference AI productivity gains.

The board's first concern is positioning.

Are we behind?

Are we exposed?

Are we missing something structural?

At this stage, the CIO must resist the temptation to present ambition as maturity. Overstating progress in Stage One creates a credibility trap later. Research consistently shows that many organizations overestimate the enterprise impact of early AI experimentation.<sup>1</sup> Boards are increasingly aware of this gap between adoption and measurable performance.

The appropriate posture in Stage One is disciplined realism. The board must understand that experimentation is occurring within defined guardrails, that informal adoption is being surfaced rather than suppressed, and that risk frameworks are evolving alongside capability. The message is not that transformation is complete. The message is that learning is structured.

---

1. McKinsey & Company. "The State of AI in 2024: Generative AI's Breakout Year." *McKinsey Global Survey on AI*, 30 May 2024, McKinsey & Company.

Stage Two introduces a different tension.

Here, artificial intelligence activity is visible across the enterprise. Pilots exist. Use cases are documented. Vendor partnerships have been announced. The board has heard updates for several quarters.

Yet financial and operational metrics remain largely stable.

This is where skepticism begins to surface.

Directors may not articulate it bluntly, but the underlying question emerges: if this technology is transformative, why has the organization not moved?

The CIO's responsibility in Stage Two is structural candor.

Artificial intelligence does not automatically produce enterprise-level performance change. Decades of research on technology adoption confirm that complementary organizational redesign is required for measurable productivity gains.<sup>2</sup> If approval layers remain intact and decision rights unchanged, machine-supported analysis accelerates preparation but not outcomes.

The board must understand that the constraint is not tool capability but authority design.

This is politically delicate. It implies that transformation requires cross-functional realignment. It may require changes in governance committees, escalation pathways, and executive decision cadence. It may challenge historical power structures.

Yet without this candor, the board will continue to evaluate AI maturity through superficial metrics. Number of use cases. Percentage of employees trained. Vendor contracts signed. None of these guarantee structural movement.

Stage Three changes the board conversation again.

When at least one recurring, material decision pathway has been redesigned and authority recalibrated, the board's focus shifts toward durability and governance.

Is the new decision structure auditable?

Are risk thresholds clearly defined?

---

2. Brynjolfsson, Erik, and Lorin M. Hitt. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives*, vol. 14, no. 4, 2000, pp. 23–48.

Is vendor dependency manageable?

Are gains sustainable under regulatory scrutiny?

At this stage, artificial intelligence is no longer treated as innovation theater. It becomes part of operating model oversight.

The CIO's presence in the boardroom evolves accordingly. You are no longer reporting on experimentation. You are explaining how authority has shifted, how decision latency has compressed, and how governance has adapted.

The political dynamics also change.

Artificial intelligence introduces transparency into knowledge work. When analytical preparation becomes faster and more accessible, inefficiencies in coordination become more visible. Directors may begin asking sharper questions about management layers, decision bottlenecks, and capital allocation cadence.

The CIO must navigate this carefully. The goal is not to expose executive peers. The goal is to frame structural redesign as enterprise maturation rather than individual shortcoming.

Another dimension of the CIO–Board relationship concerns risk asymmetry.

Boards are acutely sensitive to downside exposure. Data breaches, biased model outputs, regulatory penalties, and reputational harm carry significant weight. At the same time, directors fear strategic stagnation. Falling behind competitors is also a risk.

The CIO must therefore present AI maturity as balanced progression. Guardrails evolve alongside capability. Governance frameworks mature in tandem with authority redesign. References to established frameworks such as the NIST AI Risk Management Framework provide directors with recognizable anchors.<sup>3</sup>

Board confidence increases when experimentation is framed within discipline and redesign within accountability.

The most subtle shift, however, is conceptual.

Historically, boards evaluated technology primarily through the lens of cost, resilience, and risk containment. Artificial intelligence

---

3. National Institute of Standards and Technology. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. U.S. Department of Commerce, 2023.

forces boards to consider technology as a determinant of operating structure.

If intelligence is abundant and near-instant, what does that imply for planning cadence?

For capital deployment?

For escalation authority?

For managerial layers?

These are no longer purely management questions. They are governance questions.

The CIO's credibility in the AI era depends not on technological fluency alone, but on the ability to articulate how intelligence abundance changes the structure of the firm.

The failure mode in the CIO–Board relationship is misalignment of narrative.

If AI is framed as transformative without structural evidence, trust erodes.

If AI is framed as risky without strategic ambition, momentum stalls.

If AI updates focus on tools rather than authority, directors lose interest.

Boards do not ultimately care about model architecture.

They care about how the firm decides, competes, and protects itself.

Artificial intelligence sits at the intersection of those concerns.

The CIO who understands that shift moves from operational presenter to structural advisor.

And in an era where intelligence is abundant, structural clarity is what boards seek most.

## The CIO and the CFO

No relationship will shape the trajectory of artificial intelligence in your enterprise more than the one between you and the CFO.

For decades, the CIO–CFO dynamic has been structured around discipline and constraint. Technology proposals were evaluated through capital expenditure frameworks. Business cases were built on cost reduction, risk mitigation, or incremental revenue enablement.

Multi-year depreciation schedules aligned with infrastructure lifecycles. Operating expense was scrutinized for predictability. The CFO's primary concern was variance control.

Artificial intelligence complicates this structure because it does not fit neatly into traditional capital logic.

In Stage One, AI appears as experimentation. Small subscriptions. Pilot budgets. Innovation sandboxes. From the CFO's vantage point, this can look like uncontrolled proliferation. SaaS creep. Shadow spending. Unclear ROI. The financial risk is not existential but accumulative.

At this stage, your responsibility as CIO is to frame AI investment as bounded optionality rather than open-ended expense. The CFO must see experimentation as disciplined learning, not as discretionary drift. Clear budget envelopes, transparent reporting, and explicit guardrails signal control. Without that structure, the CFO will rationally constrain exploration.

Yet the deeper tension begins in Stage Two.

As pilots proliferate and AI capability becomes visible across functions, the CFO will begin asking a harder question: where is the financial impact?

This is where many CIOs falter.

Productivity anecdotes do not satisfy capital governance. If analysts draft faster but headcount remains static, the CFO sees no measurable change in cost structure. If models produce insights but approval chains remain unchanged, revenue acceleration is not visible. Artificial intelligence risks being categorized as marginal productivity software rather than structural leverage.

Research on digital productivity has repeatedly demonstrated that technology alone does not produce measurable performance gains unless accompanied by complementary organizational change.<sup>4</sup> The CFO's skepticism in Stage Two is often rational. Without workflow redesign, financial metrics remain stable.

---

4. Brynjólfsson, Erik, Wang Jin, and Sebastian Steffen. "Do IT Capabilities Still Drive Productivity and Innovation in the New Economy?" *Stanford Digital Economy Lab*, 15 Mar. 2024.

Your responsibility at this stage is translation.

You must connect AI investment not to tool adoption, but to structural movement in decision cycles, working capital efficiency, risk exposure, and revenue velocity. If machine-supported underwriting compresses approval timelines, that affects revenue recognition. If fraud detection reduces loss rates through redesigned escalation thresholds, that affects margin stability. If supply chain exception management accelerates, inventory holding costs shift.

The conversation must move from cost of licenses to cost of latency.

Latency has financial consequences. Slow decisions tie up capital. Prolonged approval cycles delay revenue capture. Manual reconciliation inflates operational expense. Artificial intelligence reduces analytical preparation time. If authority structures adjust accordingly, financial cycles can compress.

Stage Three introduces a more profound shift.

When analytical preparation becomes near-instant, traditional annual planning cycles begin to look misaligned. Historically, strategy and capital allocation operated on annual cadences because scenario modeling and forecasting were time-intensive. Extensive manual analysis constrained planning frequency.

Artificial intelligence alters that constraint.

Scenario modeling can be continuous. Sensitivity analyses can be generated in hours rather than weeks. Forecast revisions can incorporate real-time signals. The CFO may begin to expect shorter planning cycles and more dynamic reallocation of capital.

This is not merely a tooling question. It is a governance redesign.

If planning remains annual while analytical capability is continuous, the enterprise absorbs intelligence without exploiting it. If planning becomes rolling and iterative, capital velocity increases.

The CIO must therefore collaborate with the CFO to examine whether financial governance structures reflect informational scarcity or informational abundance.

Old model budgeting optimized for predictability.

New model budgeting may need to optimize for adaptability.

This does not imply recklessness. It implies structural responsiveness.

Artificial intelligence forces a reconsideration of how capital is evaluated, deployed, and redeployed. The CIO who frames AI as an expense center will remain constrained. The CIO who frames AI as a lever on decision velocity and capital efficiency becomes indispensable.

## The CIO and Risk

Risk is where artificial intelligence becomes real.

For decades, enterprise risk models were built on procedural control. Decisions moved through defined checkpoints. Documentation demonstrated diligence. Approval chains reduced uncertainty by slowing action. Compliance was evidenced by traceability and manual oversight.

Artificial intelligence destabilizes that architecture.

In Stage One, the risk function's first instinct is defensive. Concerns about data leakage, hallucinated outputs, intellectual property exposure, and regulatory noncompliance dominate early conversations. In regulated industries, these concerns are not hypothetical. They are existential.

The temptation at this stage is prohibition.

Yet research on shadow IT behavior demonstrates that when employees perceive tools as productivity enhancing, outright bans do not eliminate usage. They displace it.<sup>5</sup> Informal experimentation continues in private environments, increasing opacity. The CIO's responsibility in Stage One is therefore not maximal restriction but controlled legitimacy. Guardrails must be defined clearly enough to satisfy risk leadership while preserving visibility into behavior.

This requires partnership. Risk leaders must understand that artificial intelligence is already embedded in daily work, whether formally sanctioned or not. The CIO must articulate how data boundaries, human verification requirements, and accountability standards mitigate exposure without suppressing experimentation.

---

5. Györy, Attila, et al. "Shadow IT and the Human Factor: Understanding Informal Technology Use in Organizations." *MIS Quarterly Executive*, vol. 21, no. 2, 2022.

Stage Two introduces a more complex challenge.

At this stage, artificial intelligence is visible and embedded in defined use cases. The question is no longer whether AI can be used. The question becomes whether AI-supported decisions can act.

Historically, escalation existed because information quality was uneven. Human review compensated for analytical uncertainty. When machine-supported analysis improves signal quality, escalation thresholds must be recalibrated. But recalibration is politically and institutionally difficult.

Frameworks such as the NIST AI Risk Management Framework emphasize explainability, accountability, and traceability as foundational to trustworthy AI systems (National Institute of Standards and Technology). In practice, this means that decision pathways must embed auditability directly into workflow design. If a machine-supported decision influences credit approval, underwriting, pricing, or clinical prioritization, its inputs, logic, and human oversight must be reconstructable.

The CIO must therefore shift risk discussions from “Can we use AI?” to “Under what authority can AI-supported decisions operate autonomously within defined tolerances?”

Stage Three brings an even more profound shift.

When authority begins to move closer to the edge of the organization, risk governance must adapt accordingly. Autonomy increases. Escalation layers thin. Machine-supported decisions occur more frequently and more quickly.

In this environment, risk is not reduced by slowing the system. It is managed by clarifying thresholds and accountability.

High reliability research demonstrates that performance in complex systems improves when actors operate with clear authority within defined boundaries, supported by rapid feedback loops.<sup>6</sup> Artificial intelligence increases the speed of those loops. The CIO must ensure that governance frameworks evolve accordingly.

The ultimate failure mode in this dimension is assuming that

---

6. Weick, Karl E., and Kathleen M. Sutcliffe. *Managing the Unexpected: Sustained Performance in a Complex World*. 3rd ed., Wiley, 2015.

compliance frameworks designed for informational scarcity remain sufficient in an era of abundant cognition.

Artificial intelligence does not eliminate risk. It changes its structure.

The CIO's role is to redesign governance so that speed and accountability coexist.

## The CIO and Talent

Talent is where artificial intelligence becomes political.

Boards and CEOs may publicly emphasize innovation, but privately they ask a simpler question: does this reduce headcount?

The honest answer is more complex than either enthusiasm or fear suggests.

In Stage One, artificial intelligence augments individuals. Analysts draft faster. Developers navigate code more easily. Product managers synthesize research more efficiently. The visible impact is productivity at the margin. Headcount remains unchanged.

At this stage, the CIO's responsibility is literacy. Teams must learn how to interrogate output, validate reasoning, and recognize limitations. Research suggests that performance gains from AI tools are amplified when users understand both strengths and constraints.<sup>7</sup> Education must therefore focus not on prompts alone, but on judgment.

Stage Two introduces structural tension.

As pilots accumulate and productivity anecdotes multiply, executive leadership begins to ask whether workforce levels should adjust. Mishandled, this conversation creates cultural resistance. Employees interpret AI as displacement rather than redesign.

The deeper reality is that artificial intelligence shifts the composition of work more than the volume of work. Tasks centered on information gathering and formatting diminish in importance. Tasks

---

7. Brynjolfsson, Erik, Danielle Li, and Lindsey R. Raymond. "Generative AI at Work." *The Quarterly Journal of Economics*, vol. 140, no. 2, May 2025, pp. 889–942.

centered on interpretation, orchestration, and cross-functional coordination increase.

Research on digital transformation emphasizes that technology complements skilled labor rather than uniformly replacing it.<sup>8</sup> Organizations that treat AI solely as a labor reduction mechanism often undermine the very redesign required for performance gains.

Stage Three alters the managerial layer itself.

In organizations built for informational scarcity, supervisors often existed to aggregate and transmit knowledge upward. When information becomes widely accessible, that aggregation role weakens. Management layers designed primarily for reporting coordination may flatten.

This does not eliminate leadership. It changes its function.

Managers shift from supervising task execution to orchestrating decision processes. They must interpret probabilistic output, challenge machine-generated recommendations, and coordinate rapid cross-functional action.

The CIO must anticipate these shifts. Workforce planning cannot remain incremental. Role definitions must evolve. Career pathways must reflect the movement from information production to judgment design.

The most subtle talent failure mode is cultural misframing. If artificial intelligence is positioned as an efficiency weapon, resistance rises. If it is positioned as a coordination redesign tool, adaptation becomes possible.

The CIO sits at the intersection of capability and culture.

## The CIO and Enterprise Architecture

Enterprise architecture has always been the quiet determinant of what is possible.

In most organizations, architecture evolves incrementally. Systems accumulate. Acquisitions introduce redundancy. Data models prolifer-

---

8. Iansiti, Marco, and Karim R. Lakhani. *Competing in the Age of AI: Strategy and Leadership When Algorithms and Networks Run the World*. Harvard Business Review Press, 2020.

ate. Integration layers multiply. Each decision is rational in isolation. Over time, the aggregate becomes complex, brittle, and deeply path dependent. The CIO inherits this inheritance of history.

For years, modernization efforts have aimed to simplify that inheritance. Cloud migration programs, API strategies, data lake initiatives, and identity governance reforms were justified as necessary investments in agility and resilience. Yet in many firms, these programs were experienced as cost centers rather than competitive levers. Architecture was treated as foundational plumbing rather than strategic infrastructure.

Artificial intelligence alters that perception.

In Stage One organizations, architectural limitations appear first as inconvenience. Teams experiment with language models or analytics tools and discover that access to high-quality data is inconsistent. Developers find that machine-assisted tooling performs well when codebases are modular and poorly when legacy monoliths dominate. Analysts attempting to generate real-time insights discover that critical data sets are locked behind batch processes designed for overnight reconciliation rather than continuous access.

At this stage, architecture constrains quietly. The organization is still experimenting. Limitations are frustrating but not yet decisive. The CIO's role is observational and diagnostic. Where does experimentation repeatedly stall because data cannot be accessed cleanly? Where do integration bottlenecks slow adoption? Where do identity and access models prevent responsible scaling? These friction points signal where architectural redesign will eventually become unavoidable.

Stage Two brings architectural reality into sharp focus.

As pilots expand and business leaders seek operational integration, the limits of legacy architecture become structural barriers. Machine-supported analysis is technically feasible, but it cannot be embedded into core workflows because data remains fragmented across domains. Latency persists because pipelines were designed for periodic reporting rather than event-driven action. Data lineage is unclear, complicating explainability and regulatory review.

Research on digital productivity has long emphasized that the performance effects of technology depend on complementary organi-

zational and infrastructural investments (Brynjolfsson and Hitt). Artificial intelligence intensifies this dependency. Models can generate analysis at extraordinary speed, but if underlying data contracts are inconsistent or if integration pathways are brittle, the system as a whole remains slow.

In many enterprises, Stage Two exposes a mismatch between analytical capability and architectural readiness. The organization may possess sophisticated models, but it lacks the ability to deliver inputs and capture outputs in a way that meaningfully alters decision pathways.

This is where the CIO's responsibility becomes structural rather than incremental.

Architectural conversations must shift from modernization for efficiency to architecture as an enabler of authority redesign. If the ambition is to shorten credit decision cycles from days to hours, then data ingestion, validation, and model integration must operate within that temporal window. If the ambition is to automate aspects of supply chain exception management, then event-driven architectures must support near-real-time visibility across systems that were previously reconciled only periodically.

The tension is not merely technical. It is historical. Many legacy systems were designed around batch processing because compute was expensive and real-time coordination was unnecessary. Approval layers compensated for informational delay. Artificial intelligence reduces the cognitive cost of analysis. If architecture does not reduce the structural delay of data movement, the organization remains trapped between new capability and old rhythm.

Stage Three requires an even deeper architectural evolution.

Once authority begins to shift and machine-supported decisions are embedded in operational workflows, architecture must support not only speed but accountability. Data lineage must be explicit. Model inputs and outputs must be traceable. Integration must be resilient under load. Explainability cannot be an afterthought layered onto opaque pipelines. It must be designed into the system itself (National Institute of Standards and Technology).

In this stage, architecture becomes inseparable from governance.

The CIO must ensure that machine-in-the-loop workflows are auditable, that event streams are reliable, and that decision pathways are observable. This often requires rethinking identity models, logging infrastructure, and API abstraction layers. It may require modularizing legacy systems so that intelligence can be inserted without destabilizing core operations.

There is also a strategic dimension to architectural maturity.

Artificial intelligence capabilities evolve rapidly. Vendor ecosystems shift. Pricing structures fluctuate. An architecture that is tightly coupled to a single proprietary stack may accelerate short-term deployment but create long-term rigidity. Conversely, excessive abstraction may slow integration and reduce performance.

The CIO must balance integration depth with optionality. Architectural modularity becomes not merely an engineering preference but a strategic hedge against capability volatility. In an environment where model performance and cost curves change frequently, switching costs can become competitive liabilities.

The failure mode in this dimension is subtle but common.

Organizations often pursue surface modernization while leaving foundational data contracts untouched. They migrate infrastructure to the cloud but preserve fragmented domain ownership. They implement APIs without clarifying semantic consistency. They deploy AI tools on top of data lakes that lack governance discipline. The result is analytical activity without operational coherence.

Artificial intelligence does not eliminate architectural weakness. It amplifies it.

If cognition becomes abundant but data remains inconsistent, trust erodes. If analysis becomes instantaneous but integration remains slow, authority does not shift. If decision pathways rely on opaque pipelines, regulatory and reputational risk increases.

Enterprise architecture, in the age of artificial intelligence, is no longer a background modernization agenda. It is the substrate upon which decision velocity, accountability, and competitive advantage rest.

For the CIO, this is not simply a technical mandate. It is a structural one.

Architecture determines whether intelligence changes how the firm decides or merely how it drafts.

## The CIO and Vendor Strategy

Artificial intelligence has destabilized the vendor landscape more quickly than any technology cycle in recent memory. For decades, CIOs optimized for stability, integration coherence, and predictable cost curves. Enterprise contracts were negotiated over multi-year horizons. Platform decisions were made carefully and infrequently. Switching costs were high, but capability velocity was moderate.

That equilibrium no longer holds.

Model performance evolves quarterly. Pricing structures fluctuate with compute markets. Vendors bundle generative capabilities into core suites. Infrastructure providers compete with application providers. Specialized startups claim breakthrough differentiation. Meanwhile, your enterprise architecture must remain coherent.

In Stage One, the vendor risk is premature consolidation. Boards and executive peers may push for a “standard AI platform” before real use patterns are understood. Locking into a single ecosystem too early constrains discovery and may misalign architecture with emerging workflow realities. At this stage, optionality is strategic. Your task is to design experimentation portfolios without fragmenting security or data governance.

In Stage Two, the vendor landscape becomes a strategic battlefield. Use cases proliferate across business units. Procurement pressure increases. The risk is duplication and shadow contracts. But the deeper risk is architectural incoherence. If AI capability is bolted onto siloed systems without integration into core data flows, intelligence remains superficial.

Here, vendor strategy must be anchored in workflow ambition. The question is not which model performs best on benchmarks. It is which vendor strategy supports decision pathway redesign. If your ambition is to compress credit decisioning to near-real time, your vendor architecture must support low-latency data pipelines, explainability, and integration into operational systems. If your ambition is strategic

scenario modeling, your architecture must support large-scale simulation and cross-domain data ingestion.

Research on platform strategy suggests that ecosystem positioning determines long-term leverage more than individual feature superiority (Iansiti and Lakhani). In the AI era, vendor choice is not simply procurement. It is structural alignment.

In Stage Three, vendor strategy shifts again. Capability velocity remains high, but your internal operating model begins to mature. Now the risk is dependency. If your authority structures and workflows rely heavily on proprietary tooling, switching costs become governance constraints. Optionality must be preserved deliberately. API abstraction layers, modular integration design, and contractual flexibility become strategic tools.

CIOs must now negotiate not only price and uptime guarantees, but roadmap influence, data portability, model transparency, and escalation responsiveness. Vendor relationships become partnerships in operating design rather than software supply arrangements.

Artificial intelligence compresses not only analytical cycles but vendor evaluation cycles. The old rhythm of annual RFPs and multi-year lock-in must coexist with rapid capability iteration. Your job is to prevent architectural chaos while preserving strategic flexibility.

Vendor strategy in the AI era is no longer about consolidation alone. It is about designing a capability portfolio that aligns with your maturity stage and your authority ambition.

## The CIO and the CEO

The CEO relationship may be the most consequential shift of all.

Historically, many CIOs operated as stewards of reliability. The mandate was clear: maintain uptime, modernize infrastructure, protect data, reduce cost. Strategic influence varied by organization, but in many firms the CIO was not viewed as an architect of competitive advantage.

Artificial intelligence changes that dynamic.

In Stage One, the CEO is typically curious and externally influenced. Headlines describe transformative potential. Competitors

announce initiatives. The CEO wants assurance that the company is not falling behind. The pressure is directional rather than operational. Your responsibility is to temper hype without dampening ambition. You must articulate where experimentation is occurring, what guardrails exist, and how the organization is learning.

In Stage Two, the tone shifts. The CEO expects movement. They want to know why productivity anecdotes have not translated into visible performance change. They may question whether the technology organization is moving fast enough. At this stage, your influence depends on reframing the narrative. The constraint is not deployment speed. It is structural redesign.

You must be explicit that artificial intelligence will not generate competitive advantage until authority structures evolve. This is an uncomfortable message. It implies cross-functional change. It implies executive alignment beyond IT. But without that clarity, AI becomes a recurring disappointment rather than a transformation lever.

Research on digital transformation repeatedly shows that executive alignment across functions is necessary for structural impact (Brynjolfsson and Hitt). The CIO cannot carry redesign alone. The CEO must sponsor authority recalibration across the enterprise.

In Stage Three, the relationship deepens further. Artificial intelligence becomes embedded in operating rhythm. Decision cycles compress. Scenario modeling accelerates. The CEO begins to rely on faster feedback loops. Strategy discussions become more dynamic.

At this stage, the CIO shifts from technology steward to operating model partner. Conversations with the CEO focus less on tool capability and more on authority design, capital velocity, and competitive asymmetry. You are no longer simply reporting on AI initiatives. You are shaping how the enterprise decides.

There is also a political dimension.

Artificial intelligence introduces visibility into knowledge work that was previously opaque. When analytical preparation becomes faster, inefficiencies in coordination become more visible. This can create tension across executive peers. The CIO must navigate these dynamics carefully. The goal is not to expose weakness but to redesign process.

The CEO will ultimately judge AI success not by the number of models deployed, but by whether the organization feels faster, clearer, and more deliberate.

If you cannot translate AI maturity into operating maturity, the CEO relationship remains tactical. If you can, it becomes strategic.

## **FAILURE MODES AT THE CIO LEVEL**

Several predictable patterns derail maturity.

One is mistaking vendor acquisition for transformation. Purchasing advanced systems without altering authority structures produces limited enterprise impact.

Another is pilot proliferation. Activity substitutes for redesign. Energy diffuses across disconnected initiatives.

Another is governance overreach. Lengthy policy frameworks that define tool usage without redefining authority create compliance without movement.

Another is protecting legacy approval layers in the name of risk without recalibrating thresholds in light of improved information quality.

Another is sequencing error. Research consistently demonstrates that technology adoption without complementary organizational change yields muted productivity gains (Brynjolfsson and Hitt; Stanford Digital Economy Lab). When authority redesign precedes legitimacy, resistance rises. When governance precedes behavioral visibility, experimentation disappears underground.

Most subtle is the failure to recalibrate managerial cadence. If decision forums remain update forums, artificial intelligence accelerates formatting rather than outcomes.

## **THE CORE SHIFT**

For decades, firms were organized around informational scarcity. Knowledge was difficult to gather and slow to transmit. Authority accumulated where information aggregated.

Artificial intelligence alters that scarcity structure.

Intelligence becomes abundant and near-instant. The bottleneck shifts from cognition to coordination.

In this environment, competitive advantage depends less on generating analysis and more on designing authority deliberately.

Stage One requires visibility.

Stage Two requires structural courage.

Stage Three requires systemic recalibration.

Across all stages, the CIO's responsibility is consistent.

Not to deploy intelligence.

But to redesign authority in light of it.

In an era of abundant cognition, disciplined authority becomes the scarce resource.

And no executive is better positioned to reshape it than the CIO.

# EPILOGUE

## A DAY IN THE LIFE OF A CIO IN 2030

IT IS 6:45 A.M., and before the first meeting of the day, the CIO opens the enterprise operations console.

Ten years earlier, this ritual would have meant scanning overnight incident summaries, checking whether any production systems had degraded, and reading through a cascade of emails summarizing what had already happened. The posture was reactive. The goal was to ensure nothing had broken.

In 2030, the console does not summarize yesterday. It reflects operating posture in real time.

Application latency is displayed as a stability band rather than a static metric. Cloud consumption variance is shown against a rolling twelve-week forecast that recalculates continuously as demand signals change. Cyber anomaly detection confidence scores update dynamically, calibrated against evolving behavioral baselines. Model drift indicators are visible across underwriting, pricing, and fraud detection systems.

Overnight, two microservices supporting customer authentication restarted automatically after anomaly thresholds were crossed. The remediation occurred within predefined authority parameters. The system logged the trigger, the confidence interval, and the corrective action. No human escalation was required.

The CIO reviews the audit trace not because there is a crisis, but because governance is continuous. The question is not “Did something fail?” It is “Are the thresholds still appropriate?”

At 8:00 a.m., the CIO joins the weekly executive operating forum.

The meeting used to be an exercise in presentation. Slides were distributed in advance. Each executive summarized performance in their domain. Questions focused on whether numbers reconciled and whether data was current.

Now, every participant sees the same live enterprise model. Revenue performance is simulated against multiple pricing and demand assumptions. Margin compression scenarios are recalculated automatically when commodity price inputs shift. Workforce allocation models update as hiring or attrition signals change.

The CEO opens with a focused question about regional pricing posture in a market experiencing competitive pressure. The AI-supported pricing engine has recommended a more assertive adjustment, supported by elasticity modeling and inventory exposure forecasts.

The CFO verifies capital buffer implications immediately. The risk officer confirms that no regulatory constraints are implicated by the adjustment. The CIO is asked whether infrastructure elasticity modeling supports the projected transaction surge.

Instead of promising to “get back with data,” the CIO rotates the simulation view. Peak load tolerance is displayed alongside historical variance. The infrastructure has sufficient headroom.

The group debates appetite, not analytics. Within minutes, the decision is made. Parameters are updated in the orchestration layer. The system implements the pricing change within predefined guardrails.

The meeting moves on.

At 9:30 a.m., the CIO meets privately with the Chief Information Security Officer.

A privileged access anomaly has been flagged by behavioral detection systems. The signal exceeds historical variance but remains below automatic account suspension thresholds. In prior years, this might have triggered immediate manual review and potentially unnecessary disruption.

Now, the conversation focuses on calibration.

They review false positive rates over the past quarter. They examine how similar anomalies resolved historically. They discuss whether the risk tolerance boundary should narrow slightly in light of recent external threat intelligence.

The adjustment is not reactive. It is structural. A threshold moves by a small percentage. The system logs the governance change, including who authorized it and why. Accountability is explicit.

The discussion ends not with a patch but with a recalibrated parameter.

At 11:00 a.m., the CIO convenes the monthly technology investment review.

In the 2020s, this session often revolved around PowerPoint decks and static ROI calculations. Teams sought capital approval based on projected cost savings or efficiency improvements. Financial modeling required weeks of preparation and was quickly outdated by shifting assumptions.

In 2030, every material initiative exists as a dynamic object within the enterprise financial model. Projected deployment cadence, technical debt reduction, operational resilience impact, and revenue enablement potential are continuously simulated.

A modernization initiative requests acceleration. The CIO does not ask whether the numbers are correct. He asks what structural constraint is being removed by acceleration.

The architecture team explains that compressing deployment frequency in a legacy system unlocks faster integration with machine-supported decision workflows in customer onboarding. The finance model updates in real time to reflect cash flow implications and margin sensitivity.

The CFO and CIO review the trade-offs together. The question is not whether the project is interesting. It is whether accelerating it changes authority velocity elsewhere in the firm.

The decision is made. Capital allocation shifts. The enterprise model reflects the change immediately.

After lunch, the CIO meets with enterprise architecture leaders.

The discussion centers on API exposure and data latency across

core systems. Artificial intelligence agents now operate continuously across underwriting, logistics, and customer service workflows. Architectural delay has become visible as decision delay.

A map of the firm's decision pathways is displayed. Some workflows still rely on nightly batch data reconciliation. These represent pockets of structural drag. When authority attempts to compress in those domains, it encounters latency inherited from older system designs.

The roadmap discussion is not framed as modernization for modernization's sake. It is framed as removing friction from specific decision loops. The architecture function has evolved from infrastructure custodian to velocity enabler.

Later in the afternoon, the CIO reviews board preparation materials.

A decade earlier, board updates focused on project milestones and cybersecurity posture. Today, the board expects clarity about operating design.

The report includes metrics that did not exist in prior eras: average decision latency in underwriting; escalation frequency in automated compliance monitoring; override rates in machine-supported workflows; model drift stability over time; vendor concentration risk exposure.

The board no longer asks whether the company "has AI." That question is obsolete. Instead, directors want to know whether authority boundaries are clear, whether escalation logic is enforced, and whether the firm can pivot without architectural lock-in.

The CIO reviews the narrative carefully. It emphasizes governance maturity, structural coherence, and risk calibration. The tone is neither promotional nor defensive. It reflects an operating model that has internalized intelligence as infrastructure.

In the late afternoon, a vendor strategy session begins.

Cloud providers continue to bundle reasoning services deeply into their ecosystems. Pricing models fluctuate. Model performance curves shift rapidly.

The CIO's questions focus on interoperability and exit cost. How

portable are inference workloads? How modular is the architecture? What contractual clauses protect against dependency drift?

Vendor management has evolved from feature comparison to structural optionality.

Before the day ends, a live alert surfaces from the customer service domain.

An AI-supported resolution engine has autonomously issued compensation credits for a surge in service delays caused by a regional weather disruption. The total credit exposure remains within preauthorized financial limits.

The CIO reviews the override and anomaly log. No drift. No threshold breach. The system acted within the boundaries leadership defined.

Ten years earlier, hundreds of manual reviews would have accumulated. Supervisors would have been required to approve individual credits. Customer dissatisfaction would have lingered.

Now, routine action occurs automatically. Leadership intervenes only when patterns exceed tolerance.

As the CIO closes the console in the evening, the difference between past and present is not that artificial intelligence dominates the enterprise. It is that informational delay no longer defines it.

Meetings are structured around commitment rather than updates.

Authority is encoded explicitly rather than implied informally.

Systems operate continuously within defined boundaries.

Human leaders focus on judgment, calibration, and responsibility.

The CIO's role has not diminished. It has intensified.



# AFTERWORD

This book was published as an ebook on purpose.

Not because print is obsolete, and not because the ideas here are temporary. It is because the *timing* matters. The competitive value of understanding a technological shift is highest while the shift is still unevenly understood.

Within a few years, much of what you read here will feel obvious. Leaders will talk naturally about decision flows, machine-assisted judgment, and redesigned workflows. Vendors will package these concepts into products. Consultants will summarize them into frameworks. Business schools will teach them as standard operating practice.

At that point, these ideas will still be correct.

**They simply will no longer be an advantage.**

Right now, they are.

Artificial intelligence will not transform organizations because a model becomes slightly better, or because a new interface appears. It will matter because some leaders change how their organizations operate before others do. Early clarity creates disproportionate results. Late clarity creates catch-up programs.

This is why the goal of this book was never to help you “use AI.”

It was to help you *organize around it*.

Many conversations about AI drift quickly toward extremes. Some treat it as magic. Others treat it as an existential threat. Both reactions are emotional, and emotion is rarely useful for management.

AI is neither a miracle nor a catastrophe inside an organization.

It is a capability.

Its value appears only when it helps people make better decisions, faster decisions, or more consistent decisions. Without human judgment, accountability, and coordination, it does not create business outcomes. It produces output.

Do not let the cultural noise around AI distract you from the managerial reality. Organizations do not compete on who adopts a tool first. They compete on who learns to operate differently.

Begin there.