

# Flatten the AI J-Curve



# FLATTEN THE AI J-CURVE

—  
*Your Unfair Advantage in the Race  
to Enterprise Adoption*

David Luria  
Corso & Alexander, Publisher

# Flatten the AI J-Curve: Your Unfair Advantage in the Race to Enterprise Adoption

Copyright © 2026 by David Luria. All rights reserved.

No part of this publication may be reproduced, distributed, or transmitted in any form or by any means, including photocopying, recording, or other electronic or mechanical methods, without the prior written permission of the publisher, except in the case of brief quotations embodied in critical reviews and certain other noncommercial uses permitted by copyright law.

For permission requests, speaking inquiries, or bulk purchase information, contact the publisher at: david@flattenthej.com

ISBN [Paperback]: 978-x-xxx-xxxxx-x

ISBN [Hardcover]: 978-x-xxx-xxxxx-x

eISBN: 978-x-xxx-xxxxx-x

*The J-Curve, Kill Zone, Squad Model, Protocol Engineering, Janitor Agents, Toy Mode, Stepping Stone Projects, Frozen Middle, and Capacity Not Time are trademarks of Corso & Alexander, LLC. All rights reserved.*

Disclaimer: This book is intended for informational and educational purposes. The author and publisher make no representations or warranties with respect to the accuracy or completeness of the contents of this work. The strategies and frameworks described herein may not be suitable for every organization. Results referenced from third-party research and publicly available case studies are reported as documented; the author makes no guarantee of similar outcomes. This book does not constitute legal, financial, or professional consulting advice. Readers should consult qualified professionals before implementing any significant organizational or technology initiative.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1



Corso & Alexander, LLC

[flattenthej.com](http://flattenthej.com)

First Edition, 2026

*To Resilience and Strength.*

*To two dear friends  
who showed us all how it's done  
while they were here.*

**Bob Kaufman**  
**Anne Kahn**

# Contents

Introduction..... 00

## **PART I: THE REALITY CHECK**

Chapter 1: The Imagination Gap..... 00

Chapter 2: The J-Curve..... 00

Chapter 3: Stepping Stone Projects ..... 00

## **PART II: THE THREE-PHASE MIGRATION**

Chapter 4: Escaping Toy Mode..... 00

Chapter 5: The API Mandate ..... 00

Chapter 6: Process Destruction ..... 00

## **PART III: EXECUTION**

Chapter 7: Building the Foundation ..... 00

Chapter 8: Driving Adoption ..... 00

Chapter 9: Making the Business Case..... 00

## **APPENDICES**

Appendix A: Digital Asset Library ..... 00

Appendix B: Glossary of Terms ..... 00

About the Author ..... 00

Acknowledgments..... 00

# Personal Opening

This book is about me.

I need to say that up front, because every business book pretends otherwise. They all claim to be about you—your challenges, your transformation, your Monday morning. And this one will get there. But it starts here, with the reason I wrote it, because that reason is the only thing that makes the rest of it honest.

In 2022, I had stage 4B cancer. It went undiagnosed for eight months—eight months of symptoms that got explained away, deferred, rescheduled—until the night it put me in a hospital bed for an emergency splenectomy. They airlifted me from my home in Frederick, Maryland, to shock trauma in Baltimore. The surgery went sideways. I spent three weeks in a coma. The doctors told my wife I had a couple of days to live. My kids came to say goodbye.

And then one doctor's order saved my life.

She asked a simple question: Did we do a biopsy on the spleen? The answer was no. It was the height of COVID. Too many patients, not enough time, and he's not going to make it anyway. She said: Get it done. That single order is the reason I'm writing this sentence, because the biopsy revealed stage 4B blood cancer—in my bones, in my blood, everywhere. It had been there for months. Without that one question from one doctor who refused to let the situation dictate the standard of care, there would be no diagnosis. No treatment. No book.

With an extraordinary wife, extraordinary medicine, and the prayers of thousands of people I will never be able to properly thank, I made it. I've been in remission for four years.

\* \* \*

I tell you this not for sympathy. I tell you because of what happened next.

It was with no small sense of irony that I found myself, barely back on my feet, staring down what felt like another existential threat—this one to my career, my expertise, everything I'd spent twenty years building. The image of AI replacing me was like a scene from some bad pirate movie: you see the ship on the horizon, impossibly far away, and you tell yourself there's time. There's always time. Until there isn't.

But here's what coming face to face with mortality does: it brings things into sharp relief, fast. You develop a gift for seeing sense, rather than nonsense. An almost physical intolerance for wasted time, wasted motion, wasted potential. You see where wisdom stands out from noise. And perhaps more than anything, you look for opportunities to build resilience. Not the inspirational poster kind. The operational kind. The kind that means when the next crisis hits—and it will—you don't just survive it. You're stronger because of it.

This book is that opportunity.

Not because I have all the answers. But because I've been in the place where the answers matter more than the theory—where you don't have the luxury of waiting for perfect information, where the cost of inaction is not a missed quarter but something far worse. I wrote this book the way I've lived the last four years: with urgency, with clarity about what actually matters, and with zero patience for the kind of hedging and incrementalism that kills more organizations than any technology disruption ever has.

\* \* \*

*“The most important question isn’t ‘What will change?’ It’s ‘What will stay the same?’ Because the things that never change are what you build a strategy on.”*

*— Morgan Housel, Same as Ever*

\* \* \*

## **Introduction**

*People always want things faster, better, cheaper—and they always resist seeing how radically new tools can deliver that by destroying old models. The unchanging behavior is the strategy.*

Enterprise leaders ask endlessly what AI will change. Almost no one asks what won’t: our addiction to legacy processes, our fear of killing sacred cows, our preference for incrementalism over imagination. That unchanging behavior is why most AI initiatives die in the Kill Zone. This book exploits what never changes to help you win.

I’ve spent twenty years driving technology adoption inside financial services organizations where compliance makes every decision twice as hard, risk committees can kill a program with a single memo, and boards demand quarterly ROI on initiatives that need eighteen months to mature. I’ve managed enterprise transformations as a PMP-certified program manager, dissected failing processes as a Lean Six Sigma Black Belt, and coached reluctant organizations through painful transitions as a certified change management practitioner.

What I’ve learned across all of it is simple and brutal: people don’t change, even when the tools do. They resist, they hedge, they protect what they know. The technology shifts. The human behavior doesn’t.

This book uses that unchanging truth—drawn from Morgan Housel’s *Same as Ever*—to show you how to win the enterprise AI race while everyone else is still playing with toys.

Housel recounts a story about Jeff Bezos. Not long after he launched Amazon, Bezos kept getting the same question: “What’s the next big thing going to be?” He realized almost nobody asked the more important question: “What is not going to change?” People would still want things faster, better, cheaper in five years, ten years, twenty. So he built his entire business on that premise.

Each chapter in this book is linked to one of Housel’s observations about human behavior because it provides a guiding narrative throughout the book: people do not change. So let’s use that.

Here is the pattern I keep seeing. A company announces an AI strategy. They pick a vendor, spin up a pilot, and hand everyone a chatbot license. For the first two months, usage spikes. People summarize emails, rewrite Slack messages, generate bullet points for slide decks. Leadership sees the adoption metrics and declares victory. Then month three arrives.

The novelty wears off. The hard work of integrating AI into actual workflows begins, and suddenly everything slows down. People are doing their regular jobs and learning new tools simultaneously—a phenomenon I call double-keying—and productivity drops. By month four, the CFO is asking questions. By month five, the CIO is updating their resume. By month six, the board kills the program. Not because AI failed. Because nobody told them it was supposed to get worse before it got better.

That is the J-curve. And surviving it is what this book is about. The companies that push through the dip—NVIDIA’s research shows fifty-two percent of them—come out the other side with operational efficiency gains and capacity increases that their competitors, the ones who panicked and pulled the plug, will spend years trying to replicate. The Kill Zone is a filter. It separates the organizations that understand

how adoption actually works from the ones still chasing hype cycle promises. This book puts you on the right side of that filter.

A quick word about what this book is not. It is not a prompt engineering manual. It is not a breathless commentary on the dystopian impacts of artificial intelligence on society. It is not a deep dive into open-source versus closed-source model architecture, or a primer on how transformer technology works under the hood.

This is not a traditional change management book, either, because it focuses on things that do *not* change. You will not find much about training plans, readiness assessments, or impact analyses. I leave all that to others.

It is written specifically for leaders in regulated industries—financial services, healthcare, insurance, energy—where the stakes are compliance, not convenience, and where the margin for error is razor thin. Where every technology deployment needs a sign-off from legal, a blessing from risk, and a defensible audit trail before a single model touches production data. If your organization has ever killed a promising initiative because no one could answer the Chief Risk Officer’s questions in a committee meeting, this book is for you.

*This is a playbook for managers who simply want to win.*

It is a reference book you can turn to when you are staring at a failed pilot, a demoralized team, or a board that wants to see ROI by next quarter. It is built on observed principles of human behavior that do not change and applies those principles to what actually works in enterprise AI adoption right now. Throughout the book, a simple decision framework recurs: Value equals Utility divided by Cost.

Every AI initiative, every vendor pitch, every process redesign can be pressure-tested against this formula. If the utility is not clear enough to articulate in one sentence, or the cost is hidden behind vague implementation timelines, the initiative is not ready. This formula is your filter for cutting through the noise.

The book lives at the intersection of business management, operations, technology adoption, and change management. Because I have spent the last decade working in financial services, many of the examples skew toward that industry—banks, asset managers, insurers, the organizations where audit trails, risk committees, and Frozen Middle paralysis make every technology decision twice as hard. My book, my rules. But the frameworks apply to any regulated enterprise where the stakes are real.

A word about how this book was written. Writers using ghost writers is a pattern that does not change—from Roman senators to modern CEOs, the person with the ideas has relied on someone else to put words on the page. I used artificial intelligence—specifically, large language models—to help write this book.

I mention this because transparency matters, and because it would be hypocritical to write a book about enterprise AI adoption without using the technology myself. Every idea in it is mine. Original. Distilled from the work of others. Pattern-matched from one domain to be relevant in another. But the raw material comes from twenty years of driving technology adoption in organizations that would rather not change. The AI did the heavy lifting. I drove the truck.

*The most expensive software is the one that gets built twice.*

*The second most expensive is the one you use to summarize emails.*

That line captures the central tension of this book. Most companies are spending millions on AI tools that their employees use to rewrite Slack messages and summarize meeting notes. That is Toy Mode—the digital equivalent of buying a Formula 1 car to drive to the grocery store. A company paying thirty dollars per seat per month for five thousand employees burns \$1.8 million a year on what amounts to an expensive autocomplete. This book is about what happens when you stop playing with the toy and start running the machine.

Every chapter that follows is designed to make the J-curve survivable—and to turn it from a career-ending risk into a competitive weapon. The frameworks are sequenced so that what you learn in one chapter feeds directly into the next. Skip nothing. Even if your organization is already past the Kill Zone, the execution chapters will show you how to accelerate the climb.

Here is how the book is structured, and why the sequence matters.

Part I—The Reality Check—starts by confronting the Imagination Gap: why your board is almost certainly solving the wrong problem by optimizing existing processes instead of reimagining them entirely. Chapter 2 introduces the J-Curve itself—the predictable productivity decline that creates the Kill Zone where most executives panic and pull the plug on initiatives that were actually working.

Chapter 3 gives you the antidote: Stepping Stone Projects, fast wins you can deliver in thirty days that build political capital and momentum while the organization rides through the dip.

Part II—The Three-Phase Migration—maps the path from trivial AI usage to genuine transformation. Chapter 4 tackles Toy Mode and introduces the Squad Model—a department-level certification framework built on Dunbar’s research showing that peer accountability only works in groups of about fifteen.

Each squad progresses through three tiers: Gamma for sandbox learning, Beta for approved workflows, and Alpha for full agent deployment. Chapter 5 makes the case for the API Mandate—moving from humans chatting with bots to bots chatting with bots, which is where the real leverage lives. Chapter 6 is about Process Destruction: the discipline of asking whether a process should exist at all before you waste time automating it.

Part III—Execution—is where the rubber meets the road. Chapter 7 covers the unglamorous but essential foundation work: Janitor Agents that clean your data and Protocol Engineering that replaces the

chaos of individual prompt writing with enterprise-grade, repeatable structures. Chapter 8 addresses the hardest problem in any transformation—getting people to actually use the thing—through radical transparency and the unfreezing of the Frozen Middle, that layer of middle management where AI initiatives go to die. Chapter 9 closes with the business case: how to frame AI investment as capacity expansion rather than cost cutting, and how to prove it in a 90-Day Sprint.

Every chapter includes downloadable digital assets—calculators, templates, checklists, decision frameworks—so that you can take what you read and put it to work Monday morning.

The companion site, [flattenthej.com](http://flattenthej.com), houses four interactive tools that go beyond static PDFs: a project selection wheel that stress-tests which AI initiative to fund first, a process destruction wheel that challenges whether your workflows should exist at all, pre-built executive communication scripts for the three most dangerous moments in the J-curve, and the Protocol Vault—a library of ready-to-deploy AI protocols that demonstrate Protocol Engineering in action.

You fill in the form. The system runs the protocol. You never write a prompt. That is the future of enterprise AI, and you can use it before you finish the book.

If you're tired of million-dollar chat pilots that deliver nothing but summarized emails, keep reading. If your CIO just greenlit an AI strategy and nobody has mentioned what happens at Month 4, you need this book before the budget meeting. If your middle managers are quietly sabotaging adoption because they think AI is coming for their jobs, Chapter 8 was written for that exact conversation. The path out of Toy Mode is uncomfortable, predictable, and winnable. And it starts with understanding what never changes.

# **PART I: THE REALITY CHECK**

*“Imagination is more important than knowledge.  
For knowledge is limited, whereas imagination  
embraces the entire world.”*

*—Albert Einstein, “What Life Means to Einstein”  
(interview, 1929)*

## Chapter 1

# The Imagination Gap

*Why Your Board Is Solving the Wrong Problem*

In *Same as Ever*, Morgan Housel describes a pattern that repeats across every major technological shift: people don't reject new technology because they can't understand it. They reject it because they can't imagine a world where it matters. By the time the imagination catches up, the winners have already built the new world.

Enterprise leaders today see AI clearly—they pilot chat tools, run document summarizers, approve budgets. What they can't yet see is the transformed operating model where legacy intermediaries vanish and systems handle end-to-end work. That failure of imagination keeps most firms trapped in Toy Mode while competitors quietly build the future.

The telephone? “Why would I need to talk to someone who isn't in the room?” The internet? “It's a fad for academics and hobbyists.” In each case, the technology itself wasn't the barrier. The barrier was imagination. People could see what the technology did. They couldn't see what it would *change*.

Here's what makes this pattern dangerous: by the time you can imagine the world the technology creates, someone else is already living in it. And they're your competitor. The companies that won during those transitions weren't the ones with the best technology. They were the ones who could imagine—and then execute—a different business model before their competitors could even picture it.

That's the gap that kills companies. Not a technology gap. An imagination gap. And it's happening again. Right now. With AI.

## A Thousand Surveys Telling the Same Story

Two of the largest surveys in financial services published in the last twelve months paint a picture that should keep every executive awake at night—not because the news is bad, but because the gap between ambition and execution has never been wider.

NVIDIA's *State of AI in Financial Services* report—surveying 839 respondents across banking, insurance, fintech, and capital markets—found that 65% of financial services organizations are now actively using AI, up from 45% just one year earlier.

The InvestOps 2026 survey of 200 global buy-side operations leaders found 70% actively deploying AI solutions. The money is flooding in: nearly 100% of NVIDIA's respondents plan to maintain or increase AI spending in 2026, with 44% expecting budgets to rise more than 10%.

That sounds like an industry on the move. It's not.

Because when you dig into what “deploying AI” actually means, you find this: 65% of buy-side firms expanded operational teams to handle compressed timelines like T+1 settlement rather than investing in automation. When regulatory change arrived—the single most predictable kind of change in financial services—two-thirds of firms threw people at the problem instead of building systems that could absorb the change without proportional cost increases.

Sixty-three percent lack unified data across front, middle, and back office. Nearly half cite manual data reconciliation as a top challenge. Sixty percent are implementing manual workarounds as temporary solutions. NVIDIA found that data-related challenges remain the single biggest barrier to AI adoption, cited by 40% of respondents. Only 17% have achieved full straight-through processing without manual intervention.

The pattern repeats across regulated sectors: healthcare providers add staff for new documentation rules instead of AI ingestion; insurers reconcile claims manually rather than deploying agentic exception handling; energy operators optimize reporting instead of reimagining asset management workflows.

Meanwhile, the firms that *have* pushed through the hard part are seeing results that make the gap alarming. NVIDIA found that 83% of

firms that persisted through early implementation challenges now report positive ROI on their AI investments. Eighty-nine percent say AI is helping both increase annual revenue and decrease annual costs. The top three impacts: operational efficiencies (52%), improved employee productivity (48%), and improved customer experience (37%).

The firms that pushed through are getting rich. The firms that didn't are getting left behind. And the distance between them grows every quarter. This is what the imagination gap looks like in hard data. You know where AI should go. You're stuck keeping it in the toy box.

## **The Industry Everyone Complained About**

Jamie Dimon had been complaining about proxy advisers for years. So had every other asset management CEO.

The firms—Glass Lewis and Institutional Shareholder Services—formed a de facto duopoly in corporate governance. They provided research, voting recommendations, and infrastructure to help investment firms cast thousands of shareholder votes each year. Large asset managers paid them millions in annual fees. Their voting recommendations carried enormous influence over shareholder decisions. Their business models created obvious conflicts of interest.

Everyone knew it. Everyone paid anyway. In spring 2025, Dimon went public at an industry gathering: proxy advisers are “incompetent” and “should be gone and dead, done with.” Strong words. But words without action are just complaints. Every CEO in that room had said similar things in private. They all kept writing checks to Glass Lewis and ISS.

The infrastructure was too embedded. The regulatory requirements too complex. The risk of switching too high.

It was easier to complain than to change.

Then in January 2026, JPMorgan did something the industry thought was impossible.

## **“We Don't Need Them At All”**

JPMorgan Asset Management—\$7 trillion in client assets, one of the world’s largest investment firms—cut all ties with proxy-advisory firms. Effective immediately. Not “reduced reliance.” Not “supplemented with internal research.” Complete elimination.

The unit built an internal AI platform it called Proxy IQ. The system analyzes data from more than 3,000 annual company meetings and provides voting recommendations directly to portfolio managers. No external advisers. No intermediary layer. No checks to Glass Lewis and ISS.

According to the memo announcing the move, JPMorgan believes it’s the first large investment firm to entirely stop using external proxy advisers.

Here’s what makes this interesting: the proxy advisory industry didn’t lack technology. Glass Lewis and ISS are sophisticated firms with their own data analytics and research capabilities. They could have built Proxy IQ themselves. They had the expertise, the data, the client relationships. But they couldn’t imagine eliminating their own business model.

JPMorgan could. Because they weren’t trying to improve proxy voting.

They were asking a different question entirely: “What if the advisers don’t need to exist at all?” That’s not a technology question. That’s an imagination question.

## **Making the Horse Run Faster**

Most firms in JPMorgan’s position would have asked: “How do we make proxy voting more efficient? How do we get better recommendations faster? How do we reduce our reliance on external advisers?” Those are reasonable questions. They lead to reasonable projects. You’d get a pilot program, test AI-assisted research, supplement external advisers with internal analysis, gradually reduce costs while maintaining the safety net of the incumbents. You’d make the horse run faster.

JPMorgan asked: “What if we eliminate the horse entirely?”

That’s the imagination gap. And the data shows it’s endemic. Look at where firms are *actually* deploying AI in 2026: the top use case

across financial services is customer experience and engagement at 42%, followed by document processing at 40%. Both important. Neither transformational. The industry's top spending priority for 2026—cited by 41% of respondents, up from 26% the year before—isn't identifying new AI use cases. It's optimizing AI workflows that already exist.

Companies aren't building new capabilities. They're polishing old ones.

### **“It Expands Who Can Build”**

The imagination gap isn't limited to financial services. Spotify's engineering organization achieved a 90% reduction in time spent on complex code migrations using AI. Most enterprises would have stopped at that 90% efficiency number, put it in a board presentation, and called it a win. Spotify saw the bigger picture.

Spotify's Chief Architect and VP of Engineering, Niklas Gustavsson, described the real transformation: engineers can now describe and push code changes in plain language directly from Slack on their phones. That shift, as Gustavsson put it, doesn't just save time—it expands who can build and dramatically increases how quickly the organization can move.

Read that again. Most firms would headline “90% faster migrations.” Spotify asked: What if anyone in the company could ship code from their phone without writing a line? That's the leap from polishing horses to eliminating them. The first framing optimizes an existing capability.

The second creates an entirely new one that didn't exist before—a capability that changes who participates in building software, not just how fast the existing builders work. That's the difference between Little-to-Little thinking and End-to-End imagination.

Housel identified the typical progression people go through with world-changing technology: *I've never heard of it → I've heard of it but don't understand it → I understand it but don't see how it's useful → I see how it could be fun for rich people, but not me → I use it, but it's just a toy → It's becoming more useful → I use it all the time → I could not imagine life without it → Seriously, people lived without it? → It's too powerful and needs to be regulated.*

This progression described the telephone, the automobile, the internet, and the smartphone. It will describe AI. But most enterprises today are stuck at “I use it, but it’s just a toy.” They’ve deployed ChatGPT to 5,000 employees. They’ve spent \$1.8 million on licenses (\$30/seat × 5,000 × 12 months). And 91% of those employees are using it to rewrite emails and summarize meetings. That’s not transformation. That’s expensive typing practice.

## **Little-to-Little vs. End-to-End**

Here’s the framework that explains why some companies win and others wonder what happened.

**Little-to-Little:** Individual employees doing individual tasks slightly faster. Sarah uses ChatGPT to draft her emails. Tom uses it to summarize meeting notes. The company spends \$30 per seat per month and gets marginally faster typing. This is where 90% of enterprise AI sits today.

**Middle-to-Middle:** Workflow integration. Teams start coordinating through AI. Some processes change. But humans are still in the loop for every decision, creating bottlenecks. You’ve upgraded from a toy to a tool, but you haven’t changed how work actually gets done.

**End-to-End:** Full process transformation. Systems talking to systems. AI handling the routine. Humans managing exceptions. This is where JPMorgan landed—and where you get transformational change, not incremental improvement.

If JPMorgan had stayed at Little-to-Little, they would have deployed AI to help analysts research proxy votes faster. If they’d moved to Middle-to-Middle, they would have built tools integrating with Glass Lewis data, automating some analysis while keeping external advisers in place. They went End-to-End. Proxy IQ doesn’t make the old system faster. It makes the old system irrelevant.

## **The Next Wave Is Already Here**

While most firms are still optimizing document processing and email workflows, the next wave is already arriving. NVIDIA found that 42% of financial services firms are now using or assessing agentic AI—

systems that don't just respond to prompts but autonomously reason, plan, and execute complex tasks. Twenty-one percent have already deployed them in production.

One in five financial services firms already has AI agents operating autonomously in their workflows. Not chatbots. Not summarizers. Agents that make decisions, take actions, and learn from outcomes. The firms deploying agents aren't asking "How can AI help my employees work faster?" They're asking "What work can AI do without my employees being involved at all?"

That's the leap from Little-to-Little to End-to-End. And it's happening right now, in your industry, at companies you compete with. If you're still debating whether to expand your ChatGPT pilot, you're not just behind. You're solving a problem that no longer exists.

## **The Competitor You Don't See Coming**

The competitor you should worry about isn't the one with more capital or more people. It's the one without your constraints.

Imagine a consultant walked into your operations today—someone who'd never seen your business before. They wouldn't ask "How can we make this faster?" They'd ask "Why does this process exist? Why does this require 14 people? Why does this intermediary exist at all?" The outsider wouldn't do your processes faster. They'd skip your processes entirely.

They'd offer something you can't match—not because they have more people or more capital, but because they're not carrying 15 years of accumulated workarounds and legacy decisions.

That competitor exists. Right now. In your industry. Sometimes it's a fintech startup with a team of twelve, building in six months what your IT department says would take three years. Sometimes it's an existing competitor who got serious about transformation 18 months ago while you were still forming steering committees. The most dangerous version is the company from an adjacent industry that realizes they can serve your customers better because they don't have your constraints.

Glass Lewis and ISS had everything—the client relationships, the data, the expertise. They could have built Proxy IQ themselves. But

they couldn't imagine obsoleting their own business model. JPMorgan could.

Because they asked a different question: "What if we don't need them at all?"

Now ask yourself: What's your version of that question? What intermediary in your business exists because "that's how it's always been done"? What vendor relationship continues because nobody's questioned whether it needs to exist? If you can't answer that question, someone else will.

## **Closing the Gap**

If you're reading this book, you're probably stuck somewhere between "I use it, but it's just a toy" and "It's becoming more useful for me."

That's where most companies are. You've deployed AI. People are using it. You're getting some value.

But here's the problem: you're improving at 10% while your competitors are improving at 40%. The gap widens every quarter. Your competitor—the one who jumped to End-to-End transformation 18 months ago—isn't comparing themselves to their past performance. They're comparing themselves to what's now possible. And they're building capabilities you can't match by optimizing existing processes.

The Imagination Gap isn't something you close in a single strategic planning session. It's a muscle you build by repeatedly asking the hard question: "What would we do if we weren't carrying all this legacy?" And then actually doing it—even when it's uncomfortable.

The rest of this book is about how to make that transition—and how to survive the inevitable productivity dip that happens along the way.

Because here's what nobody tells you: the path from Little-to-Little to End-to-End goes through a valley. Productivity drops before it rises. Costs increase before they decrease. Your board will panic. Your middle managers will resist. Your best employees will threaten to quit.

This isn't a bug. It's a feature. The productivity dip is how you know transformation is actually happening. If everything stays

smooth and comfortable, you're not changing anything important. You're just making the horse run faster.

That valley—the place where productivity drops and boards panic—has a name. I call it the J-curve. And surviving it is what separates the winners from the 68% of initiatives that die in the Kill Zone. Turn the page.

\* \* \*

## **Closing the Imagination Gap**

Five questions for your next executive meeting:

1. If a new competitor entered your market tomorrow with unlimited AI capability, what would they do that you currently can't?
2. What percentage of your current processes exist because of historical constraints that AI removes?
3. Where are you on the Little-to-Little → Middle-to-Middle → End-to-End spectrum? Where are your competitors?
4. What would it cost to be wrong about AI's impact on your industry? What would it cost to be right but too slow?
5. Who in your organization has the authority to kill a process entirely—not automate it, eliminate it?

The Start Here tool at [flattenthej.com](http://flattenthej.com) takes these diagnostic questions further. Two card types—Utility and Cost—test both business impact and implementation complexity through an interactive exercise designed for leadership team meetings. Spin the wheel, discuss the card with your team, repeat. Forty-five minutes of honest conversation about where your organization actually stands.

\* \* \*

## Chapter 1 Takeaways

**The Imagination Gap is the real barrier.** Companies don't fail at AI because of technology limitations. They fail because they can't imagine what AI makes possible beyond their current capabilities. Glass Lewis and ISS could have built Proxy IQ. They couldn't imagine obsoleting their own business model.

**Little-to-Little → Middle-to-Middle → End-to-End.** Individual tasks lead to workflow integration, which leads to process transformation. Most companies are stuck at Little-to-Little. Winners are building End-to-End. The gap widens every quarter.

**The data is overwhelming.** Eighty-three percent of firms that pushed through early challenges report positive ROI. Eighty-nine percent report both revenue increases and cost decreases. Meanwhile, 65% of firms threw headcount at T+1 instead of automating and 60% are implementing manual workarounds.

**The agentic wave is already here.** Forty-two percent of financial services firms are using or assessing AI agents. Twenty-one percent have deployed them in production. If you're still debating chatbot pilots, you're solving yesterday's problem.

**The diagnostic question.** Not "How can AI make us efficient?" but "What would a competitor unconstrained by our legacy do differently?"

*Next: Chapter 2—The J-Curve: Why You Will Get Worse Before You Get Better*

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for technology adoption patterns, Housel's progression framework, and

principles of unchanging human behavior in response to innovation.

**Pitcher, Jack.** “JPMorgan Cuts All Ties With Proxy Advisers in Industry First.” The Wall Street Journal, January 7, 2026. JPMorgan Asset Management’s \$7 trillion in client assets; elimination of external proxy advisers Glass Lewis and Institutional Shareholder Services; deployment of internal AI platform Proxy IQ analyzing 3,000+ annual company meetings; Jamie Dimon’s spring 2025 statement that proxy advisers are “incompetent” and “should be gone and dead, done with”; December 2025 Trump administration executive order calling for regulatory probe of proxy advisers; Glass Lewis announcement to discontinue “benchmark” voting recommendations starting 2027.

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents across financial services, fielded August–September 2025. Referenced for active AI usage (65%, up from 45%), positive ROI among adopters (83%), revenue and cost impact (89%), data challenges as top barrier (40%), spending priorities (41% optimizing existing workflows), budget increases (44% expecting >10% growth), agentic AI adoption (42% using or assessing, 21% deployed), document processing as leading use case (32%), customer experience and engagement (42%), and top AI impacts on business operations (operational efficiencies 52%, employee productivity 48%, customer experience 37%).

**Dell’Acqua, F., McFowland III, E., Mollick, E., 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, September 2023. Study of 758 BCG consultants found 40% higher quality results when AI used within its capabilities; 19 percentage points lower performance when used outside capabilities.

**McKinsey Global Institute.** “The Economic Potential of Generative AI: The Next Productivity Frontier.” June 2023. Estimates \$2.6–4.4 trillion annually in value creation across 63 use cases.

**McKinsey & Company.** “The State of AI: Global Survey 2025.” Finds 80%+ of organizations see no tangible EBIT impact from AI; only 21% have fundamentally redesigned workflows.

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside Operations Leaders.” InvestOps Report 2026. Survey of 200 global buyside institutions (\$10B to \$500B+ AUM) across asset management, pension funds, and insurance firms. Key findings: 70% actively deploying AI solutions; 65% expanded operational teams for T+1; 63% lack unified data across front, middle, and back office; 60% implementing manual workarounds; 49% cite manual data reconciliation as top challenge; 57% cite cultural resistance as most significant challenge; 55% prioritize competitive differentiation through innovation; only 17% achieved full straight-through processing.

**Gustavsson, Niklas.** Chief Architect and VP of Engineering, Spotify. Quoted in Anthropic Enterprise Briefing, “The Briefing: Enterprise Agents,” February 2026. Referenced for 90% reduction in engineering time on complex code migrations and expansion of who can build software through plain-language code changes via Slack.

*“Success is not final, failure is not fatal: it is the courage to continue that counts.”*

*—Commonly attributed to Winston Churchill*

## **Chapter 2**

# **The J-Curve**

*Why You Will Get Worse Before You Get Better*

In *Same as Ever*, Morgan Housel writes about the universal pattern of progress: things always get worse before they get better. The baby learning to walk falls repeatedly before taking confident steps. The company reorganizing its structure experiences chaos before clarity. The athlete rebuilding their swing performs terribly before breaking through to the next level.

This isn't a bug—it's how transformation works. The dip is the price of admission—no skipping, no shortcuts, just preparation or panic.

That's individual learning. Now scale it to 500 people trying to transform how they work while the business keeps running, the Board keeps watching quarterly numbers, and your competitors are making the same bet with different timelines.

Welcome to the J-curve.

## **The Pattern You Can't Avoid**

Picture Sarah, your star claims processor. Month 1, she's thrilled with the AI demo—it nailed the three test cases in the boardroom, and she's imagining what her team could do. Month 4, she's double-keying everything, running the AI's output against her manual work, correcting mistakes the demo never made, and her throughput is down 30%. Your board wants answers. Sarah wants her old workflow back.

If this is not your first rodeo with enterprise transformation, you've seen it. The enthusiasm at kickoff. The confident projections. The budget approvals. Then, three months in, everything falls apart.

Productivity tanks. Users are frustrated. Executives panic. The CFO drafts termination memos.

The pattern appears when you plot productivity over time: sharp drop, painful trough, eventual recovery. It looks like the letter J.

Individual adoption curves are smooth S-curves trending upward. Enterprise adoption curves have a valley. You go down before you go up. Not because the technology is bad. Not because your team is incompetent. Because organizational learning works fundamentally differently than individual learning.

The companies that survive don't avoid the dip—that's impossible if you're doing real transformation. They plan for it. They make the dip less severe and the recovery faster. The difference between dropping 10% versus 30%, between recovering in 6 months versus 18.

Let me be direct: implementing genuine AI transformation guarantees a productivity drop. Not a plateau. Not a slight dip. A measurable, painful decline that shows up in quarterly numbers and triggers panic at the Board level.

This is the Kill Zone. Most companies don't survive it—not because the technology fails, but because executives lose their nerve when productivity is down and the Board is demanding answers.

The firms that make it through are the ones who planned for it.

## **The Evidence Is Everywhere**

In 2025, researchers from MIT, the University of Colorado Boulder, the U.S. Census Bureau, and Stanford published the most comprehensive study to date on AI's productivity effects. Analyzing data from two waves of U.S. Census Bureau surveys covering tens of thousands of manufacturing firms between 2017 and 2021, they found something that will make every executive who's lived through an enterprise technology rollout nod in grim recognition.

AI adoption reduced productivity by an average of 1.33 percentage points in the short term. That's the average. When the researchers corrected for selection bias—accounting for the fact that firms expecting higher returns are more likely to adopt early—the short-run negative impact was dramatically larger, reaching up to 60 percentage points before eventual recovery. The firms that pushed through

eventually outperformed their non-adopting peers, but only after four or more years of sustained investment.

Think of it as a three-act story. Act one: the short-term productivity drop hits, averaging 1.33 percentage points but reaching catastrophic levels for the most ambitious adopters. Act two: sustained investment in training, process redesign, and organizational restructuring—the intangible assets that don’t show up on anyone’s balance sheet. Act three: the firms that survived outperform everyone who sat on the sidelines. Four-plus years to get there.

This pattern isn’t new. Stanford economist Erik Brynjolfsson documented it across general purpose technologies going back a century. When factories installed electric motors in the early 1900s, they kept their old layouts designed for steam power. It took two to three decades—and complete factory redesigns—before productivity jumped. The same thing happened with computers in the 1980s.

The pattern is called the Productivity J-Curve. It’s driven by a simple reality: transformative technologies require massive investments in intangible assets—training, process redesign, organizational change—that don’t show up in productivity statistics until years later. Brynjolfsson and colleagues estimated that adjusting for these intangible investments yields true productivity levels nearly 16% higher than official statistics suggest. The work you’re doing in the trough is creating real value. It just doesn’t show up in the numbers yet.

## **Why Demos Lie**

In February 2026, researchers at Princeton published a study that may be the single most important paper for understanding why the J-curve exists at a technical level. Stephan Rabanser, Sayash Kapoor, and Arvind Narayanan—the team behind the widely cited “AI Agents that Matter” research—tested 14 AI models from OpenAI, Google, and Anthropic across 500 benchmark runs and found something that should reframe how every executive thinks about AI deployment.

Capability and reliability are fundamentally different dimensions. And nearly two years of rapid capability gains produced only modest reliability improvements.

Think aviation: planes need 99.999% reliability to fly passengers without a human at every decision point. The Princeton researchers argued that autonomous operation in high-stakes business contexts requires a similar standard—three to five “nines” of performance, meaning 99.9% to 99.999% accuracy. Current AI agents aren’t close. Consistency scores ranged from 30% to 75%—agents that solve a task on one attempt frequently fail on repeated attempts under identical conditions.

Most models couldn’t distinguish their correct answers from incorrect ones better than random chance. Bigger models weren’t uniformly more reliable—scaling up actually hurt consistency in some cases, as larger models with richer behavioral repertoires showed more run-to-run variability. All three major AI providers clustered together on reliability scores, confirming this is an industry-wide limitation, not a single-vendor problem.

This is the technical mechanism behind the J-curve. Your VP saw the AI nail a claims processing demo in the boardroom. She was watching a capability test—a single run under controlled conditions with a clean dataset and a carefully constructed prompt. When 500 claims processors started using that same AI on real data with real edge cases, messy inputs, and a thousand variations in how they phrased their requests, they were running a reliability test. Different test, different results. That’s Month 4.

Demos test capability. Production tests reliability. The gap between the two is the dip.

The Princeton research also explains why throwing more money at the problem doesn’t close the gap. Scaling up models improves some reliability dimensions but can hurt others. You can’t buy your way past the J-curve with a bigger model or a more expensive vendor. You have to build the organizational infrastructure—structured inputs, quality layers, certification frameworks—that compensates for reliability limitations the models themselves haven’t solved. That’s what the rest of this book is about.

## **Watching the J-Curve in Real Time**

Here's what makes 2026 different from past technology waves: we can watch the J-curve happening in real time, across multiple industries, and the data confirms what executives feared.

Two major surveys capture where the industry stands right now, and together they paint a picture of an entire sector caught mid-transformation.

NVIDIA's 2026 *State of AI in Financial Services* report surveyed 839 respondents across banking, insurance, fintech, and capital markets. The headline numbers sound triumphant: 65% of organizations are actively using AI, up from 45% in 2024. Nearly 90% are either deploying AI solutions or running pilots. Some 73% of leaders believe AI is important to their company's future.

But look beneath the surface and you see the J-curve's signature. Eighty-nine percent of respondents say AI is helping both increase revenue and reduce costs—yet 40% cite data-related issues as their top challenge, up from 33% the year before. Another 35% say they lack the AI experts and data scientists to execute. Implementation difficulties rank as a persistent barrier at 28%.

The money is flowing in—nearly 100% of respondents plan to maintain or increase AI spending in 2026, with 44% expecting budgets to rise more than 10%—but the operational reality is messy, expensive, and slower than anyone projected.

The InvestOps 2026 survey tells a complementary story from the buy-side. Among 200 global operations leaders managing anywhere from \$10 billion to \$500 billion in assets, 65% expanded operational teams to handle compressed settlement timelines rather than investing in automation. When T+1 settlement arrived, two-thirds of firms threw people at the problem instead of building systems that could absorb change without proportional cost increases.

That's the J-curve in action. You're adding headcount to keep pace because your operating model can't absorb changes without adding staff. Sixty-three percent lack unified data across front, middle, and back office. Nearly half—49%—cite manual data reconciliation and validation as a top challenge. Sixty percent are implementing manual workarounds as temporary solutions. Only 17% have achieved full straight-through processing without manual intervention.

Translation: most firms know they need to transform. They're spending money. But they're stuck in the trough, adding complexity and cost while waiting for productivity gains to arrive.

The NVIDIA data reinforces this tension. The top spending priority for 2026 isn't identifying new AI use cases—it's optimizing AI workflows and production cycles that already exist, cited by 41% of respondents, up from 26% the year before. Companies aren't scaling new capabilities. They're trying to make current ones actually work. That's what the trough looks like from the inside.

## **What the J-Curve Looks Like from the Inside**

The MIT/Census research tells you the J-curve is real. The industry surveys tell you firms are in it right now. But neither captures what it feels like inside the organization when the numbers go red.

Here's the pattern, drawn from the documented mechanisms the MIT researchers identified and the operational realities the InvestOps and NVIDIA surveys captured. I'm describing a composite, but if you've been through an enterprise AI deployment, you'll recognize every beat.

A 12-person team in a lending department deploys AI-assisted underwriting. The promise: faster approvals, better risk assessment, fewer errors. The AI analyzes financials, flags risks, generates preliminary recommendations. Underwriters focus on judgment calls and client relationships.

Four months in: disaster. Processing time is slower than before. Team output has declined significantly. Underwriters spend more time correcting the AI's mistakes than they ever spent doing the work manually. Error rates are up. Client complaints about delays are up. The CFO drafts a memo recommending termination of the program.

The Board meeting is brutal. "We spent hundreds of thousands and made things worse."

But the VP holds the line. "This is exactly what we predicted at the outset. Productivity will recover. We're on track."

The CFO doesn't believe it. Neither do two Board members. But they give it 90 more days.

By month 9, the team hits breakeven. Processing time returns to baseline. Underwriters are proficient. The AI's recommendations are accurate most of the time, and when wrong, the team spots errors immediately.

By month 14, transformation is complete. The same 12-person team is processing multiples of their original volume with zero headcount additions. The bottleneck shifts from underwriting to business development—they can't generate enough qualified leads to keep the team busy.

The CFO, who wanted to kill the program in Month 4, presents the results at the annual shareholder meeting as the centerpiece of the bank's digital transformation strategy.

This is the J-curve. Every element of it—the initial decline, the executive panic, the eventual recovery exceeding baseline—matches what the MIT researchers documented at scale. The firms that pushed through eventually outperformed their non-adopting peers. But “eventually” means four-plus years of sustained investment. That's the part nobody wants to hear.

## **The Mechanism: Why Productivity Always Drops**

The J-curve isn't a failure of technology. It's a feature of organizational learning. Four forces drive the decline, and understanding them is the difference between planning for the trough and being blindsided by it.

### **Double-Keying.**

During training, staff must do the work manually AND supervise the AI simultaneously. An underwriter who used to spend 3 hours analyzing a loan application now spends 2 hours analyzing it manually plus 2 hours reviewing what the AI produced and correcting its errors. Workload doesn't decrease—it increases by roughly 1.5x. This is the mechanism that drove Sarah's frustration in Month 4. She isn't failing. She's doing two jobs at once.

This isn't optional. You can't trust AI output until you've verified it against manual work hundreds of times. The AI is learning what “good” looks like by watching humans. But humans are doing double

the work while learning happens. The productivity drop isn't a bug. It's the cost of training.

The InvestOps data confirms this pattern. Sixty-five percent of firms responded to accelerating market requirements by expanding operational teams rather than accelerating automation. They're adding people because the AI isn't ready to work unsupervised—but they can't afford to stop manual work either. You're caught in the middle. Double-keying. Paying for both.

### **The Friction of Unlearning.**

The MIT researchers found something counterintuitive: older, more established firms experienced significantly larger short-term productivity losses than younger firms. The reason isn't that older firms have worse technology. It's that they have more accumulated knowledge, process design, and organizational structure built around the old way of doing things.

Your best underwriter has 20 years of experience. She knows which financial ratios matter for which industries. She knows which documentation is critical and which is noise. She can spot a bad loan in 30 seconds by pattern-matching against thousands of previous applications.

Now you're asking her to supervise an AI that doesn't know any of that yet. Worse, you're asking her to trust a system that occasionally contradicts her hard-won judgment. She doesn't resist because she's afraid of change. She resists because her expertise is being devalued—and she's not wrong. The AI's judgment is worse than hers right now. It won't always be. But right now, it is.

The friction of unlearning is highest among your most valuable people. The ones you can least afford to alienate are the ones most resistant to change. Not because they're luddites—because they're experts.

### **Process Breakage.**

Every enterprise has decades of accumulated workarounds. The spreadsheet that reconciles data between two systems that were never designed to talk to each other. The email chain that serves as an approval workflow because the official system doesn't handle

exceptions. The PDF that someone prints, signs, scans, and emails back because nobody automated the signature process.

None of that is documented. It's just "how things work."

When you deploy AI, those workarounds break. The AI can't access the PDF in someone's email. It doesn't know about "what we did last time." It needs structured inputs to produce structured outputs. Suddenly you're debugging 15 years of accumulated workarounds. Every edge case that someone handled manually now requires a decision: Do we build this into the AI? Do we standardize it? Do we eliminate it?

These are good questions to answer. But answering them takes time. While you're answering them, productivity drops.

InvestOps found that 49% of firms cite manual data reconciliation as a top operational challenge and 60% are implementing manual workarounds as temporary solutions. The NVIDIA survey corroborates: data-related issues—privacy, sovereignty, data scattered across disparate locations—are the number one challenge, cited by 40% of respondents. Firms are discovering their processes are worse than they thought. The workarounds are multiplying faster than the automation.

### **Measurement Chaos.**

Your KPIs were designed for the old world. "Loans processed per day" made sense when underwriters did everything manually. Now what are you measuring? Loans where the AI generated a preliminary assessment? Loans where underwriters accepted the AI recommendation without modification? Loans where the AI was overruled? Loans processed faster than baseline? Loans processed with fewer errors?

All of these matter. None of them existed in your old system. Your dashboard shows "loans processed per day" declining, but it doesn't show "underwriters reviewing AI output and catching errors before they reach clients."

You're measuring the wrong things. The numbers look bad because you're tracking yesterday's metrics during tomorrow's transformation.

Only 18% of buy-side firms have adopted predictive analytics with decision support. Only 20% operate on a single platform enabling immediate decision-making. Most firms are flying blind through the J-curve because their measurement systems weren't built for this transition.

## **The Kill Zone: Months 3–6**

Here's when most CIOs get fired. Budget spent, productivity down, Board panicking. The initiative looks like failure precisely when transformation is beginning.

Month 3: Users are certified but not proficient. They know how to use the tool but don't trust it yet. They're still double-checking everything, so workload is at its peak. The AI is making mistakes on edge cases nobody anticipated. Support tickets are flooding in.

Month 4: The CFO pulls usage reports. Only 40% of certified users are actively using the system. The other 60% have reverted to old workflows because "it's faster." Productivity is down from baseline. Error rates are up. Client complaints about delays are up.

Month 5: Board meeting. The CEO asks pointed questions. "When will we see results?" The CIO points to the original timeline showing breakeven at Month 9. The CFO counters with current trajectory showing breakeven at Month 14. Two Board members suggest pulling the plug.

Month 6: Make-or-break moment. Either the executive team holds the line and gives it 90 more days, or they kill the program and write off the investment.

This is where most AI programs die. Not because they were doomed to fail—because leadership didn't plan for the dip and panicked when it arrived.

The NVIDIA data shows why the stakes keep rising. Forty-two percent of financial services firms are now using or assessing agentic AI—systems that autonomously reason, plan, and execute complex tasks. Twenty-one percent have deployed them in production.

But the challenges with agents mirror the J-curve pattern exactly: 34% cite performance and reliability issues, 33% cite lack of internal skills to manage them, and 30% cite data-related problems. The firms

entering the agentic era without having survived the J-curve of basic AI adoption are going to hit a wall twice as hard.

## **The Survival Guide**

If you're going to make it through the J-curve, you need four things in place before the trough arrives.

### **The Pre-Mortem: Set Expectations Before the Dip.**

Before you deploy, tell the Board explicitly: "Productivity will decline in Months 3–6. If it doesn't decline, we're not actually changing anything meaningful."

Frame it as the cost of transformation. "We're going to have a bad quarter. That's the price of building capability that will compound for years."

Get this on record in Month 0. When Month 4 arrives and the CFO wants to kill the program, you pull out this presentation and say: "As predicted." This is your insurance policy. Without it, when things get bad, leadership will interpret the dip as failure rather than progress.

### **The North Star Metric: Ignore "Time Saved" During the Trough.**

The wrong metric during the J-curve is "hours saved per employee." This will be negative. You're not saving time—you're investing time in training the AI and building proficiency.

The right metrics during the trough are leading indicators: Is the AI's output getting more accurate week over week? Are error rates declining? What percentage of users have completed Gamma certification? How many have graduated to Beta? What's the weekly active usage rate?

Don't measure ROI in the trough. You're not getting ROI yet. Measure learning. If learning happens, ROI will follow.

### **Protected Budget: Lock in 18 Months Before You Start.**

Six-month pilots die in the Kill Zone. The budget runs out in Month 6, right when you'd normally break even in Month 9.

Lock in 18 months of funding before you start. Tell the Board: “We will not know if this worked until Month 12 at the earliest. If we stop at Month 6, we will have spent the money and gotten nothing.”

This requires conviction. You’re asking for money with no early exit. Most leaders aren’t willing to make that ask. That’s why most AI programs fail. The MIT research is clear: companies that push through need four or more years before early adopters outperform non-adopters. The upside is real, but it’s not fast.

The good news from NVIDIA’s survey: 83% of financial services firms that have pushed through now report seeing a return on investment in their AI use cases. Eighty-nine percent say AI is helping both increase annual revenue and decrease annual costs. The other side of the J-curve is real. But only if you survive the trough.

### **Leading Indicators: Track Certification, Not Output.**

During the J-curve, track adoption behavior, not business outcomes. How many people completed Gamma certification within their squad? How many have graduated to Beta? What’s the weekly active usage rate? How many support tickets are filed? High ticket volume in Month 3 is good—it means people are using the system and finding edge cases.

Output metrics will be bad during the trough. Behavior metrics tell you if you’re on track to make it out.

## **The Communication Framework**

What you say, and when, determines whether you survive the J-curve.

### **Month 0: The Board Presentation.**

“We are implementing AI-assisted underwriting in the commercial loan department. Based on research from MIT and the U.S. Census Bureau studying tens of thousands of manufacturing firms, and NVIDIA’s survey of 839 financial services institutions showing that the top challenge remains data integration and skills gaps even among the most advanced adopters, we expect productivity to decline in the short term as the team learns the system and we refine workflows.

This is normal and expected. Productivity will return to baseline within 9–12 months and should significantly exceed our current capacity by 18–24 months. We have locked in an 18-month budget to ensure we reach the productivity gains on the other side of the learning curve.”

This is your insurance policy. When Month 4 arrives and productivity is down, you pull out this presentation and say: “As predicted.”

### **Month 4: When the Dip Arrives.**

“As predicted in our Month 0 presentation, productivity has declined in the short term. This is within the expected range based on similar enterprise AI implementations. 87% of squad members are now Gamma-certified. Error rates have declined 22% from Month 2. We are on track for breakeven in Q4.”

Notice what you’re not saying: “We’re sorry” or “We’re working to fix this” or “We didn’t expect this.” You’re not apologizing for the dip—you predicted it. You’re reporting progress on the metrics that matter, not defending failure.

### **Month 10: Recovery Begins.**

“Productivity has returned to baseline. Early wins appearing in the commercial lending team—we processed 30% more volume last month with the same headcount. Three squad members have achieved Beta certification and are building approved automated workflows. We’re expanding the pilot to the consumer lending team in Q1.”

Now you start talking about outcomes. Not before.

## **The Psychological Game**

The J-curve isn’t a technical problem. It’s a psychological problem. Executives panic when short-term metrics decline, even when long-term strategy is sound.

Your job isn’t to avoid the dip—it’s to survive it. Set expectations. Track the right metrics. Communicate relentlessly. Lock in the budget before the pressure arrives.

The firms that make it through don't just recover. They outperform with 30% capacity increases, as NVIDIA's survivors report. Fifty-two percent of firms that pushed through cite operational efficiency as one of the top ways AI improved their business, and 48% report boosted employee productivity—that second number more than doubled from 22% the year before. The gains accelerate once you're past the trough.

Surviving the dip also means asking the right questions mid-trough. Before you automate a broken process, ask the E-S-A questions: Can we eliminate this step entirely? Can we standardize what remains? Only then should you automate. Chapter 6 builds the full framework, but the instinct needs to start here—in the Kill Zone, when the temptation is to throw more technology at workflows that shouldn't exist in the first place.

## **The Gap Is Widening**

The InvestOps data shows the separation happening now. After two consecutive years prioritizing operational efficiency, 55% of firms have shifted their top priority to competitive differentiation through innovation. Portfolio management and trading innovation jumped from 32% to 56% in a single year.

Technology initiatives supporting this shift—advanced client self-service, reduction of manual processes through automation, portfolio management and trading innovation—each jumped approximately 20 percentage points from 2025.

The firms that survived their J-curve 18–24 months ago aren't just more efficient. They're building new capabilities. They're launching products competitors can't match. They're responding to regulatory changes without adding headcount.

The firms still stuck in the trough are falling behind at an accelerating rate.

This is why the J-curve matters so much. It's not just a temporary inconvenience. It's a permanent sorting mechanism. The firms that get through it pull away. The firms that don't get stuck in an expanding gap they may never close. Your competitor is launching AI-driven products while you're still reconciling data manually—that's not a gap. It's a chasm.

Chapter 3 shows you how to build momentum during the trough—fast wins that create political cover while the organization navigates the Kill Zone.

\* \* \*

## Digital Assets

The Don't Panic tool at [flattenthej.com](http://flattenthej.com) provides three pre-written executive communication scripts—one for each critical J-curve milestone. Enter your email to unlock all three as downloadable Word documents.

**Month 0: “The Insurance Policy.”** Lock in 18-month budget and get the productivity dip on record before deployment begins. Includes three slides: J-curve visual showing expected trough and recovery, leading indicators you'll track instead of ROI, and the budget ask with protected timeline. This presentation is your shield when Month 4 arrives.

**Month 4: “Hold the Line.”** Kill Zone survival script. No apology, no hedging. Report leading indicators—certification rates, error rate trends, usage patterns. Point back to the Month 0 presentation: “As predicted.” Includes a “What You Must Not Say” section covering the phrases that trigger Board panic.

**Month 10: “The Victory Lap.”** Story first—find a Sarah, the person whose work was transformed. Then the 30% capacity number. Then the expansion ask. The victory lap isn't just celebration. It's your pitch for the next deployment, baked into the success story of the first one.

*Access all three scripts at [flattenthej.com](http://flattenthej.com).*

\* \* \*

## Chapter 2 Takeaways

**The J-curve is real and inevitable.** MIT and Census Bureau research across tens of thousands of firms confirms: AI adoption reduces productivity by 1.33 percentage points on average in the short term, with some firms experiencing declines up to 60 percentage points. This isn't failure—it's the documented cost of organizational transformation.

**Demos test capability. Production tests reliability.** Princeton researchers found that 18 months of rapid capability gains produced only modest reliability improvements, with consistency scores ranging from 30% to 75%. The gap between what AI does in a demo and what it does in production is the J-curve's technical mechanism.

**The Kill Zone is Months 3–6.** This is when executive panic peaks, usage drops, and programs get killed. The firms that survive don't avoid the dip—they plan for it, communicate through it, and fund past it.

**Track learning, not output.** During the trough, the right metrics are squad certification completion rates, error rate trends, and usage patterns—not ROI. ROI comes later. If learning happens, ROI follows.

**Lock in 18 months before you start.** Six-month pilots die in the Kill Zone. The budget runs out right when breakeven would arrive. Protected funding is the single most important survival factor.

**The gap is permanent.** Firms that survive the J-curve are shifting from efficiency to innovation. Firms still in the trough are falling behind at an accelerating rate. The J-curve isn't a temporary inconvenience—it's a sorting mechanism.

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for universal pattern of progress: transformations always involve temporary setbacks before gains.

**Churchill, Winston S.** Epigraph attribution. The quote “Success is not final, failure is not fatal: it is the courage to continue that counts” is commonly attributed to Churchill, though its precise origin is disputed.

**McElheran, Kristina, Mu-Jeung Yang, Erik Brynjolfsson, and Zachary Kroff.** “Microfoundations of the Productivity J-Curve(s).” U.S. Census Bureau Center for Economic Studies Working Paper CES-WP-25-27, 2025. Primary source for J-curve evidence: AI adoption reduced productivity by 1.33 percentage points on average among tens of thousands of U.S. manufacturing firms (2017–2021 Census Bureau data); selection-bias-corrected estimates show short-run losses up to 60 percentage points; older/more established firms experienced significantly larger short-term losses; early adopters showed stronger growth over four-plus years.

**Brynjolfsson, Erik, Daniel Rock, and Chad Syverson.** “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 13, no. 1 (January 2021): 333–372. Referenced for historical J-curve pattern across general purpose technologies (electric motors, computers); estimated that adjusting for intangible investments yields true productivity levels nearly 16% higher than official measures.

**McElheran, Kristina, J. Frank Li, Erik Brynjolfsson, Zachary Kroff, Emin Dinlersoz, Lucia S. Foster, and Nikolas Zolas.** “AI Adoption in America: Who, What, and Where.” *Journal of Economics & Management Strategy* 33, no. 2 (2024): 375–415. Referenced for U.S. Census Bureau Annual Business Survey data on AI adoption patterns across 850,000 firms.

**Rabanser, Stephan, Sayash Kapoor, Peter Kirgis, Kangheng Liu, Saiteja Utpala, and Arvind Narayanan.** “Towards a Science of AI Agent Reliability.” arXiv:2602.16666, February 2026. Princeton researchers tested 14 models from OpenAI, Google, and Anthropic across 500 benchmark runs. Found that nearly two years of rapid capability progress produced only modest reliability gains. Decomposed reliability into four dimensions (consistency, robustness, predictability, safety) borrowing from nuclear and aviation safety engineering. Consistency scores ranged from 30% to 75%; most models couldn’t distinguish correct from incorrect predictions better than chance; scaling up improved some dimensions but hurt consistency. Aviation targets one catastrophic error per billion flight hours; researchers argue autonomous high-stakes operation requires 3–5 nines (99.9%–99.999% accuracy).

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents, fielded August–September 2025. Referenced for: 65% actively using AI; 40% cite data-related issues as top challenge; 35% lack AI experts; 28% cite implementation difficulties; 44% expect budgets to rise 10%+; 83% report positive ROI; 89% say AI increases revenue and decreases costs; 52% cite operational efficiency gains; 48% report boosted employee productivity; 42% using or assessing agentic AI; 21% deployed agents in production.

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside Operations Leaders.” InvestOps Report 2026. Referenced for: 65% expanded operational teams for T+1; 63% lack unified data; 49% cite manual data reconciliation as top challenge; 60% implementing manual workarounds; only 17% achieved full straight-through processing; 55% shifted top priority to competitive differentiation; portfolio management and trading innovation jumped from 32% to 56%.

**Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond.** “Generative AI at Work.” Quarterly Journal of Economics 140, no. 2 (May 2025): 889–942. Referenced for supporting evidence on AI productivity effects across skill levels.

*“The most effective way to do it, is to do it.”*

*—Amelia Earhart, advice to women in aviation, circa 1930s*

## **Chapter 3**

# **Stepping Stone Projects**

*How to Get a Win in 30 Days*

Morgan Housel describes a pattern that shows up everywhere, from ice ages to investing: things that compound rarely feel like they’re compounding when they’re happening. You plant a bamboo seed. For the first year, nothing visible happens. Year two, still nothing. Years three, four, five—no growth above ground. Then in year six, the bamboo shoots up 90 feet in six weeks. The growth didn’t start in year six. It started on day one. You just couldn’t see it because the plant was building roots.

This is how momentum works in organizations. The visible wins that leaders celebrate—the 400% capacity increase, the eliminated cost center, the new revenue stream—didn’t start when they became visible. They started with something small that nobody noticed at the time. People underestimate small things that compound and overestimate big things that don’t. They want the dramatic gesture, the bold initiative, the transformational program. They discount the quiet accumulation of marginal gains.

Earhart didn’t wait for someone to design a perfect flight plan across the Atlantic. She climbed in the cockpit and flew. Chapter 2 explained why productivity drops during real transformation—the J-curve is inevitable. This chapter is about what you do while you’re in the trough: you stack small wins until the Board stops panicking and the organization starts believing. You’re not trying to transform everything at once. You’re building bamboo roots.

Imagine your Board in Month 4, staring at a 15% productivity dip, ready to kill the budget. This chapter hands you the wins to change their minds.

## **The Problem with Grand Plans**

Most AI initiatives die because they try to boil the ocean. The 18-month roadmap. The cross-departmental steering committee. The enterprise architecture review. The change management consultants who bill \$400/hour to create PowerPoints nobody reads. By the time you've finished planning, the Board has lost patience and the budget is gone.

The data backs this up from every angle. Everyone is moving—the InvestOps 2026 survey of 200 global buy-side leaders found 70% actively deploying AI solutions, NVIDIA's State of AI in Financial Services report found 65% of organizations actively using AI (up from 45% in 2024), and nearly 90% are either deploying or running pilots. Yet most organizations struggle with the same pattern: ambitious plans that collapse under their own complexity. Nearly a third of firms in the InvestOps survey have noted that their AI deployments remain limited to a handful of use cases, hindered by fragmented data, legacy systems, talent gaps, and cultural inertia. Your competitor is already past this stage. You're still planning.

You know how this plays out. Months 1–3: formation of working groups, stakeholder interviews, requirements gathering. Months 4–6: vendor selection RFP, three finalists, arguments about build versus buy. Months 7–9: pilot scope definition, political arguments about whose budget and whose department goes first. Months 10–12: pilot finally launches, designed by committee to offend nobody and help nobody. Months 13–18: results are “inconclusive,” CIO presents revised plan, Board has moved on to the next crisis. You just spent 18 months and \$2M to prove nothing.

NVIDIA's data reveals where the money is actually going: the top spending priority for financial services firms in 2026 isn't identifying new AI use cases—it's optimizing AI workflows and production cycles that already exist, cited by 41% of respondents, up from 26% the year before. Companies aren't scaling bold new capabilities. They're trying to make current ones actually work.

The alternative: pick one thing, build it fast, show it works, repeat.

## **The Person Who Makes It Work**

Every organization that successfully navigates AI transformation has one person—often unrecognized, rarely thanked—who makes it

happen. Not the CIO. Not the Chief Digital Officer. Not the consultants. A middle manager who understands both the business process and the technology well enough to iterate fast. Someone who moves at the speed of problems, not approvals.

Call this person the Business AI Disruptor. The title doesn't matter. The role does.

This is not IT's job. IT builds infrastructure. IT maintains systems. IT keeps the lights on. That's critical work. But IT doesn't feel the daily pain of the accounts payable process that takes 47 manual steps. IT doesn't know which 12 of those steps exist only because someone retired in 2009 without documenting a workaround. The Business AI Disruptor lives in that process. Knows where it breaks. Knows which steps are theater and which actually matter. And knows the political landscape well enough to get things done without asking permission.

Look for the ops manager who's already hacking ChatGPT at 11 PM because they're tired of copying data manually. The best candidates are at the department-level business management layer: operations managers who own process flows, senior analysts who touch data all day, team leads who hear complaints from frontline staff. People close enough to feel the pain but senior enough to have discretionary time.

They understand the difference between process documentation (fiction) and actual workflow (truth). They don't need convincing about the technology; they need permission to scale it. They have credibility with both leadership and frontline staff. And they don't wait for permission—the best Business AI Disruptors have a track record of small acts of intelligent disobedience.

At one insurer, it was the compliance lead who quietly automated weekly regulatory reports—saving 22 hours per week before anyone in leadership noticed. That's the profile.

Give them protected time—even just 10 hours per week initially—connect them with IT for technical support, and point them at a Stepping Stone Project. You'll know within four weeks if you picked right.

## **What Makes a Good Stepping Stone**

Not all projects are equal. Pick one with four characteristics.

**One team, one process.** No cross-departmental dependencies. No waiting for other people to finish their part. The entire project sits within one manager's control. The moment you introduce dependencies, you introduce delay. Delay kills momentum. Momentum is the only thing that matters during the J-curve.

**Visible results.** Leadership needs to see it work. Frontline staff need to see it work. A win that nobody hears about doesn't build momentum. If the CFO asks "What have we accomplished?" at the next Board meeting, someone needs to answer: "The compliance team eliminated 22 hours per week of manual data entry. They're now processing 40% more cases with the same headcount." Specific. Measurable. Understandable by non-technical executives.

**Room to fail.** If this project fails, nobody gets fired. The stakes are high enough to matter but low enough to experiment. If people believe their jobs are on the line, they'll sandbag. They'll only propose initiatives they're 100% certain will succeed—which means incremental improvements, not transformational changes. You want intelligent risks. That requires downside protection.

**Repeatable.** Success in one team becomes a template for ten. Don't pick a one-off that solves a problem unique to a single person. Pick something where, if it works, you can copy it to similar teams in other departments. The goal isn't one win. The goal is one win that creates a pattern for many wins.

Where to start: find tasks that are high value (moves the needle if you fix them) and high pain (everyone hates doing them). That's your first Stepping Stone. High value doesn't mean "saves \$10M annually." It means "if this worked better, leaders would notice and staff would thank you." High pain means "everyone complains about it and we've built workarounds to avoid it."

## The Value Formula

Once you've identified a handful of candidates, you need a way to compare them. Not a spreadsheet with weighted scores—that's the enterprise planning disease this chapter is trying to cure. You need a heuristic that fits on a napkin and forces honest conversation.

Here it is: Value equals Utility divided by Cost.

That's it. It's a variation on internal rate of return, simplified for a conference room conversation. Utility is the numerator—how much does this project matter if it works? Cost is the denominator—how much organizational energy does it take to find out? A project with enormous utility but enormous cost (touches six systems, requires Legal sign-off, needs a vendor integration) might score the same as a project with modest utility and almost no cost (one team, existing tools, data already accessible). The formula doesn't tell you which project is "right." It tells you which projects give you the best return on organizational energy—and organizational energy during the Kill Zone is the scarcest resource you have.

In practice, this means a project like automating compliance checks—VP-noticeable impact, single-team scope, existing data—scores far better than overhauling the general ledger, which might deliver bigger impact but touches six systems and requires Legal, Finance, and two vendor integrations. In healthcare, it means ditching the cross-department RFP for a single-team claims processing bot that demos in Week 4. The formula biases toward action, which is exactly what you need during the Kill Zone.

This isn't meant to be calculated to two decimal places. It's meant to be argued over. The value of the formula is the conversation it forces, not the number it produces. When a department head says "automating the monthly reconciliation is high utility" and someone else says "but it touches the general ledger, the trading system, and two vendor feeds—that's high cost," you're having the right conversation. You're surfacing complexity before it surfaces itself at 2 AM on a Friday.

Think of it as a relative stack ranking. You're not trying to determine the absolute value of any single project. You're trying to determine which of your five or six candidates gives you the highest ratio of impact to effort. The one that sits in the upper-right quadrant—high utility, low cost—is your first Stepping Stone. The one in the lower-left—low utility, high cost—is the project you kill before it kills your momentum.

There are two categories of questions that drive this conversation. The first tests utility: Does this project produce results that a VP would mention at the next all-hands? If we fix this, does the team get

measurably faster, more accurate, or able to handle more volume? Can you explain the benefit in one sentence to someone who doesn't work in your department? Is there a line-of-business owner who would sponsor this personally? These aren't survey questions with Likert scales. They're conversation starters. If the room can't answer them clearly, the project isn't ready.

The second category tests cost: How many systems does this touch? Name them. Who outside this room needs to approve it? How many weeks before someone can start building? Is the data already accessible, or does someone need to negotiate access? What's the worst thing that happens if the prototype breaks? Again, not a scoring rubric. A provocation. The answers reveal whether a project is genuinely low-cost or whether the team is underestimating the organizational drag that will slow them down.

Here's what the formula protects against. Every leadership team has a bias: they gravitate toward high-visibility, high-complexity projects because those feel important. The CEO's pet initiative. The cross-functional digital transformation. The customer-facing chatbot that will require six months of legal review. These projects score high on utility but catastrophically high on cost. The ratio is terrible. They're J-curve accelerators—they make the dip deeper and longer. The Value Formula redirects attention to the boring, unglamorous projects that nobody puts on a conference slide but that actually work. The compliance workflow that takes 47 steps. The reconciliation that three people spend every Friday doing manually. The report that someone rebuilds in Excel every month because the last system migration broke the automated version. These are ugly problems with beautiful ratios: high utility to the people who live with them, low cost because they sit inside one team with existing data.

The formula also serves as a kill mechanism. If you can't make a credible case for either high utility or low cost, the project doesn't make the list. This is harder than it sounds. Organizations generate AI project ideas the way they generate meeting invites—constantly, reflexively, without quality control. The Value Formula gives you a polite way to say “not yet” without killing enthusiasm. “The utility is real, but the cost profile doesn't fit a Stepping Stone. Let's revisit this when we've built the infrastructure in Phase Two.”

One more thing. The formula changes as you mature. Early in the J-curve, you need low-cost, high-utility projects—fast wins that prove

the technology works and build political cover. Once you've survived the Kill Zone and have three or four wins under your belt, you can afford to take on higher-cost projects because you've built the organizational credibility and technical infrastructure to absorb them. The denominator tolerance increases as your track record grows. What would have been a dangerous bet in Month 2 becomes a reasonable investment in Month 10.

## **When Klarna Let AI Answer the Phone**

In February 2024, Klarna deployed an AI assistant to handle customer service inquiries. Not a company-wide AI transformation. Not an 18-month roadmap. One function, one process, fast deployment.

The results in the first month alone: the AI assistant handled 2.3 million conversations—two-thirds of all incoming customer service chats. Average resolution time dropped from 11 minutes to under 2 minutes. The capacity equivalent: roughly 700 full-time agents. Klarna reported approximately \$40 million in profit improvement tied to AI efficiencies in 2024, with a 40% reduction in cost per transaction since early 2023.

This is a textbook Stepping Stone. One team (customer service). One process (incoming chat inquiries). Visible results (resolution time, volume handled, cost reduction). Repeatable (the same pattern could expand to email support, phone routing, internal help desk). Klarna didn't try to transform every function simultaneously. They picked the highest-volume, highest-pain process, proved it worked, and let the results speak for themselves.

Run it through the Value Formula: utility was massive—customer service was Klarna's single highest-volume human interaction. Cost was manageable—one channel, one input type, existing data, no cross-departmental dependencies. The ratio was obvious. That's why it worked.

The pattern mirrors what NVIDIA found across financial services: document processing ranks as the top AI use case in capital markets at 35%, and a leading use case across banking and insurance. High-volume, repetitive, manual work—exactly the kind of process that makes a perfect Stepping Stone. One visible win created more momentum than any strategy deck could have.

## Why Small Wins Work

Technology changes. Human behavior doesn't.

**People copy what works.** One colleague's success story beats ten executive memos. When leadership announces a new initiative from the top down, people are skeptical—"Here we go again, another flavor-of-the-month." When a peer shows them something that actually works? They pay attention. "Wait, Sarah automated that terrible report we all hate? How'd she do it?" Peer adoption beats executive mandate every time.

**Fear beats greed.** "Your competitor is already doing this" motivates more than "You could save 20%." Humans are loss averse—we work harder to avoid losing something we have than to gain something we don't. After two consecutive years prioritizing operational efficiency, 55% of buy-side firms now prioritize competitive differentiation through innovation—the first time innovation has overtaken efficiency as the top strategic priority. Portfolio management and trading innovation jumped from 32% to 56% in a single year. The fear of falling behind is driving faster adoption than the promise of cost savings. Frame AI adoption as a competitive threat, not an efficiency opportunity.

**Stories beat spreadsheets.** The Board remembers "Sarah automated the compliance report" longer than "We achieved 3.2x ROI." When you present quarterly results, give the CFO the numbers. But lead with the story: "The compliance team was drowning in manual data entry. Sarah built an agent that eliminated 22 hours per week of grunt work. Now the same team is processing 40% more cases and compliance is asking to expand the pilot." That story sticks. The 3.2x ROI gets forgotten the moment the slide advances. Housel's observation: facts don't change minds, stories do. Numbers prove. Stories persuade.

## The 4-Week Sprint

If you can't show something working in four weeks, you picked the wrong project. Not "fully deployed and scaled to the enterprise." Not "integrated with 15 legacy systems." Just: does it work? Can you demonstrate value to a skeptical manager in a live demo?

**Week 1: Define the problem.** What’s the input? What’s the output? What does “good” look like? Be ruthlessly specific. Don’t define the problem as “improve customer service.” Define it as: “When a customer calls with a billing question, pull their account history, identify the three most likely issues, and surface the relevant policy section to the rep—all within 10 seconds.” If you can’t describe the success criteria in one sentence, you don’t understand the problem well enough to solve it.

**Week 2: Build the ugly version.** Duct tape and prayers. Does it work at all? This is where most people get stuck—they want clean code, scalable architecture, elegant UI. Stop. You need proof the concept works. Ship the ugliest version that demonstrates value. If it works, you can clean it up later. If it doesn’t, you just saved yourself three months of polishing something nobody wanted.

**Week 3: Test with real data.** Feed it 50–100 real examples. Where does it break? Where does it hallucinate? Where does it give answers that are technically correct but useless? NVIDIA’s survey found that 40% of financial services firms cite data-related issues as their top challenge—up from 33% the year before. Week 3 is where you discover if your data is good enough to support the use case. Better to find out now than after a six-month rollout.

**Week 4: Demo to leadership.** Show, don’t tell. Don’t send a memo. Don’t create a slide deck. Do a live demo—don’t slide-deck it. Show the 20-minute old way side by side with the 30-second AI way. Watch the CFO’s eyes light up. “Here’s how we do it today: 20 minutes per case, 30 cases per day. Here’s the AI version: 30 seconds, same accuracy.” Then ask: “Should we keep building this, or try something else?” If they say keep building, you’ve got a Stepping Stone. If they say try something else, you killed a bad idea in 4 weeks instead of 18 months.

## **Political Cover During the Kill Zone**

Chapter 2 explained the J-curve: productivity drops before it rises, and most companies panic during the trough. Stepping Stone Projects are the antidote to panic.

When the Board asks “What have we accomplished?” in Month 4—when productivity is down and the CFO is questioning the

investment—you need visible wins to point to. Not ROI projections. Not “we’re building capability.” Actual working solutions that real people are using to do real work.

The firms that push through are getting rewarded. NVIDIA’s survey found that 83% of financial services firms that persisted through the early challenges now report positive returns on their AI investments. Eighty-nine percent say AI is helping both increase revenue and decrease costs. Nearly 100% plan to maintain or increase AI spending in 2026, with 44% expecting budgets to rise more than 10%.

The other side of the J-curve is real. But only if you survive the trough. One win in Month 2. Another in Month 4. Another in Month 6. By Month 8, when productivity starts recovering, you’ve got a track record of delivered value that makes it impossible for skeptics to kill the program. You’re not trying to transform the enterprise in 4 weeks. You’re building evidence that transformation is possible. Evidence beats promises.

## **What’s Next**

You’ve got your Business AI Disruptor identified. You’ve picked your first Stepping Stone. You’ve shown leadership a working solution in 4 weeks.

Now you’ve got a new problem: people are using AI, but they’re using it badly. They’re treating it like a toy—writing emails, summarizing documents, doing trivial tasks that don’t move the business forward. You’re spending \$30 per seat per month on 5,000 employees, and 91% of them are in Toy Mode. That’s \$1.8 million a year on expensive email summarization. Chapter 4 breaks you out—how to escape Toy Mode through the Squad Model, turning casual users into certified operators who drive real business value.

## Digital Assets

The Start Here tool at [flattenthej.com](http://flattenthej.com) operationalizes the Value Formula and the Stepping Stone selection criteria from this chapter into an interactive exercise you can run with your leadership team in under an hour.

Two card types drive the conversation. Utility cards test the numerator—whether a project is visible enough and impactful enough to build momentum. Questions like “If we fix this, does a VP mention it at the next all-hands?” Cost cards test the denominator—implementation complexity, dependencies, and organizational drag. “How many systems does this touch? Name them.” Together they force the honest assessment that most project selection processes skip: is this Stepping Stone worth the organizational energy it will cost?

Spin the wheel, discuss the card with your team, repeat. No scoring rubric. No weighted matrix. Just 45 minutes of structured conversation designed to surface the project with the best value ratio—high utility, low cost—that gives you the best chance of a visible win in 30 days.

*Access the Start Here tool at [flattenthej.com](http://flattenthej.com).*

## Chapter 3 Takeaways

**Stack small wins before they compound.** Momentum builds underground before it breaks through. Don't chase the dramatic gesture—accumulate marginal gains. Klarna's AI assistant handled 2.3 million conversations in its first month because they picked one process and executed fast.

**Embed a Business AI Disruptor in every business unit.** Someone from inside the business function must own rapid workflow iteration. Not IT, not change management, not consultants. An embedded operator who moves at the speed of problems.

**Select Stepping Stones with four characteristics.** One team, one process. Visible results. Room to fail. Repeatable success that becomes a template for scale.

**Score projects on a napkin, not a spreadsheet.** Value equals Utility divided by Cost. The best Stepping Stone has the highest ratio of business impact to organizational energy. Early in the J-curve, bias hard toward low-cost projects—you can afford bigger bets after you've survived the Kill Zone.

**Prove concepts in four weeks or pick a different project.** Define (Week 1), build the ugly version (Week 2), test with real data (Week 3), demo and decide (Week 4). Live-demo the old way versus the AI way. If you can't prove value in 4 weeks, the project is wrong.

**Lead with stories, back with numbers.** Peer success stories beat executive mandates. Fear of competitive loss beats promise of efficiency gains. The Board remembers "Sarah automated the compliance report" longer than "We achieved 3.2x ROI."

**Build evidence to survive the trough.** Only 17% of firms achieved full straight-through processing. The other 83% are stuck in partial automation. Stepping Stone Projects build the track record and momentum to push past the Kill Zone into real transformation.

*Next: Chapter 4—Escaping Toy Mode: The Squad Model*

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for principles on compounding, human behavior patterns, and how small things that work beat big things that don't.

**Earhart, Amelia.** Attributed advice to women pursuing aviation careers, circa 1930s. Frequently cited in collections of Earhart's public statements and correspondence. Epigraph source.

**Klarna.** Press release, February 2024: AI assistant handled two-thirds of customer service chats in first month (2.3 million conversations), resolution time from 11 minutes to under 2 minutes, equivalent of approximately 700 full-time agents. Q1 2025 press update: approximately \$40 million profit improvement tied to AI efficiencies, 40% reduction in cost per transaction since Q1 2023.

**WBR Insights and SimCorp.** "The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside Operations Leaders." InvestOps Report 2026. Referenced for AI deployment rates (70%), vendor evaluation priorities (57% stability, 54% innovation), data quality concerns (62%), predictive analytics applications (66%), competitive differentiation as top strategic priority (55%), and straight-through processing rates (17%).

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents across financial services, fielded August–September 2025. Referenced for active AI usage (65%, up from 45%), positive ROI among adopters (83%), revenue and cost impact (89%), data challenges as top barrier (40%), spending priorities (41% optimizing existing workflows), budget increases (44% expecting >10% growth), and document processing as leading use case across sectors.

**Luria, David.** The Failure of I.T. Project Management: Why It's Broken and How to Fix It. Self-published, 2022. Referenced for establishing business ownership of technology adoption.

*“There is nothing so useless as doing efficiently that which should not be done at all.”*

*—Peter Drucker, The Effective Executive (1967)*

## **Chapter 4**

# **Escaping Toy Mode**

### *The Squad Model*

Morgan Housel writes that incentives are the most powerful force in the world. People don't do what makes sense—they do what gets rewarded. Put the wrong metric in front of the smartest team in your company, and they will optimize for that wrong metric with terrifying efficiency.

Drucker saw the same thing sixty years earlier. Doing the wrong thing efficiently is worse than doing nothing—because it creates the illusion of progress while consuming the resources you need for real work. That's exactly what's happening with enterprise AI adoption in 2026. Companies are handing out generative AI tools to thousands of employees, then wondering why nobody's generating business value.

NVIDIA's *State of AI in Financial Services* report—surveying 839 respondents across banking, insurance, fintech, and capital markets—found that 65% of financial services organizations are actively using AI, up from 45% in 2024. The InvestOps 2026 survey found 70% of global investment firms actively deploying AI solutions. Nearly 100% of NVIDIA's respondents plan to maintain or increase AI spending in 2026, with 44% expecting budgets to rise more than 10%. The money is flowing. But when you look at what those solutions are actually doing, the picture gets ugly fast. Only 18% of firms have adopted predictive analytics with decision support. Only 17% achieved full straight-through processing. The rest? They're in Toy Mode—using \$30-per-seat-per-month enterprise AI licenses to summarize emails and rewrite Slack messages.

## **The \$1.8 Million Email Problem**

Picture this: a global company with 5,000 employees deploys ChatGPT Enterprise at \$30 per seat per month. Total annual cost: \$1.8 million. Eleven months later, leadership asks the obvious question: “What are we getting for that money?”

IT runs a usage audit. The data comes back: 600 users haven’t logged in for 60 days. Another 3,200 are using it exclusively for email summarization and document formatting. The most common prompt across the entire organization: “Make this sound more professional.”

The COO’s response: “We’re paying \$1.8 million a year for Grammarly.”

And if you’re in a regulated industry, it’s worse than Grammarly. Every one of those polished emails is still being manually double-keyed downstream—exactly the overlap that creates the 15% productivity dip you cannot afford. You’re paying for AI and still paying for the manual process it was supposed to replace.

A CFA Institute survey of 1,210 global charter holders found that only 35% use AI tools in their daily workflow. Among those who do, the top three use cases are writing assistance (20%), research and information gathering (18%), and coding (18%). Writing assistance. That’s the killer app. Not building predictive models. Not automating reconciliation. Not designing new investment strategies. Writing emails.

Meanwhile, the InvestOps survey shows that 60% of firms are implementing manual workarounds, 63% lack unified data across front, middle, and back office, and 49% cite manual data reconciliation as their top operational challenge. These are the problems AI could actually solve. Instead, firms are using enterprise AI to polish prose.

This isn’t an AI problem. It’s an incentives problem. If you’re the one signing the \$1.8M renewal, ask yourself: what behavior did you actually incentivize? When you give people access to powerful tools without guidance, boundaries, or consequences, they’ll use those tools for the easiest, safest, most visible tasks. Email is easy. Data reconciliation is hard. You get what you incentivize. You’ve incentivized toy usage.

## **The Cognitive Offloading Tax**

Toy usage isn't neutral. It's actively harmful. A 2024 study of 666 participants in the United Kingdom found that increased AI usage is positively correlated with cognitive offloading (72%) and negatively correlated with critical thinking (-68%). The more people use AI for simple tasks, the worse they get at thinking.

You're paying \$1.8 million per year for 5,000 employees to use AI for trivial tasks. While they're polishing emails, they're simultaneously getting worse at the higher-order thinking you actually need—the thinking that drives innovation, solves complex problems, and creates competitive advantage. Every time someone uses AI to summarize an email they should have read themselves, they're outsourcing a cognitive task that keeps their brain sharp. Do that enough times, and critical thinking atrophies. You're not just wasting money. You're degrading your workforce's capabilities.

Cultural resistance—which 57% of firms in the InvestOps survey cite as their biggest challenge—isn't irrational. It's a rational response to poorly structured change. Your regulators will not accept “we gave everyone ChatGPT” as a control framework. Your compliance team already knows this. What message are you sending when AI adoption equals email summarization? You're confirming their suspicion that this is a toy. Stop deploying toys.

## **The Real Escape: From Chatbots to Agents**

Most companies are stuck on one side of a dividing line they don't even know exists. Generative AI—the ChatGPT interface your people are typing into—responds to prompts. It's passive. You ask a question, you get an answer. You copy-paste the output somewhere. It's a smarter search engine with a text box.

Agentic AI is a fundamentally different animal. Agents reason, plan, and execute complex tasks autonomously. They orchestrate workflows, make decisions, and chain together multiple steps without waiting for a human to type the next instruction. NVIDIA's survey found that 42% of financial services firms are now using or assessing agentic AI, with 21% already deployed in production. That's nearly half the industry evaluating autonomous AI systems.

Salesforce saw this transition firsthand when they embedded AI into Slack. The result: 97 minutes per week reclaimed per employee—

not by opening a separate AI application, but by weaving intelligence into the tools people already use. Users don't switch to a chatbot interface. The AI meets them inside the workflow. Rob Seaman, EVP and General Manager for Slack, frames it as the journey toward an "agentic enterprise"—one where AI doesn't wait to be asked but anticipates and acts. That's the distinction in practice. Gamma users open a separate chat window. Beta users have AI embedded in their daily tools. Alpha users deploy agents that run without being asked at all.

Eighty-nine percent of NVIDIA's respondents say AI is helping both increase revenue and reduce costs. But that ROI isn't coming from email summarization. It's coming from document processing (the top use case at 32%), customer experience optimization (30%), and risk management. The companies generating returns are the ones that moved past the chat interface. The progression from generative to agentic is the progression from toy to transformative.

## **Why Top-Down Certification Fails**

The instinctive response to Toy Mode is centralized control. Build a corporate AI certification program. Have IT design the curriculum. Roll it out org-wide. It sounds reasonable. It doesn't work.

First, it's generic. What "AI proficiency" means in compliance is completely different from what it means in portfolio management. A single corporate test cannot distinguish between the analyst who needs to automate regulatory reporting and the trader who needs to build real-time market models.

Second, it's IT-driven. When IT owns certification, AI adoption becomes a technology project instead of a business transformation. IT designs tests around tools and features. Business leaders need tests around outcomes and workflows.

Sound familiar? Here's the one that kills the whole approach: centralized certification has no social engine. You pass a test, you get a badge, and nobody you work with daily knows or cares. There's no peer pressure. No visibility. In a company of 5,000, being uncertified is invisible. You're one name in a spreadsheet nobody reads.

## **Fifteen People Who Know Your Name**

Robin Dunbar, the Oxford evolutionary psychologist, spent decades studying human social groups. His famous “Dunbar’s number”—150—gets all the press. But the real insight is in the layers underneath. Human relationships organize themselves into concentric circles of roughly 5, 15, 50, and 150 people, each about three times the size of the last, each with decreasing emotional intensity.

The 15-person layer is what Dunbar calls the “sympathy group.” These are the people whose well-being you actively track. The people whose opinions materially affect your behavior. People devote roughly 60% of their total social effort to this inner circle of 15. Military planners discovered these numbers independently through centuries of trial and error, eventually settling on unit sizes that mirror Dunbar’s layers because cooperation broke down at other sizes. Teams of 15 show up everywhere: rugby squads, jury panels, platoon sections. It’s not accident. It’s neurology.

At 15 people, every member can track every other member’s status. You know who’s contributing and who’s coasting. You can’t hide in a group of 15 the way you can hide in a department of 150. That’s the foundation of the Squad Model.

## **How the Squad Model Works**

Every department-sized group of roughly 15 people becomes a squad. Each squad defines three tiers of AI proficiency—Gamma, Beta, and Alpha—based on what those tiers mean for their specific roles and workflows. There are no mandatory training hours. No corporate classroom. No set curriculum. Individuals train however they choose—peer coaching, online courses, self-study, vendor certifications, or simply learning by doing. The squad doesn’t prescribe the path. It only validates the destination. And critically, every person’s tier ranking is visible within their squad.

The terminology is deliberate. In investment management, alpha is the holy grail—excess return above the benchmark. Gamma is the floor. Beta is the market. Alpha is where the value lives. Nobody in financial services needs that explained. The language carries its own incentive structure.

**Gamma** is the entry point. Access to AI tools in a sandboxed environment with safe data. You’re learning, exploring, working with

generative AI as a conversation partner. But you're not connected to production systems and you're not executing business workflows. Every new user starts at Gamma regardless of title or tenure.

**Beta** means you've demonstrated competence for your specific function. A Beta-tier compliance analyst can execute approved AI workflows connected to internal data. A Beta-tier portfolio manager can run predictive models with live market feeds. Beta users aren't chatting with AI—they're executing structured protocols that produce consistent, auditable outputs. What "Beta" means is defined by each squad because the required competencies are different for every function.

**Alpha** is transformation. Alpha-tier users deploy autonomous agents, redesign processes, and build new capabilities. They've moved from generative to agentic—demonstrating not just proficiency but the judgment to know when AI should act independently and when it shouldn't. Alpha certification requires executive sponsorship and cross-functional impact. You don't reach Alpha by being individually productive. You reach it by making your entire squad more productive.

Here's what this looks like in practice. Consider a compliance analyst at a mid-size health insurer. At Gamma, she's using AI to summarize regulatory updates—useful, but her team is still manually cross-referencing every update against the company's policy library. She advances to Beta by building and validating a structured protocol that automatically maps new CMS guidance to affected internal policies, flags gaps, and generates a preliminary compliance memo for her supervisor's review. The audit trail is built into the workflow. Her supervisor doesn't review every output—she reviews exceptions. That single Beta workflow eliminates 14 hours per week of manual cross-referencing across her squad, and every step is documented for the next regulatory exam. That's not email summarization. That's transformation you can show an examiner.

This tiered progression isn't just good organizational design—it's what the research demands. Princeton researchers (Rabanser, Kapoor, Narayanan, 2026) tested 14 AI models across 500 benchmark runs and found consistency scores ranging from 30% to 75%. Agents that solve a task correctly often fail when given the identical task again. Their recommendation: reliability thresholds before moving from sandbox to production, the way aviation requires certification

before service. That's the Squad Model in academic language. When your tools produce different outputs from the same inputs a third of the time, you cannot skip the certification step. The Princeton team also recommends building an incident-reporting culture around agent failures—exactly what squad-level visibility formalizes. Every squad member sees what works and what breaks. That shared awareness doesn't happen in a 5,000-person deployment. It happens in a group of 15.

## **The Social Engine**

In a squad of 15, your tier ranking is visible to every person you work with daily. Not to 5,000 anonymous colleagues in a corporate directory. To the 14 people who sit in your meetings, review your work, and share your deadlines. Seeing whether you're Gamma, Beta, or Alpha among your 14 closest peers creates organic social pressure to improve. Nobody wants to be the only Gamma in a room full of Betas. Nobody wants to be the last person on their team still in the sandbox while their peers are automating real workflows.

This isn't shaming—it's the same peer pressure that keeps fighter pilots and transplant surgeons sharp. In your 15-person squad, nobody hides. At 15 people, we're within the cognitive limit where every member's status is trackable by every other member. The group is small enough that accountability is personal but large enough that the pressure is meaningful.

The InvestOps data reinforces this. When 57% of firms cite cultural resistance as their biggest challenge, the solution isn't another org-wide email from the CEO about "embracing AI." The solution is making adoption visible at the unit level where people actually feel accountable. In a squad, your resistance isn't hidden in aggregate statistics. It's a Gamma badge next to your name while your colleagues are Beta and Alpha.

## **The Frozen Middle Question**

If the squad is department-sized, the middle manager runs it. And if that middle manager is Gamma-tier themselves—if they're the one resisting AI—why would you let them define proficiency for everyone else?

Good question. Wrong assumption. The middle manager doesn't define the certification criteria. They facilitate the process. The certification content—what Gamma, Beta, and Alpha actually mean for a given function—comes from department heads, domain thought leaders, and Business AI Disruptors. The middle manager's job is logistics: scheduling certification windows, tracking progress, ensuring accountability.

But here's what the Squad Model does to the Frozen Middle that centralized programs can't: it makes their resistance visible. In a top-down certification rollout, a resistant middle manager just slow-walks implementation. Nobody notices for months. In a squad of 15, if the manager is Gamma while their reports are advancing to Beta and Alpha, the social dynamics invert. The squad doesn't need the middle manager to be enthusiastic about AI. It needs them to not block progress. And when blocking means being publicly outranked by your own direct reports, blocking becomes very expensive. This is the through-line to Chapter 8's deeper treatment of the Frozen Middle.

## **What Structured Deployment Looks Like**

McKinsey didn't hand 45,000 employees a ChatGPT login and hope for the best. When they launched their internal AI platform Lilli in July 2023, they built structured access from day one. The platform was purpose-built on McKinsey's proprietary corpus of over 100,000 documents, with role-based access controls and zero-trust security. They started with 7,000 users, expanded to firmwide over three months, and invested in upskilling employees on prompt engineering and content validation throughout.

The results speak for themselves: over 75% of McKinsey's 45,000 employees now use Lilli monthly, averaging 17 sessions per week. The platform processes over 500,000 prompts every month. Users report 30% time savings on research and knowledge synthesis. McKinsey has since deployed over 12,000 AI agents to automate tasks from data analysis to proposal drafting—moving well beyond the chat interface into agentic territory.

McKinsey didn't hand out logins—they built squad-level visibility and role-based controls from day one. Their middle managers didn't become blockers; they became facilitators because the metrics were visible in every 15-person pod. Compare that to the \$1.8 million

Grammarly scenario. Same technology category. Radically different outcomes. The difference isn't the AI. It's the structure. McKinsey built purpose-specific tools connected to proprietary data, phased access based on demonstrated capability, and invested in moving users from basic prompting to structured workflows to agent deployment. They didn't use the language of Gamma, Beta, and Alpha—but they followed the same progression. Sandbox learning first, structured execution second, autonomous agents third.

The Squad Model formalizes this progression for any enterprise. It makes the structure explicit, ties it to peer accountability, and ensures that the \$30 per seat per month produces real business value instead of polished emails.

## **What Gets Measured Gets Managed**

The Squad Model changes what you measure and who sees the data. Before, the metric was Daily Active Users—the vanity metric. “We have 8,000 people using ChatGPT!” Great. What are they doing with it? Silence.

After, the metrics are squad-level and outcome-based. Tokens per Complex Task measures depth of engagement—if you're using 50 tokens to summarize an email, that's Gamma behavior; if you're using 5,000 tokens to build an optimization model, that's Beta or Alpha work. Protocol Execution Rate tracks how often users are running approved, structured workflows versus typing into a blank prompt box. These metrics are published at the squad level, not the org level. Each squad sees its own distribution: how many Gammas, Betas, and Alphas, and how does it compare to last quarter. Two layers of accountability—individual within the squad, squad across the organization.

## **Enforcement Is the Model**

The Squad Model has built-in enforcement that centralized programs lack: peer visibility. But it still requires organizational commitment.

**Automated access control:** After 30 days at Gamma, the system automatically locks accounts that haven't demonstrated Beta-level engagement. No exceptions. No manual overrides without VP

approval. Technology enforces the policy so managers don't have to make uncomfortable decisions. This is your insurance against Month-4 panic—when the Kill Zone hits and leadership starts questioning the investment, automated enforcement means adoption is progressing whether individual managers champion it or not.

**Quarterly squad reviews:** Every quarter, squad-level metrics go to leadership. Not individual names—squad distributions. “Squad 7 has 3 Alphas, 8 Betas, and 4 Gammas. Squad 12 has 0 Alphas, 2 Betas, and 13 Gammas.” Leadership identifies which squads are transforming and which are stagnating—without micromanaging individuals. When the J-curve trough arrives, these quarterly snapshots become your evidence that the system is working even when the aggregate numbers look painful.

**Executive modeling:** The C-suite goes through squad certification first. Publicly. If the CEO can't achieve Beta within their own functional squad, how do you expect anyone else to take it seriously? Executive participation signals that this isn't a middle-management initiative—it's a business priority.

The InvestOps data shows why enforcement matters. The same survey that found 70% of firms actively deploying AI also found that only 17% achieved full straight-through processing. That's a 53-percentage-point gap between deployment and results. Deployment without structure produces that gap. The Squad Model, with visible peer accountability and organizational enforcement, closes it.

## What's Next

You've got structure. You've got squads. Your Beta operators are executing approved workflows. Your Alpha users are building new capabilities. Now you've got a different problem: most of your AI work still involves humans typing into chat interfaces. That's fine for Gamma learning and some Beta work. But it's a dead end for transformation.

Chat-based AI is human-in-the-loop. For 80% of enterprise processes, you don't want human-in-the-loop. You want human-on-the-loop—supervising autonomous systems that run without constant intervention. NVIDIA's data confirms the direction: 42% are already moving toward agentic AI. The industry isn't waiting.

That's what Chapter 5 is about: moving from chat interfaces to API-driven automation. From humans asking AI questions to AI systems executing business processes end-to-end. The Squad Model gets people using AI correctly. The API Mandate makes AI actually transformative.

## **Digital Assets**

The Protocol Vault at flattenthej.com demonstrates what Beta and Alpha-tier protocol usage looks like in practice—structured inputs generating professional outputs, with the prompt invisible to the user.

## Chapter 4 Takeaways

**Incentives drive everything.** People don't do what makes sense—they do what gets rewarded. AI deployment without structure produces toy usage because toy usage is easy, safe, and visible. Real value requires real structure.

**The \$1.8M email problem is universal.** 70% of firms are deploying AI. Only 18% are using it for predictive analytics. Only 17% achieved straight-through processing. The gap between deployment and value is the structure gap.

**Toy usage actively harms your workforce.** Cognitive offloading increases 72% with AI usage while critical thinking drops 68%. You're not just wasting money—you're degrading the capabilities you need most.

**The future is agentic, not generative.** 42% of financial services firms are already using or assessing agentic AI. McKinsey moved 45,000 employees from basic prompting through structured workflows to 12,000 deployed AI agents. Salesforce embedded AI into Slack and reclaimed 97 minutes per week per employee. Your certification tiers should follow the same progression.

**Certification isn't bureaucracy—it's a reliability requirement.** Princeton researchers found AI consistency scores of 30–75%. Same task, same inputs, different results a third of the time. You wouldn't deploy uncertified pilots. Don't deploy uncertified AI workflows.

**The Squad Model works because of Dunbar's 15.** Department-sized groups of ~15 people where Gamma/Beta/Alpha tier rankings are visible to every peer. Each squad defines proficiency for its own function. Social pressure drives adoption without top-down mandates.

**Enforcement sustains the model.** Automated access control, quarterly squad reviews, public executive participation. Without enforcement, even squad-level visibility fades. With enforcement, behavior changes in 90 days.

*Next: Chapter 5—The API Mandate: From Chatting With Bots to Bots Chatting With Bots*

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for principles on incentives, behavior patterns, and how structure determines outcomes.

**Drucker, Peter.** The Effective Executive. Harper & Row, 1967. Epigraph source.

**Dunbar, Robin.** Grooming, Gossip, and the Evolution of Language. Harvard University Press, 1996. Referenced for the 5-15-50-150 social layer model, sympathy group dynamics, and the neurology of peer accountability. See also “Beyond Dunbar circles” (Scientific Reports, 2022) confirming the geometric progression.

**McKinsey & Company.** “Rewiring the Way McKinsey Works with Lilli, Our Generative AI Platform.” McKinsey.com case study. Referenced for structured deployment data: 75%+ of 45,000 employees active monthly, 500,000+ prompts/month, 30% time savings, 17 sessions/week average, 12,000 AI agents deployed. Kate Smaje quoted via Bloomberg, June 2025. Initial team of 3 expanded to 70+ experts; phased rollout from 7,000 users to firmwide July–October 2023.

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents across financial services, fielded August–September 2025. Referenced for active AI usage (65%, up from 45%), agentic AI adoption (42% using or assessing, 21% deployed), positive ROI among adopters (83%), revenue and cost impact (89%), data challenges as top barrier (40%), spending priorities (41% optimizing existing workflows), budget increases (44% expecting >10% growth), and document processing as leading use case (32%).

**Rabanser, S., Kapoor, S., & Narayanan, A.** “Are AI Agents Ready for Production? A Systematic Benchmark of Reliability.” arXiv:2602.16666, 2026. N=14 models, 500 benchmark runs. Referenced for consistency scores (30–75%), reliability thresholds for production deployment, aviation certification analogy, and incident-reporting culture recommendations.

**Salesforce/Slack.** Anthropic enterprise case study. Referenced for 97 minutes/week reclaimed per employee through AI embedded in workflow tools. Rob Seaman (EVP/GM, Slack) quoted on the journey toward the “agentic enterprise.”

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage.” InvestOps Report 2026. Survey of 200

global buyside operations leaders. Referenced for AI deployment rates (70%), cultural resistance as top challenge (57%), manual workarounds (60%), unified data gaps (63%), predictive analytics expectations vs. adoption (66% vs. 18%), straight-through processing (17%), and operational team expansion (65%).

**CFA Institute.** Generative AI, Unstructured/Alt Data, and Open Source Survey. May 2024. N=1,210 global charter holders. Referenced for daily AI usage rates (35%) and top use cases (writing assistance 20%, research 18%, coding 18%).

**Abbas, M., et al.** “Cognitive Offloading and Critical Thinking: Exploring the Interplay of AI Assistance in Educational Contexts.” 2024. N=666 UK participants. Referenced for correlation between AI usage and cognitive offloading (+72%) and critical thinking (-68%).

*“Whosoever desires constant success must change his conduct with the times.”*

*—Niccolò Machiavelli, The Prince (1532)*

## Chapter 5

# The API Mandate

*From Chatting With Bots to Bots Chatting With Bots*

Morgan Housel makes an observation in *Same as Ever* about compound returns that most people understand intellectually but almost nobody applies correctly: the real power of compounding isn't visible in any single period. It's visible over time, when small advantages connect and multiply in ways that linear thinking can't predict. A 1% daily improvement doesn't look like much on Monday. It looks like a different universe by December.

That principle explains everything about why most enterprise AI deployments are stuck—and why a small number of firms are pulling away from the pack so fast that the gap will soon be unclosable. In a bank still running T+1 settlement on spreadsheets, that invisible 1% daily drag is the difference between absorbing the next regulatory wave and watching your ops headcount balloon 65% while your competitor pulls away.

Here's what compound returns look like in AI adoption: a firm that builds API-driven workflows connecting systems to systems doesn't just get faster. Each automated handoff eliminates a bottleneck that was constraining every downstream process. The next automation builds on the previous one. Capacity expands not in proportion to investment, but exponentially relative to it. Meanwhile, the firm still typing into chat interfaces gets the same linear productivity bump on Day 300 that it got on Day 1.

Machiavelli understood this five centuries ago. The firms that adapt their conduct to the times—not just their tools, but their architecture—are the ones that win. The firms that keep doing what worked last year, faster, become case studies in obsolescence.

This chapter is about making the architectural shift that separates the compounders from the linear thinkers. It's about moving from chatting with bots to bots chatting with bots.

## **The \$200 Billion Administrative Tax**

The U.S. healthcare system spends roughly \$200 billion annually on administrative work: claims processing, prior authorizations, billing disputes, and the armies of back-office staff required to manage it all. That's approximately what the country spends to treat cancer. Not because the work is inherently complex—but because the systems that handle it were never designed to talk to each other.

That's not just a healthcare problem. It's the enterprise AI problem, across every regulated industry.

The InvestOps 2026 survey of 200 global buy-side operations leaders tells the same story in financial services: 65% of firms expanded operational teams to handle compressed timelines like T+1 settlement rather than building systems that could absorb regulatory changes without adding staff. 63% lack unified data across front, middle, and back office, forcing manual reconciliation. 60% are implementing manual workarounds as temporary solutions that become permanent. 49% cite manual data reconciliation and validation as their top operational challenge.

They're spending money on technology. They're deploying AI tools. But they're deploying them at the edges—summarizing documents, generating emails, formatting reports—while their core workflows remain manual, fragmented, and expensive.

The problem isn't technology. The problem is architecture. Specifically, the architecture of human interaction with AI.

## **Human-in-the-Loop vs. Human-on-the-Loop**

Most AI deployments in 2026 are human-in-the-loop: AI generates an output, a human reviews it, the human approves or rejects it, the process moves forward. Every single transaction requires human touch.

This is safer than full automation. It's also a productivity trap.

Human-in-the-loop scales linearly with headcount. If you want to process twice as many transactions, you need roughly twice as many people reviewing AI outputs. You've made the work faster, but you haven't changed the underlying economics. Your capacity is still constrained by how many humans you employ and how many hours they work.

Human-on-the-loop inverts this model: AI executes workflows autonomously within defined parameters. Humans set the boundaries, monitor dashboards for exceptions, and intervene only when something falls outside acceptable ranges. The system doesn't wait for approval on every transaction—it runs continuously and escalates edge cases.

This is how you get from 17% straight-through processing (STP) to something closer to 80%. This is how you stop throwing headcount at regulatory changes and start building systems that absorb complexity without proportional cost increases.

The distinction isn't academic. It's the difference between incremental productivity gains and genuine capacity transformation.

## **The Health System That Stopped Playing Defense**

A denied claim lands on a billing coordinator's desk at 9:17 a.m. in St. Paul. Missing documentation. Again. Three hours and four phone calls later, the patient still doesn't know what they owe. Multiply by 5,000 visits a month and you have Allina Health in early 2025—before they flipped the architecture.

The old model worked like this: a provider would see a patient, deliver care, then submit a claim to the insurer. Days or weeks later, the claim would come back denied—missing documentation, wrong code, coverage question the billing team didn't catch at the point of care. The billing department would then spend hours gathering additional documentation, resubmitting, appealing. Multiply that by thousands of claims per month and you've got an enormous administrative machine whose primary function is cleaning up problems that could have been prevented.

Human-in-the-loop at every step. Every claim touched by multiple people. Every denial requiring manual intervention, manual research, manual resubmission. Healthcare organizations spend nearly \$20

billion annually just to overturn denied claims—and an estimated 84% of those first-time denials are avoidable.

Allina—a 12-hospital system with more than 90 clinics across Minnesota and western Wisconsin—deployed Optum Real, an AI-powered real-time claims system, initially for outpatient radiology and cardiology. The shift was architectural, not incremental. Instead of submitting claims into a void and waiting for denial notices, the system uses AI to validate coverage at the point of care—before the claim is ever submitted. It analyzes payer contracts, member benefits, and clinical documentation in real time, surfacing issues while the patient is still in the office. Claims that match standard parameters flow through automatically. Claims with potential issues get flagged immediately, with specific guidance on what documentation is needed.

The system doesn't ask permission on every transaction. It operates within boundaries defined by clinical protocols and payer rules. Humans handle exceptions, not routine validation.

Dave Ingham, Allina Health's Chief Information Officer, described the impact after processing more than 5,000 patient visits through the system: claims denials dropped meaningfully in the pilot departments. Prior authorization requests accelerated. Billing and coding staff stopped spending their days on preventable back-and-forth with insurers. Patients got clearer information about their coverage before leaving the office rather than receiving surprise bills weeks later.

But here's what matters for this book: Allina didn't shrink its billing team. The staff redeployed to higher-value work—resolving complex cases, improving the patient financial experience, handling the exceptions that actually require human judgment. Capacity expanded without headcount growth. And the productivity gain tied directly to the book's 30% capacity target: UnitedHealth Group's Q1 2025 SEC filing confirmed that its AI-powered claims processing tools delivered productivity increases exceeding 20% across revenue cycle management customers—with further headroom as Allina expands the pilot to surgical procedures.

This is what human-on-the-loop delivers at scale: systems that prevent problems rather than processing the paperwork to fix them.

## **Why Chat Interfaces Are Dead Ends for Enterprise Work**

The Allina case reveals something critical: the chat interface wasn't wrong. It just couldn't get them where they needed to go.

Typing questions into ChatGPT and copy-pasting answers is fundamentally human-in-the-loop work. The human is the integration layer between systems. AI generates value one interaction at a time, but the workflow still depends on humans to connect inputs to outputs, outputs to downstream systems, systems to decisions.

That architecture breaks at scale. Context switching costs compound: every time a human stops their primary work to prompt AI, review output, copy-paste results, and return to work, they pay attention overhead. Studies show it takes 23 minutes on average to return to full concentration after an interruption. If your workflow requires 15 AI interactions per task, you've just destroyed productivity, not enhanced it.

Knowledge doesn't transfer across sessions. Each chat session starts from zero. AI doesn't remember what you asked yesterday, what rules your organization follows, what decisions got made last week. Humans provide that continuity. That means every interaction requires human context-setting, human memory, human judgment about what's relevant. The human remains the system of record.

Error propagation requires human catch. When AI makes mistakes in a chat workflow, humans catch them downstream—sometimes. Other times, errors compound across multiple steps before someone notices. Human-in-the-loop validation at every step isn't just inefficient—it's unreliable, because humans get fatigued reviewing AI outputs and start rubber-stamping approvals.

Handoffs don't automate. You can use AI to draft an email, but you still have to send it. You can use AI to analyze data, but you still have to export results, format them, load them into the target system. Each handoff between AI and human and system creates friction, latency, and opportunities for work to stall.

Chat interfaces optimize for individual productivity—one person, one task, one interaction. API-driven automation optimizes for organizational capacity—systems executing continuously, humans supervising outcomes, value compounding across every connected

workflow. That's the Housel principle in action. Chat gives you linear returns. APIs give you compound returns. And over time, the gap between linear and compound becomes the gap between market leaders and everyone else.

## **From Operator to Architect**

The shift from chat interfaces to API-driven automation isn't just technical. It's conceptual.

Chat interfaces cast humans as operators: you provide instructions, AI executes, you review results. You're still doing the work. AI is a tool.

API-driven automation casts humans as architects: you design systems, define boundaries, set parameters. AI does the work. You supervise outcomes.

That shift requires different skills, different training, and different governance. It requires business leaders who can think in workflows, not tasks. It requires technical partners who can translate business logic into automated processes. It requires executives who can distinguish between "we made email summarization 40% faster" and "we eliminated three manual touchpoints from our core workflow."

Most firms don't have those capabilities yet. The InvestOps survey shows it clearly: 28% cite inadequate data quality, availability, or integration as a top AI challenge. 26% cite resistance to change among employees. 21% cite weak governance or risk controls. These aren't technology problems. They're adoption problems. And adoption problems only get solved when leadership invests in the boring infrastructure work that makes transformation possible.

## **What Compound Returns Actually Look Like**

To make the shift from chat to API concrete, consider how a portfolio monitoring workflow changes when you move from human-in-the-loop to human-on-the-loop.

The human-in-the-loop version: An analyst opens a dashboard every morning. Reviews overnight market movements. Identifies positions that may need attention. Pulls up relevant research. Drafts a summary for the portfolio manager. Schedules a meeting to discuss.

Portfolio manager reviews, asks questions, makes decisions. Analyst implements changes. Documents everything. An analyst missing a 2 a.m. market move doesn't get an alert—he gets a conversation with the PM about the 40 basis points he didn't protect.

That workflow takes 3–4 hours daily for one analyst covering one strategy. It scales linearly: two strategies need two analysts. Ten strategies need ten analysts. And it runs on human attention, which means it operates at the speed of calendars—daily reviews, weekly meetings, monthly reports.

The human-on-the-loop version: Market data feeds into a monitoring engine continuously. AI evaluates every position against pre-defined parameters—risk limits, concentration thresholds, performance triggers. Positions within normal parameters require no action. Positions approaching limits generate alerts with AI-generated context: what changed, why it matters, what the options are. Positions that breach limits trigger automatic rebalancing within pre-approved ranges. Only genuine edge cases—scenarios outside the system's operating parameters—escalate to humans with full context already assembled.

This runs continuously. Every trading day. Every hour if market conditions require it. The system doesn't wait for humans to pull reports and schedule meetings. It operates at the speed of markets, not the speed of calendars. The human-driven model maxes out at quarterly or monthly cadence because humans can't scale attention. The automated model runs continuously because systems don't get tired, don't get distracted, and don't need meetings to hand off work.

## **The Prerequisites Most Firms Are Ignoring**

Only 20% of firms in the InvestOps survey operate on a single platform enabling immediate decision-making. That means 80% are running fragmented systems where data lives in silos, workflows require manual handoffs, and end-to-end automation is impossible without massive integration projects. This is the real blocker to API-driven AI. It's not AI capability. It's architecture debt.

Before API-driven AI becomes viable, apply the E-S-A Framework in sequence: Eliminate → Standardize → Automate. Most firms try to skip straight to Automate. That's how you get expensive automation of

broken processes. The firms that achieved 17% full STP automation didn't skip this step. They built a unified data layer—single source of truth, common data models, real-time access across systems—and standardized their workflows before they touched an API. The NVIDIA survey confirms the data gap: privacy, sovereignty, and data scattered across disparate locations remain the number one challenge, cited by 40% of respondents across 839 financial services institutions. When front office makes a trade, middle office needs to see the position immediately, back office reconciles immediately, risk updates exposure immediately. No batch exports. No overnight processing. No manual reconciliation. Without that foundation, automation is theater.

Exception handling architecture is the governance question nobody assigns. Human-on-the-loop only works if you've designed clear escalation paths. When the AI hits a scenario outside its operating parameters, where does it go? Who sees the alert? What information do they need to make a decision? How does their decision feed back into the automated workflow? This isn't a technical question—it's a business ownership question, and it requires business owners to define the boundaries, not IT to guess at them.

Monitoring and observability close the loop. Chat-based AI fails obviously: you ask a question, you get a nonsense answer, you try again. API-driven AI fails silently. The system runs. Transactions process. And somewhere in the background, something goes wrong. Without real-time monitoring—dashboards showing throughput, error rates, exception volumes, processing latency—you're flying blind. You won't know the system is broken until downstream effects surface days or weeks later.

## **When Your AI Attacks Your Own Systems**

Most infrastructure failures are dramatic. A breach. A ransomware attack. A vendor outage that takes down trading for four hours. You know the moment it happens. People get paged. Incident response kicks in.

The infrastructure failure that AI causes looks different. It's quiet. And you caused it yourself.

Here's what happens. You deploy an AI agent to handle a workflow that your team currently processes manually. The team runs maybe 200 API calls a day—they're human-speed, working business hours, taking lunch. The agent doesn't take lunch. It doesn't sleep. And in its early weeks, before it's been tuned and optimized, it's inefficient. It makes redundant calls. It re-queries data it already fetched. It retries on soft errors instead of waiting. Within days, that 200-call-per-day workflow is generating tens of thousands of calls, all hammering the same internal APIs.

Your AI just launched a DDoS attack on your own infrastructure.

This isn't a hypothetical. It's a high-probability failure pattern in enterprise AI deployments, and it's almost never anticipated because the people designing the AI workflow and the people maintaining the API infrastructure are rarely in the same room. The data team made the data AI-ready. The development team built the agent. Nobody asked whether the underlying API layer was engineered for agent-scale velocity.

The distinction matters more than most executives realize: having AI-ready data and having AI-ready infrastructure are two entirely different problems. Clean, labeled, accessible data is necessary but not sufficient. The delivery mechanism—the APIs through which agents request, retrieve, and write that data—has to be rebuilt for a different operating model.

Traditional enterprise APIs are designed for application-scale usage. An application calls an API when a user clicks a button or submits a form. The volume is bounded by human behavior—working hours, transaction rates, process cadence. Agent-scale usage has no such bounds. An agent optimizing a portfolio doesn't wait for a user to click a button. It runs continuously, evaluating positions, checking parameters, re-querying on market moves. The same infrastructure that handles 200 daily calls from a human-operated workflow cannot absorb orders of magnitude more calls from an agent operating on the same process.

The fix is straightforward but requires deliberate design before deployment, not after the incident. Agent-aware rate limiting sets explicit ceilings on how many calls a given agent can make within a defined window—not because the agent shouldn't need more data, but because your infrastructure has limits that need to be respected and

agents, unlike humans, won't self-regulate. Batch endpoint design restructures data retrieval so agents can request large data sets in single calls rather than chaining thousands of individual queries. And before you put any agent into production, you stress test. Simulate swarm conditions. Find out what your API layer actually breaks at, not what it was designed for on paper.

The broader principle: moving from Application-First to Agent-First API architecture isn't a nice-to-have for advanced implementations. It's table stakes for anything beyond supervised chat workflows. If your Alpha users are building human-on-the-loop automation, they're building agent deployments. And agent deployments at scale will find every infrastructure weakness you didn't know you had.

Build for agents before you deploy them. The alternative is a 2 a.m. page about a system you deliberately built.

## **The Squad Model Revisited**

Chapter 4 introduced the Squad Model—the tiered certification framework that organizes AI users into Gamma, Beta, and Alpha levels within squads of roughly 15 people. Chapter 5 is where that framework becomes essential.

Gamma and Beta users operate primarily in chat interfaces. They're human-in-the-loop by design. They're learning, executing approved workflows, building competence within their squads. Peer accountability keeps them progressing.

Alpha users build and manage human-on-the-loop systems. They're designing automated workflows, configuring APIs, setting operating parameters, defining escalation rules. They're not using AI to make their own work faster. They're using AI to eliminate entire categories of work.

That's the shift from Little-to-Little (individual productivity) to Middle-to-Middle (workflow automation). That's what the Squad Model protects.

You don't want Beta operators building API integrations. They're not trained on failure modes, edge cases, security implications, or regulatory requirements for autonomous systems. You don't want

Gamma users anywhere near production workflows. But you also don't want to restrict API-driven automation to "when IT gets around to it." IT doesn't understand the business logic, doesn't feel the pain of manual workflows, doesn't have incentive to move fast on operational efficiency projects.

Alpha certification within squads creates business-owned automation. The people who know the workflows build the automation. IT provides platform, security, governance, and tooling. Business owns the outcomes. The social dynamics within squads—watching your peers advance from Beta to Alpha—creates the pull that centralized programs never generate.

NVIDIA's data confirms the direction: 42% of financial services firms are already using or assessing agentic AI—systems that don't just respond to prompts but autonomously reason, plan, and execute complex tasks. Twenty-one percent have deployed them in production. The firms building Alpha-level capability within their squads are the ones positioned to deploy agentic AI effectively. Everyone else will be scrambling to catch up.

Only 17% of firms have achieved straight-through processing. The other 83% are still paying the administrative tax in headcount and missed capacity. The Squad Model is how you close that gap—not by waiting for enterprise software vendors to build the perfect solution, but by empowering business users with the right training, the right access, and the right guardrails to build what they need.

## **The Compounding Advantage**

The firms that made it through their J-curve 18–24 months ago—the ones who built straight-through processing, unified data platforms, and API-driven workflows—aren't just more efficient. They're building capabilities competitors can't match. They're responding to regulatory changes without adding staff. They're processing 30% more volume with the same headcount. They're launching new strategies without multi-year technology projects.

They're living in the compounding returns phase of AI adoption. And they got there by stopping the chat with AI and starting the conversation between systems.

The InvestOps survey captures this inflection point: after two consecutive years prioritizing operational efficiency, 55% of firms now prioritize competitive differentiation through innovation. Portfolio management and trading innovation jumped from 32% to 56% in a single year. The firms with API-driven architecture can pivot toward those new priorities. The firms still stuck in chat-based workflows can't.

That's what Chapter 7 addresses: the foundational work—data quality, unified platforms, protocol engineering—that turns API-driven AI from concept to reality.

But before you build the foundation, you need to confront the elephant in the room: most of your processes shouldn't exist at all. Most of your workflows weren't designed. They evolved through decades of patches, workarounds, and "that's how we've always done it."

You can't automate your way out of broken processes. You need to destroy them first.

That's Chapter 6.

## Digital Assets

The tools to implement human-on-the-loop automation:

- 1. Human-in-the-Loop vs. Human-on-the-Loop Decision Matrix:** Framework for evaluating which workflows should require human approval at every step vs. human supervision of automated execution. Includes risk assessment, regulatory requirements, process maturity scoring, and recommended governance models.
- 2. API Integration Readiness Scorecard:** Assessment tool for determining whether your systems, data, and processes are ready for automated workflows. Evaluates data quality, system connectivity, process standardization, exception handling capability, and monitoring infrastructure.
- 3. Workflow Automation Opportunity Map:** Template for identifying high-value candidates for human-on-the-loop transformation. Prioritizes based on transaction volume, manual touchpoints, error rates, processing time, and strategic importance.
- 4. Exception Escalation Playbook:** Guide for designing clear exception handling rules—when does automated processing stop and require human review? What information does the human need? How do decisions flow back into automated workflows? Includes templates for 15 common exception types across finance, operations, and customer service workflows.

*Download all templates at [flattenthej.com](https://flattenthej.com).*

## What You Must Do Before Chapter 6

Compound returns beat linear gains—and the gap is irreversible once it opens. Chat-based AI delivers individual productivity improvements that scale with headcount. API-driven AI compounds value by connecting systems without human bottlenecks. The difference isn't visible in single interactions. It's visible at scale, and by the time it's visible, it's too late to close without a multi-year rebuild.

Only 17% of firms have achieved straight-through processing. The other 83% are still paying the administrative tax in headcount and missed capacity. Chat interfaces are the reason. Individual productivity tools scale linearly with headcount. The firms in that 17% made one architectural decision: they stopped deploying AI as a tool that humans operate and started deploying it as a system that humans supervise. That decision is available to every executive reading this book. It requires infrastructure investment, process work, and governance design—not new AI capability.

The Squad Model is your execution mechanism. Alpha certification within squads creates business-owned automation with proper training and guardrails. This is how you close the gap between “IT will get to it eventually” and “business builds what business needs, within guardrails IT sets.” Unified data, standardized processes, exception handling architecture, and real-time monitoring aren't optional. Without these foundations, API-driven AI is impossible. Build them before you automate—or you'll spend twice as much automating chaos.

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for compound returns framework and principles of unchanging human behavior in response to innovation.

**Machiavelli, Niccolò.** The Prince, trans. W. K. Marriott. London: J. M. Dent & Sons, 1908; originally published 1532. Chapter XXV. Epigraph source.

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside

Operations Leaders.” InvestOps Report 2026. Referenced for data on straight-through processing rates (17%), unified data gaps (63%), manual workarounds (60%), operational team expansion (65%), manual reconciliation challenges (49%), single-platform operation (20%), AI adoption challenges (28% data quality, 26% resistance to change, 21% governance), and strategic priority shift toward competitive differentiation through innovation (55%, with portfolio management and trading innovation rising from 32% to 56%).

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents across financial services, fielded August–September 2025. Referenced for data on data-related challenges as top barrier (40%), agentic AI adoption (42% using or assessing, 21% deployed), and industry-wide AI deployment patterns.

**Olsen, Emily.** “Optum Launches AI System to Speed Medical Claims.” Healthcare Dive, October 22, 2025. Referenced for Allina Health pilot of Optum Real; 12-hospital system across Minnesota and Wisconsin; pilot in outpatient radiology and cardiology; over 5,000 visits processed; reduction in claims denials and administrative errors; Dave Ingham (CIO) quoted on operational efficiency and patient experience improvements.

**Palmer, Katie.** “UnitedHealth’s ‘Optum Real’ Uses AI to Speed Up Medical Claims.” Bloomberg / Detroit News, October 21, 2025. Referenced for Allina Health pilot details including deployment since March 2025; meaningful reduction in claims denials; prior authorization acceleration; Dave Ingham quoted on staff redeployment; \$200 billion annual U.S. healthcare administrative spending figure; Sandeep Dadlani (CEO, Optum Insight) quoted on platform strategy.

**Minemyer, Paige.** “HLTH25: Optum Unveils New AI-Powered Claims Processing Platform.” Fierce Healthcare, October 22, 2025. Referenced for Allina Health expansion plans to surgical procedures; Puneet Maheshwari (SVP, Optum Real) quoted on ecosystem momentum.

**Optum.** “Optum Reinvents Claims and Reimbursement Process to Eliminate Complexity and Administrative Waste.” Press release, October 21, 2025. Referenced for 84% avoidable first-time denial rate; \$20 billion annual cost to overturn denied claims; McKinsey estimate of \$150–\$300 billion in potential administrative cost savings from AI and automation.

**UnitedHealth Group.** “First Quarter 2025 Results.” SEC Filing, April 2025. Referenced for confirmation that AI-powered claims processing tools increased productivity by over 20% for revenue cycle management customers; Optum Insight revenues of \$4.6 billion with \$32.9 billion revenue backlog.

**Morse, Susan.** “Optum Deploys AI for Real-Time Claims Processing.” Healthcare Finance News, October 23, 2025. Referenced for real-time coverage validation capabilities and multi-payer platform design.

*"It's not the daily increase but daily decrease; hack away at the unessential."*

—Bruce Lee, *Tao of Jeet Kune Do* (1975)

## Chapter 6

# Process Destruction

*Don't Pave the Cow Path*

Morgan Housel makes a deceptively simple observation in *Same as Ever*: complexity is a tax that compounds silently. Organizations add a process here, an approval layer there, a reporting requirement somewhere else—and each addition seems reasonable in isolation. But the cumulative cost is enormous. Not in dollars, though dollars are part of it. In attention. In speed. In the ability to change direction when the market demands it.

Bruce Lee understood this intuitively. The martial artist's path to mastery wasn't adding techniques—it was stripping away everything that didn't serve the objective.

Here's how this applies to AI adoption in regulated industries: most firms automate their existing processes without ever asking whether those processes should exist in the first place. They take broken workflows, manual handoffs, redundant approvals, and fifteen years of accumulated workarounds—and they make AI do them faster. That's not transformation. That's paving the cow path with high-speed internet.

## The Automation Trap

The InvestOps 2026 survey of 200 global buy-side operations leaders tells this story in stark numbers: while 51% of firms report streamlined workflows with minimal manual touchpoints, only 17% have achieved full straight-through processing. The other 83% are living in what Dean McIntyre of SimCorp calls hybrid workflows—partially automated systems that still depend on manual intervention at critical junctures. They automated the easy parts and left the messy parts alone.

It gets worse. When T+1 settlement compressed timelines across financial services, 65% of firms responded by expanding operational teams rather than rethinking their workflows. They threw headcount at a structural problem. They paved the cow path with more bodies instead of asking whether the path needed to exist at all.

This is the automation trap: the instinct to make existing processes faster rather than questioning whether they should exist. Teams default to automating what's easy or what's annoying. That leads to what I call *automating donkey work*—perfecting low-value, time-consuming tasks that are neither critical to business outcomes nor risky if they fail. You end up with beautifully automated nonsense that still requires a human to watch it, audit it, and defend it in the next regulatory exam.

AI doesn't fix broken processes. It exposes them. And if you're not prepared to destroy the processes that don't add value, your AI adoption will plateau at Toy Mode regardless of how much you invest.

## **The Two Questions That Matter Most**

Before you inventory processes for AI automation, you need a way to cut through the noise. Most departments have dozens or even hundreds of different workflows. Staring at that list triggers paralysis: where do we even start?

Start with two litmus test questions that force prioritization:

**Question 1: What is the most important thing we do?** This identifies the processes that directly drive revenue, customer satisfaction, regulatory compliance, or competitive advantage. These are the workflows where even small improvements compound across the organization. These matter.

**Question 2: What is the riskiest thing we do?** This identifies the processes where errors create material harm: regulatory violations, financial losses, customer attrition, reputational damage. These are the workflows that could blow up if ignored. These matter too.

Everything else? That's donkey work. It chews up time but doesn't move the needle. It might be frustrating, but it's not consequential.

When you overlay these two questions on your process inventory, three categories emerge. High-Impact or High-Risk processes are critical work that drives outcomes or prevents disasters—start here. Moderate-Impact or Moderate-Risk processes are legitimate work that supports operations but isn't make-or-break—second priority. Low-Impact and Low-Risk processes are donkey work. The temptation is to automate this first because it's easy. Resist that temptation. You're just encoding trivia into expensive systems.

These two questions give executives a pragmatic triage method: focus on work that genuinely matters to business outcomes or could genuinely blow up if mishandled. Everything else can wait—or be eliminated entirely.

## **The Destruction Value Formula**

The Two Questions sort your processes into three buckets. But within the high-impact, high-risk bucket, you'll still have a dozen candidates competing for attention. You need a way to decide which processes to put through the E-S-A framework first.

Chapter 3 introduced the Value Formula for Stepping Stone project selection: Value equals Utility divided by Cost. The same heuristic applies here, adapted for process destruction. Instead of evaluating which project to build, you're evaluating which process is worth the organizational energy to destroy and rebuild.

For process destruction, Utility means the capacity or efficiency gain from eliminating or transforming the process. How much time does this workflow consume? How many people touch it? How much capacity would be freed if it ran at half the manual effort? If this process disappeared tomorrow and was replaced by something clean, what would your team be able to do that they can't do now? That last question is the one that matters most, because it reframes the conversation from cost reduction to capacity expansion—the “Capacity, Not Time” argument from Chapter 9.

Cost means the organizational disruption of changing it. How many systems does this process touch? Who loses budget or headcount if it's eliminated? Does it have regulatory requirements that constrain what you can change? How much institutional knowledge is embedded in the current workflow that would need to be

documented before destruction? And critically: who will fight you? Every process has defenders. The question is whether the defense is rational (regulatory requirement, genuine risk) or political (headcount justification, empire protection).

The formula works the same way it does for Stepping Stone selection: it's a conversation tool, not a calculation. You're forcing your squad to articulate, out loud, exactly why a process is worth the disruption of destroying it. When an operations manager says "Eliminating the monthly reconciliation frees 200 hours a month" and someone else says "But it touches four systems, two vendor feeds, and requires sign-off from compliance," you're having the right conversation. You're surfacing the real cost before you commit.

Here's what this protects against: the same bias that plagues project selection. Teams gravitate toward destroying the most annoying process—the one that generates the most complaints—rather than the one with the best destruction ratio. The most annoying process might be deeply embedded, politically protected, and connected to six legacy systems. The destruction cost is enormous. Meanwhile, a less annoying but equally wasteful process sits inside one team with clean data and no political defenders. The ratio is better. Start there.

The formula also protects against a subtler mistake: treating donkey work destruction as a win. Eliminating a low-value process feels productive. Your team celebrates. But the capacity freed is minimal because the process didn't consume meaningful resources in the first place. You've optimized the irrelevant. The Destruction Value Formula forces you to quantify the numerator honestly: if the utility of destroying this process is small, the ratio is bad regardless of how easy it is to kill.

Now you've got prioritized processes. You know which ones are important, which are risky, which are donkey work, and which have the best destruction ratios. You're ready for the framework that determines what to do with them.

## **The E-S-A Framework: Eliminate, Standardize, Automate**

Most firms approach AI adoption backwards. They ask, “What can we automate?” The correct question is, “What can we eliminate?”

The E-S-A framework forces discipline. Three steps, in strict sequence. Skipping ahead is the most expensive mistake you can make.

**Step 1: Eliminate.** Does this process need to exist at all?

This is the question most firms never ask. Only 12% of firms start their AI assessment by questioning whether a process should exist at all. Every organization carries processes that survived long past their original purpose—reports nobody reads, approvals that duplicate other approvals, reconciliation steps that exist because two systems were never integrated, compliance checks for regulations that expired years ago. These processes persist because no one has the authority, the incentive, or the visibility to kill them.

The InvestOps data illustrates exactly what happens when you skip this step. 60% of firms are still duct-taping manual workarounds they swore were temporary. Those workarounds become the process. And then someone proposes automating the workaround, which locks dysfunction into your systems permanently. You’ve now spent six months building AI around something that shouldn’t exist.

Elimination requires executive air cover. Middle management protects processes that justify headcount and budgets. Without top-down mandate, sacred cows survive regardless of cost. The question “Does this output need to exist?” threatens people whose job is producing that output. That’s why it takes a COO or division head to ask it and mean it.

**Step 2: Standardize.** Can we remove 80% of the variation?

This is the step most firms skip entirely. They jump from “this process exists” to “let’s automate it” without ever examining the chaos underneath. And chaos is the right word. Most enterprise processes have accumulated years of exceptions, special cases, and one-off workarounds that were never formalized. Five people do the same task five different ways. Ten systems store overlapping data with conflicting formats. The process looks simple on a flowchart and looks like spaghetti in practice.

AI can’t learn from chaos. Machine learning requires patterns, and patterns require consistency. If your commercial loan review process

has 47 different variations depending on which officer handles it, which system feeds the data, and which regional regulations apply, no AI model will produce reliable results. You'll get garbage out because you fed garbage in—just faster.

Standardization means defining exactly what the process is, which steps are required versus habitual, which data inputs are necessary versus nice-to-have, and which decision points require human judgment versus rule-based execution. It means reducing 47 variations to 3 or 4. It means creating the clean, consistent substrate that AI needs to actually learn.

The InvestOps survey shows the gap: 63% of firms lack unified data across front, middle, and back office. That's not a technology problem. That's a standardization problem masquerading as a technology problem. You can't build straight-through processing when every department defines "complete" differently.

**Step 3: Automate.** Now, and only now, what can AI do?

Automation deployed on a clean, standardized process is transformative. Automation deployed on a messy, unstandardized process is a catastrophe with a technology budget.

This is where the 17% who achieved full straight-through processing separated from the 83% still stuck in hybrid workflows. They didn't deploy better AI. They spent 18 to 24 months eliminating waste and standardizing logic before deploying anything. When they finally automated, the AI had clean inputs, consistent patterns, and clear decision rules to learn from. The results compounded because the foundation was sound.

The sequence matters. Automating waste makes waste permanent. You encode dysfunction into your systems, and now it's locked behind code that no one wants to touch because "the system works." Except it doesn't. It just works faster at producing things no one needs.

Do this in the wrong order once and you will spend the next two years explaining to your board why the shiny new AI system is slower and more expensive than the old one.

**360,000 Hours to Seconds: JPMorgan's COiN**

The E-S-A framework sounds straightforward in theory. Here's what it looks like when a firm applies it to a process that everyone assumed was permanent.

JPMorgan Chase's lawyers and loan officers were spending 360,000 hours annually interpreting commercial loan agreements. To put that in perspective, 360,000 hours equals more than 41 years of continuous work. At scale, the bank was processing 12,000 new wholesale contracts per year, each requiring manual review to identify key clauses, risk provisions, and compliance requirements. The work was mind-numbing, expensive, and prone to errors—loan-servicing mistakes stemmed directly from human error in interpreting these contracts.

The obvious move: use AI to help lawyers read contracts faster. Speed up the existing process. Pave the cow path.

JPMorgan didn't do that. They applied E-S-A, whether they called it that or not.

**Eliminate:** They asked which elements of manual review were actually necessary versus habitual. Did every clause need a lawyer's eyes, or had the process expanded beyond its original purpose? Decades of practice had turned contract review into a sprawling exercise where lawyers reviewed entire documents even when the relevant provisions were concentrated in a handful of sections. Much of the review was redundant—checking standard boilerplate that hadn't changed across thousands of similar agreements.

**Standardize:** This was the critical step. They identified approximately 150 key attributes that mattered across commercial credit agreements—the specific clauses, terms, and risk provisions that actually drove decisions. Instead of 12,000 contracts with infinite variation, they created a structured framework: 150 attributes, consistently defined, consistently extracted. The chaos of “every lawyer reads it their own way” became a systematic, repeatable classification.

**Automate:** With standardized attributes defined, JPMorgan built COiN—Contract Intelligence—using machine learning and image recognition to extract those 150 attributes from each agreement. The system processes 12,000 contracts in seconds. Compliance-related errors dropped by approximately 80%. Loan-servicing mistakes driven by human interpretation fell dramatically.

The result wasn't just efficiency. It was capacity transformation. Legal teams stopped spending their days on routine document analysis and redirected to negotiation strategy, complex advisory work, and the judgment-intensive tasks that actually required their expertise. The bank didn't eliminate lawyers. It eliminated the work that was wasting their talent.

Run it through the Destruction Value Formula: the utility was massive (360,000 hours of annual capacity, reduced errors, faster processing), and the cost, while significant (standardizing 150 attributes across 12,000 contract types), was contained within one function. The ratio was obvious. JPMorgan has since expanded this approach across its operations, with its LLM Suite reaching 200,000 employees and investment bankers creating five-page pitch decks in 30 seconds. The bank's \$18 billion annual technology budget isn't buying faster versions of old processes. It's buying the elimination of processes that shouldn't exist and the transformation of processes that should.

## **10 Weeks to 10 Minutes: Novo Nordisk**

JPMorgan is a financial services story. Here's what E-S-A looks like in a regulated healthcare environment where the stakes are measured in patient outcomes, not processing speed.

Novo Nordisk's clinical study documentation process consumed more than 10 weeks from start to finish. For a pharmaceutical company, clinical documentation isn't bureaucratic overhead—it's the critical path between a drug's development and its availability to patients. Every week of delay in documentation is a week that a treatment isn't reaching the people who need it. The process had accumulated decades of manual handoffs, review cycles, approval gates, and formatting requirements. Multiple teams touched each document. Each handoff introduced delay. Each review cycle added weeks.

Louise Lind Skov, Novo Nordisk's Director of Content Digitalization, put it in pure Bruce Lee language: they didn't automate the existing documentation process. They destroyed it and rebuilt from the ground up. The result: documentation that previously took 10 weeks now takes 10 minutes.

Apply the E-S-A lens. **Eliminate:** They challenged which review cycles, formatting steps, and approval gates were genuinely required by regulatory bodies versus which had been added by internal teams over the years as institutional habit. In heavily regulated industries, the assumption is always that every step exists for a reason. Often it doesn't. **Standardize:** They created consistent templates, structured inputs, and defined exactly what information was needed at each stage—eliminating the variation that made every document a bespoke production. **Automate:** With clean, standardized inputs, AI could generate documentation that met regulatory requirements without the cascading human handoffs that had consumed weeks.

Skov described the impact not as operational efficiency but as fundamentally transforming how medicines reach patients. That's the Capacity, Not Time reframe in action. The conversation wasn't "we saved 10 weeks of staff time." It was "patients get treatments faster." Same transformation, completely different narrative—and the narrative that gets executive budgets renewed.

The Destruction Value Formula made this an obvious priority: enormous utility (faster drug-to-patient timelines in a company whose entire business model depends on drug development speed), and while the cost was real (regulatory validation, change management across clinical teams, documentation standard creation), the ratio was compelling. The process was high-impact, high-risk, and high-destruction-value. Exactly where E-S-A should be applied first.

## **Why Partial Automation Is Worse Than No Automation**

The 30% of firms in the InvestOps survey operating "adequate" hybrid workflows—balanced automation with routine manual processes—face a problem that's invisible on dashboards. Their systems handle routine cases beautifully. When transactions fit expected patterns, everything flows smoothly and the executive dashboard shows green across the board.

Then an edge case appears. A non-standard transaction. An exception that doesn't fit the automated rules. The system stops, the human intervenes, and the entire workflow reverts to manual processing for that case. If edge cases are rare, this is manageable. But in complex regulated environments, edge cases aren't rare. They're

constant. And every manual intervention reintroduces the variability, the errors, and the latency that automation was supposed to eliminate.

This is automation-induced brittleness. The system works perfectly until it doesn't, and when it doesn't, it fails harder than the fully manual process because nobody remembers how to handle the exception. The institutional knowledge that once lived in experienced operators has been lost to attrition and retirement. The automated system didn't capture it because it only needed to handle the standard path.

In a regulated environment one missed exception isn't a dashboard blip—it's a FINRA letter, an OCR audit finding, or a patient-safety incident that lands on the front page of the Journal. The Frozen Middle knows this. They've seen it. That's why they quietly keep the old manual override alive "just in case."

Dean McIntyre of SimCorp describes the consequence directly: for the 83% still operating hybrid workflows, T+1 settlement exposed the brittleness of manual handoffs. T+0 will make partial automation untenable, demanding the integration most firms have yet to build.

The fix isn't better automation. It's going back to Step 1 and Step 2. Eliminate the exceptions that shouldn't exist. Standardize the ones that should. Then automate with clear escalation paths for genuine edge cases—not the manufactured complexity of a process that was never cleaned up.

## **Squads as Process Destruction Units**

The Squad Model from Chapter 4 was designed for structured AI adoption. But squads serve a second, equally important function: they're the natural unit for process destruction.

Here's why. Process elimination requires deep operational knowledge. You can't decide what to kill from a conference room with a whiteboard. You need the people who actually do the work to tell you which steps are necessary, which are habitual, and which exist because of a regulation that expired seven years ago. That knowledge lives in squads—the 15-person teams closest to the workflows.

Alpha-certified users within squads are especially critical for this work. They've progressed through Gamma (learning) and Beta

(approved workflows) and have the operational context plus the technical capability to assess which processes are candidates for E-S-A treatment. A Gamma user doesn't have enough context to know what's unnecessary. A Beta user executes within existing workflows. An Alpha user can step back and ask whether the workflow itself should exist.

The process destruction exercise is designed as a squad-level effort. Each squad documents its top processes, applies the Two Questions (important? risky?), runs the Destruction Value Formula to prioritize within the high-impact bucket, and puts each surviving process through E-S-A. The squad's Business AI Disruptor facilitates; Alpha users provide technical assessment; Beta and Gamma members contribute operational reality.

This bottom-up approach solves the political problem that kills top-down process elimination. When a COO announces "we're going to eliminate unnecessary processes," middle management circles the wagons around everything in their domain. When a squad of 15 people who actually do the work identifies that three of their twelve processes serve no purpose, the evidence is harder to argue with. The elimination recommendation comes from the people closest to the work, not from consultants or executives who haven't touched the workflow in a decade.

NVIDIA's data reinforces the direction: 42% of financial services firms are already using or assessing agentic AI—systems that autonomously execute complex tasks. But agentic AI deployed on messy processes produces autonomous chaos. The firms that will succeed with agents are the firms that used squads to clean their processes first. You don't give an autonomous system the keys to a workflow that shouldn't exist.

## **What Remains Is Faster**

After two consecutive years prioritizing operational efficiency, 55% of firms in the InvestOps survey now prioritize competitive differentiation through innovation. Portfolio management and trading innovation jumped from 32% to 56% in a single year. Julia Cloud of Deloitte Global captures the inflection: streamlined execution is now the norm, underpinned by earlier simplification and standardization efforts that removed many reconciliation and touchpoint frictions.

That's the payoff of process destruction. The firms that did the hard work of elimination and standardization aren't just more efficient. They're free. Free to pursue new strategies, new markets, new products—because they're not spending 60% of their operational capacity on manual workarounds and hybrid workflows that drain talent and attention.

Bruce Lee was right. Hack away at the unessential. What remains is faster, more powerful, and more adaptable than what you started with.

You've now got the framework for structured AI access (Chapter 4, the Squad Model), the architecture for system-to-system automation (Chapter 5, the API Mandate), and the methodology for cleaning your processes before AI touches them (Chapter 6, E-S-A). The next question is foundational: what infrastructure needs to be in place before any of this works at scale?

That's Chapter 7—the boring, essential, career-defining work of building the technical foundation that makes AI transformation possible. Janitor Agents. Protocol Engineering. The infrastructure decisions that separate firms who deploy AI from firms who deploy AI that actually works.

## Digital Assets

The Permission Slip tool at [flattenthej.com](http://flattenthej.com) operationalizes the E-S-A Framework and the Destruction Value Formula from this chapter into an interactive exercise you can run with your leadership team.

Two card types drive the conversation. Existence cards challenge whether a process should exist at all—questions like “Why does this process exist? Can anyone in this room answer without guessing?” and “If this process disappeared tomorrow, who would notice and what would break?” Redesign cards challenge how it should be rebuilt: “What would you have to stop believing about this process to redesign it completely?” and “If you were building this function from scratch today, would this process be in it?”

Spin the wheel, discuss the card with your squad, decide. The tool is designed for the same 45-minute leadership conversation as the Start Here wheel from Chapter 3, but focused on destruction rather than selection. Where Start Here asks “What should we build?” Permission Slip asks “What should we kill?”

*Access the Permission Slip tool at [flattenthej.com](http://flattenthej.com).*

## Chapter 6 Takeaways

**Start with the two litmus test questions.** What's the most important thing we do? What's the riskiest thing we do? These questions prevent teams from defaulting to automating donkey work—low-value tasks that chew up time but don't drive outcomes or prevent disasters.

**Destruction Value equals Utility divided by Cost.** Use the same napkin heuristic from Stepping Stone selection to prioritize which processes to put through E-S-A first. High utility (massive capacity gain) and manageable cost (contained disruption) means start here. Don't waste E-S-A effort on processes that are merely annoying.

**Complexity is a tax that compounds silently.** Most processes survive because no one questions why they exist. AI doesn't fix broken processes—it exposes them. 60% of firms are still duct-taping manual workarounds they swore were temporary.

**The E-S-A framework forces discipline. Eliminate first. Then Standardize. Only then Automate.** JPMorgan reduced 360,000 hours of annual contract review to seconds—not by automating the old process, but by standardizing 150 key attributes and building AI on the clean substrate. Novo Nordisk compressed 10 weeks of clinical documentation to 10 minutes by destroying the old process entirely.

**Partial automation is worse than no automation.** Hybrid workflows handle routine cases beautifully and collapse on exceptions. In regulated environments, one missed exception is a FINRA letter, an audit finding, or a patient-safety incident. The 83% still operating hybrid systems face automation-induced brittleness that T+0 settlement will make untenable.

**Squads are natural process destruction units.** Alpha-certified users have the operational context and technical capability to identify what should be eliminated. Bottom-up recommendations from the people closest to the work are harder to argue with than top-down mandates from consultants.

*Next: Chapter 7—Building the Foundation: Janitor Agents and Protocol Engineering*

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for complexity-as-tax principle and organizational behavior patterns.

**Lee, Bruce.** Tao of Jeet Kune Do. Santa Clarita, CA: Ohara Publications, 1975. Epigraph source.

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside Operations Leaders.” InvestOps Report 2026. Referenced for straight-through processing rates (17% fully automated, 51% streamlined, 30% hybrid, 2% manual), operational team expansion for T+1 (65%), manual workarounds (60%), unified data gaps (63%), and strategic priority shift (55% competitive differentiation, portfolio innovation 32% to 56%). Dean McIntyre quoted on hybrid workflow brittleness and T+0 implications. Julia Cloud quoted on standardization and simplification as foundations for innovation.

**JPMorgan Chase.** COiN (Contract Intelligence) platform, deployed 2017. Referenced via multiple sources including Bloomberg, ABA Journal, and Harvard Business School. 360,000 annual hours of manual contract review reduced to seconds; 12,000 commercial credit agreements processed per year; approximately 150 key attributes extracted per contract; compliance errors reduced approximately 80%; loan-servicing mistakes reduced. LLM Suite reaching 200,000 employees by 2025; investment bankers creating five-page decks in 30 seconds; \$18 billion annual technology budget.

**Son, Hugh.** “JPMorgan Software Does in Seconds What Took Lawyers 360,000 Hours.” Bloomberg, February 28, 2017. Primary source for COiN deployment data and operational impact.

**Waldron, Derek.** Interview with McKinsey, October 2025. Referenced for LLM Suite user data and technology budget context.

**Novo Nordisk.** Clinical study documentation transformation. Referenced via Anthropic case study. Louise Lind Skov, Director of Content Digitalization, quoted on transforming how medicines reach patients. Documentation cycle reduced from 10+ weeks to 10 minutes through process destruction and AI-powered rebuild.

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents, fielded August–September 2025. Referenced for agentic AI adoption (42% using or assessing, 21% deployed in production).

*"The way to become rich is to put all your eggs in one basket and then watch that basket."*

*—Andrew Carnegie, address, 1885*

## **Chapter 7**

# **Building the Foundation**

*Janitor Agents and Protocol Engineering*

Morgan Housel wrote about how the best investments aren't the ones that sound exciting---they're the ones built on boring, reliable foundations. A company with clean books beats a company with a charismatic CEO. A diversified portfolio beats a hot stock tip. The infrastructure you can't see determines whether the structure you can see will still be standing when conditions change.

Carnegie understood this instinctively. Put all your eggs in one basket and watch that basket. The metaphor isn't about diversification---it's about focused attention on the thing that matters most. In enterprise AI, that basket is your data infrastructure. The models get the headlines. The infrastructure determines whether those models produce reliable results or expensive hallucinations.

That principle applies directly to every AI deployment decision you'll make this year. The most critical choices aren't which model to deploy or which vendor sounds most impressive in the pitch. They're the unsexy architecture decisions that determine whether your AI investments will scale beyond pilots or collapse under production load.

This chapter is written for executives who need to approve infrastructure budgets without having a computer science degree. You don't need to configure servers. You need to understand what you're buying, why it matters, and which constraints are negotiable versus which ones will kill your deployment if ignored.

## **The Infrastructure Problem in Hard Numbers**

The InvestOps 2026 survey of 200 global buy-side operations leaders reveals the foundation problem clearly: 63% of investment firms lack a unified, real-time data layer across front, middle, and back office. Another 62% cite data quality and integrity as their top operational challenge. Forty-three percent struggle with data fragmentation and duplication across multiple systems and vendors. These aren't AI problems---they're infrastructure problems that AI exposes and amplifies.

The consequences of getting this wrong aren't theoretical. In May 2024, the CFTC fined JPMorgan \$200 million for misreporting critical data fields---a data quality failure at one of the world's most sophisticated financial institutions. Aggregate regulatory fines are at historic levels: the SEC issued nearly \$5 billion in enforcement actions in 2023 alone. As AI-powered models expand from credit scoring to fraud detection to automated reporting, every data inconsistency becomes a potential compliance violation at machine speed.

SimCorp's Hugues Chabanis explains the core problem: AI trained on fragmented data just amplifies the inconsistencies firms spent decades cleaning up. Natural language prompting works on clean, unified data with complete lineage. Without it, prototypes succeed in controlled environments, but production systems fail when a single hallucination removes user confidence. Accuracy is key.

One hallucination. That's all it takes to destroy months of adoption work. The Bain & Company survey found that 39% of executives cite quality and accuracy concerns---specifically hallucinations---as a top reason slowing GenAI adoption. The fix isn't better prompts. It's better infrastructure.

Every month you delay unified data and Protocol Engineering is another month your J-curve stays steep. The 15% productivity dip that should last eight weeks stretches into four months. The Kill Zone

widens. Boards panic. Initiatives die. Fix the foundation first and you compress the valley; delay it and you widen the grave.

You already know clean data matters. You've heard the single-source-of-truth sermon from every consultant who's walked through your door. So we're not rehashing it here. The real value lies in the novel infrastructure concepts specific to AI adoption and flattening the J-curve: agent chains that automate data remediation, confidence scoring that quantifies trust, and Protocol Engineering that eliminates blank-box chaos. That's where we're spending our time.

## **How Information Actually Moves**

Before you approve another infrastructure investment, you need to visualize how information actually moves through your organization. Not the org chart version where data flows neatly from system to system. The reality version where information gets stuck, duplicated, contradicted, and manually reconciled by people who know which spreadsheet to trust.

Every piece of information goes through four stages before AI can do anything useful with it.

**Stage 1: Data Sources.** Where information originates---CRM systems, ERP platforms, external APIs, document repositories, supplier feeds. The InvestOps survey shows 43% of firms struggle with fragmentation across multiple systems. That fragmentation didn't happen by accident. It happened because each department bought the tool that solved their immediate problem without thinking about how it connects to everything else.

**Stage 2: Processing Layers.** Where raw data gets cleaned, transformed, and validated before AI touches it. This is the layer most firms skip entirely---which is why their AI hallucinates. This is where Janitor Agents come in, and we'll cover the architecture in detail next.

**Stage 3: AI Touch Points.** Where models interact with information---classification, generation, analysis, decision-making. Critical insight: if Stage 2 is broken, Stage 3 will fail catastrophically. The most sophisticated model in the world can't compensate for contradictory data.

**Stage 4: Output Delivery.** How AI-enhanced information flows back to business systems and users---automated reports, API calls, dashboard updates, exception alerts. If the first three stages are solid, Stage 4 is straightforward. If they're not, you get what the InvestOps survey describes: 60% of firms implementing manual workarounds as temporary solutions. Manual workarounds mean your infrastructure is broken. That's not a training problem. That's an architecture problem.

April Wilcox, Director of Investment Services at CalSTRS, captures the priority: the foundation of reliable predictive analytics is structured, well-governed data. That's why they're focusing on pulling everything into unified environments where it can be formatted, cleansed, and translated into a single, consistent language. Only then can AI models generate insights that are truly trustworthy and actionable.

## **The Five-Stage Agent Chain**

The original Janitor Agent concept---small, cheap AI models cleaning data and tagging metadata before the main model sees it---works for simple use cases. But in production, at enterprise scale, data cleaning isn't a single task. It's a sequence of decisions. And sequences are best handled by chains of specialized agents, each doing one thing well.

Picture your operations floor at 4:17 p.m. on quarter-end Friday. A pricing feed from a European vendor arrives 43 seconds late. One desk is already on the phone with Legal. Another is manually keying corrections into three different systems. Now imagine that same feed is greeted by five specialized agents working in sequence---before any

human even notices. That's not science fiction; that's the difference between surviving the Kill Zone and becoming a case study in what went wrong.

The most effective data integrity architecture deploys five specialized agents in sequence. Each has a narrow mandate, clear inputs and outputs, and full audit logging. No single agent has the authority to both identify a problem and fix it---that separation of concerns is what makes the chain trustworthy.

**Agent 1: Triage.** Scans incoming data for anomalies. Doesn't fix anything---flags. Missing fields, format inconsistencies, values outside expected ranges, duplicate records, timestamp conflicts. Think of it as the intake nurse in an emergency room: assess severity, assign priority, route to the right specialist. Translation for the budget meeting: this agent replaces the three FTEs who currently spend their days chasing format errors.

**Agent 2: Analysis.** Takes flagged data and determines the nature of the issue. Is this a formatting problem fixable automatically? A source conflict requiring reconciliation? A missing data point requiring enrichment from another system? Or a genuine anomaly suggesting upstream corruption? The analysis agent classifies and determines whether automated intervention is safe or whether the issue needs human eyes. Translation for the budget meeting: this is the senior analyst who used to spend four hours per exception---compressed into seconds.

**Agent 3: Decision.** Selects the appropriate remediation action based on the analysis classification. Formatting issues route to auto-correction. Source conflicts apply reconciliation rules based on confidence scoring. Missing data triggers enrichment queries to authoritative sources. Genuine anomalies escalate to human review. The decision agent operates within governance rules---it can't invent new remediation strategies. It picks from a pre-approved menu of actions. Translation for the budget meeting: this is your rule book---

automated, consistent, and never distracted by a fire drill on the other desk.

**Agent 4: Data Change.** Executes the selected remediation within strict governance boundaries---no unauthorized changes, no overwriting without logging, no modifications to data classified as "human review required." Every change is versioned. The original data is preserved. If the remediation turns out to be wrong, you roll back without losing anything. Translation for the budget meeting: every change is reversible, every action is documented, and your auditors will love you for it.

**Agent 5: Audit.** Logs every action taken by every other agent in the chain. What was flagged, classified, selected, changed, and what the before/after state looks like. This creates a complete, immutable audit trail that satisfies regulatory requirements and gives executives full traceability. When your compliance team asks "what happened to this data point?" the answer is documented, timestamped, and attributable. Translation for the budget meeting: this is how you answer the regulator's question before they finish asking it.

The key insight for executives: you're not replacing the people who manually reconcile data today. You're giving them a system that handles the 90% of cases that are routine so they can focus on the 10% that genuinely require human judgment.

## **Confidence Scoring: Quantifying Trust**

Underpinning the agent chain is a confidence scoring model. Every data point that enters your system gets assigned a trust score based on three factors: the reliability of its source, its historical accuracy, and its consistency with other data points in the same domain.

A pricing data point from Bloomberg with a ten-year track record of 99.8% accuracy and no conflicts with other sources might score 0.97 out of 1.0. A manually entered spreadsheet value from a regional office, updated inconsistently and conflicting with two other sources,

might score 0.34. The agent chain uses these scores to make reconciliation decisions. When two sources conflict, the higher-confidence source wins---automatically, transparently, and with full documentation of why.

Confidence scores aren't static. They update over time as sources prove themselves reliable or unreliable. A vendor feed that was accurate for three years but started producing errors last quarter will see its score degrade. A newly integrated system that consistently matches authoritative sources will see its score increase. This creates a self-improving data quality engine that gets smarter the longer it runs.

For executives, confidence scoring answers a question that has historically required tribal knowledge: "Which system do I trust?" Instead of relying on the operations analyst who's been there fifteen years and knows that Bloomberg is right for fixed income but the internal system is better for FX, the confidence scoring model codifies that institutional knowledge into a quantitative framework. When that analyst retires, the knowledge doesn't walk out the door.

Here's why external confidence scoring matters more than you think. Princeton researchers tested 14 AI models across 500 benchmark runs and found that agents' self-reported confidence carries almost no signal. Most models couldn't distinguish their own correct predictions from incorrect ones any better than random chance. Your AI doesn't know when it's wrong. That's the empirical case for an independent quality layer---Janitor Agents with externally computed confidence scores---rather than trusting the model's own assessment of its output. If you're relying on an AI system to flag its own hallucinations, you're relying on the one entity least equipped to catch them.

Any data flagged with low confidence or ambiguous outcomes is not acted on automatically---it gets kicked out for human review. This is the human-on-the-loop principle applied to data integrity. The agent chain handles high-confidence, routine remediation autonomously. Edge cases get routed to humans. Automation handles

the volume; humans handle the judgment calls. The audit agent documents everything.

## **Protocol Engineering: From Chaos to Control**

Here's the uncomfortable truth: your employees are terrible at writing prompts. That's not a training problem. That's a design problem.

The average ChatGPT Enterprise session involves 1.4 prompts. One attempt, maybe one follow-up, done. No iteration. No refinement. No structured thinking about what information the model needs to produce accurate results. You can train employees on prompt engineering best practices all you want. They won't use them. Because typing into a blank box is cognitively expensive, and humans optimize for effort, not quality. Remember the \$1.8M annual Toy Mode tax we calculated in Chapter 1---five thousand seats at thirty dollars a month, each generating marginally useful email rewrites. That's not an adoption problem. It's an infrastructure problem wearing adoption's clothes.

The enterprise solution isn't better training. It's better infrastructure. Protocol Engineering is the answer.

Instead of asking users to craft creative prompts in blank boxes, you give them structured input fields---like Mad Libs---that generate optimized prompts on the backend. Users never see the prompt. They see a form. The system handles everything else.

**Bad approach (blank-box prompting):** "Write a compliance memo about this trade exception."

**Good approach (Protocol Engineering):** User fills in structured fields: Trade ID (auto-populated from system), Exception Type (dropdown: Settlement Delay / Pricing Discrepancy / Counterparty Issue / Documentation Gap), Business Impact (dropdown: Client-Facing / Internal Only / Regulatory Reportable),

Resolution Deadline (date picker), and Additional Context (optional text field, max 200 characters).

The system executes Protocol #307 on the backend: generate compliance memo for the specified trade with the selected exception type, use the standard template, include regulatory citations from the compliance database, set the resolution deadline, flag for Legal review if regulatory reportable, include client communication template if client-facing, and cross-reference similar exceptions in the past 90 days to detect patterns. If more than three similar exceptions are found, escalate to Risk Committee with trend analysis.

User sees: "Compliance memo generated. Flagged for Legal review---exception is regulatory reportable."

The user provided four pieces of information. The system handled regulatory citations, template selection, escalation logic, pattern detection, and routing---all automatically. That's Protocol Engineering.

The same architecture works across regulated industries. In insurance, the structured fields auto-pull policy number, claim type, and state-specific citation tables. In healthcare, they pull patient encounter codes, formulary references, and HIPAA-compliant disclosure language. The form changes; the principle doesn't. Users fill in what they know. The system handles what they shouldn't have to look up.

The Princeton research validates this approach with hard data. Their study found that AI agents handle server crashes and technical failures gracefully---but break when instructions are rephrased with the same meaning. Same task, same intent, different words, different results. That's the empirical case for locking input structures. Every freestyle prompt introduces the exact variability the research shows degrades performance. Protocol Engineering isn't bureaucratic overhead. It's a reliability requirement backed by evidence. If agents produce inconsistent outputs when humans phrase requests

differently, the solution isn't training humans to phrase requests consistently---that's a losing battle. The solution is removing the human from the phrasing equation entirely.

Why this matters for executives: consistency---every compliance memo follows the same template logic, one protocol executed correctly every time. Auditability---when regulators ask "how did your AI generate this memo?" you point to Protocol #307 with logged inputs and documented decision trails. The InvestOps survey shows 48% of firms prioritize regulatory and compliance automation as a top AI application. You can't automate compliance if you can't explain what the system did and why. Control---you control what AI can and can't do by controlling which protocols exist and who can execute them.

This maps directly to the Squad Model from Chapter 4: Gamma users view and practice with protocols in sandbox mode. Beta users execute standard protocols against production data. Alpha users modify existing protocols and deploy new ones. No certification, no access. And when regulations change, you update Protocol #307 centrally. Every memo generated after that change is compliant. You don't need 200 analysts to individually learn new rules.

## **The Constraints That Actually Matter**

Every infrastructure conversation hits constraints. Budget, legacy systems, regulation, security. The question isn't whether constraints exist---they always do. The question is which are actually binding versus which are institutional inertia pretending to be constraints.

In regulated industries, some infrastructure decisions aren't yours to make. Financial services firms face data residency requirements, seven-year minimum audit trail retention, and model explainability mandates from SEC, FINRA, and OCC. Healthcare organizations deal with PHI encryption requirements under HIPAA and sixty-day breach notification timelines. Insurance companies navigate actuarial review

of AI-based pricing models and state-varying data privacy frameworks. If your CTO can't tell you definitively whether the proposed architecture satisfies all applicable regulatory requirements, you're not ready to approve the budget.

Legacy systems present a different challenge. True binding constraints must be integrated: core transaction systems, regulatory reporting that can't go dark during migration, customer-facing systems with contractual SLAs, and systems with decades of unmigrateable history. False constraints should be killed: reporting tools twelve people use out of habit, departmental databases duplicating authoritative sources, and workarounds that became permanent because nobody had authority to end them. Budget for integrating the first category. Kill the second. The most expensive infrastructure is the stuff you maintain because you're afraid to make a decision.

On budget, think in three tiers. Minimum viable ( $\$300K$ -- $\$800K$  first year) gets you past pilots but caps around 500--1,000 users. Production-scale ( $\$1.2M$ -- $\$3.5M$ ) adds the agent chain, confidence scoring, Protocol Engineering, and enterprise security for 1,000--5,000 users. Enterprise transformation ( $\$5M$ -- $\$15M$ +) supports multi-region deployment, real-time data sync, custom model hosting, and end-to-end transformation for 5,000+ users. The InvestOps survey shows vendor priorities: 57% cite stability as top criterion, 54% prioritize innovation access. Firms want established vendors who deliver cutting-edge capabilities.

## **The Foundation Value Formula**

Chapter 3 introduced the Value Formula for Stepping Stone project selection: Value equals Utility divided by Cost. Chapter 6 adapted it for process destruction decisions. The same heuristic applies here, adapted for infrastructure investment prioritization---because the budget tiers above don't tell you where to start. They tell you what it costs. You still need to decide what to buy first.

For infrastructure decisions, Utility means the number and breadth of AI use cases an investment unblocks. A unified data layer scores high because every downstream application benefits---the agent chain, confidence scoring, Protocol Engineering, reporting, and any future use case you haven't imagined yet. An integration with a single legacy system scores lower---it solves one bottleneck for one workflow. The question to ask: if we build this, how many teams can do something with AI that they couldn't do yesterday?

Cost means the same thing it means in every other application of the formula: organizational energy. How long does the implementation take? How many existing systems does it disrupt during migration? Does it require vendor contracts, procurement cycles, Legal review? Does it create dependencies on scarce technical talent you don't have? How much ongoing maintenance does it demand once deployed? Infrastructure cost is particularly treacherous because it compounds. A system that takes six months to deploy and requires two full-time engineers to maintain has a far higher real cost than its sticker price suggests. The engineers you assign to maintenance are engineers you can't assign to building new capabilities.

Here's how the formula changes the conversation. Without it, infrastructure investment debates default to two modes: the CTO presents a wish list organized by technical elegance, or the CFO demands an ROI calculation that nobody can honestly produce for foundational work. Both modes fail. The wish list ignores organizational reality. The ROI calculation forces fabricated precision on inherently uncertain investments. The Value Formula replaces both with a pragmatic question: which investment unblocks the most AI capability per unit of organizational energy?

Run your infrastructure backlog through the formula and patterns emerge. The unified data layer almost always wins---massive utility because every team benefits, and while the cost is real, it's a one-time migration that reduces maintenance burden long-term. The agent

chain scores high utility for firms with significant data reconciliation volume---if your operations team spends 40% of their time reconciling data across systems, automating that process unblocks capacity across the organization. Protocol Engineering scores high utility and relatively low cost because it's a software layer that sits on top of existing infrastructure rather than replacing it.

The investments that score poorly are the ones that sound impressive in vendor demos but solve narrow problems. A custom model hosting environment for one department's specific use case. A multi-region deployment when your operations are concentrated in two locations. An enterprise-grade orchestration platform when you have twelve people using AI in production. These are Tier 3 investments masquerading as Tier 1 priorities. The ratio exposes them: narrow utility, enormous cost.

The formula also protects against a common infrastructure trap: building for a future that may never arrive. Executives love to approve architectures designed for 10,000 users when the current deployment serves 200. That's not visionary planning---it's sunk cost creation. The Value Formula forces you to evaluate utility based on what the investment unblocks now, not what it theoretically supports in three years. Build for the next phase of the J-curve, not for the final state. You can always expand infrastructure. You can't recover the eighteen months you spent building a platform that's 95% unused.

Like the Stepping Stone and Destruction formulas, this isn't a calculation---it's a prioritization conversation. When the CTO says "we need a unified data platform" and the head of operations says "we need the agent chain for reconciliation," the formula gives you a structured way to ask: which one unblocks more downstream value per dollar and per month of implementation time? The answer might be "both, in sequence"---unified data first because the agent chain needs clean data to operate on. The formula surfaces dependencies that technical roadmaps often bury.

## Infrastructure Meets the Squad Model

The agent chain, confidence scoring, and Protocol Engineering don't deploy themselves. They need governance---and that governance maps directly to the Squad Model from Chapter 4.

**Gamma-certified users** interact with protocols in sandbox environments. They're learning which protocols exist, how inputs map to outputs, and what the system can and can't do. They see confidence scores but can't act on low-confidence data. They're building fluency without risk.

**Beta-certified users** execute standard protocols against production data. They run Protocol #307 to generate compliance memos. They query production systems through secure integration layers. They see and act on the agent chain's outputs for routine cases. But they operate within pre-approved boundaries---no custom protocols, no modifications to existing ones, no overrides on confidence scoring decisions.

**Alpha-certified users** are the infrastructure operators. They modify protocol logic when regulations change. They adjust confidence scoring weights when new data sources come online. They design new agent chain rules for novel data types. They deploy new integrations. Alpha users are rare---you might have five in a squad of fifteen---and they carry the responsibility that comes with production access.

This tiered access model ensures that infrastructure investments are protected by the same governance that protects AI usage. The agent chain enforces data quality rules---but those rules are set by Alpha users and executed by Beta users within squad-level accountability. The Business AI Disruptor from each squad coordinates with the infrastructure team to ensure protocols match operational reality.

## The Path Forward

You now have the map: the Four-Stage Information Flow shows where data breaks, the Five-Stage Agent Chain automates remediation with full auditability, Confidence Scoring quantifies trust because your AI can't, Protocol Engineering eliminates blank-box chaos, and the Foundation Value Formula tells you what to build first. The technology works. The question is whether people will use it.

Chapter 8 addresses that problem: radical transparency through AI Leaderboards, and how to flip the Frozen Middle from resisters into champions.

Infrastructure enables transformation. Adoption delivers it.

## **Digital Assets**

The Protocol Vault at [flattenthej.com](http://flattenthej.com) is a live demonstration of Protocol Engineering in action. Ten protocols across four categories---Sales, HR, Legal, and Vendor---show exactly how structured inputs replace blank-box prompting at enterprise scale.

Users fill in form fields. The system executes the protocol on the backend via AI. Users never see the prompt. They see a professional document generated from their structured inputs---a prospect research brief, a job description with bias-flagged language, a contract review summary, a vendor RFP customized to their industry and regulatory environment.

This isn't a downloadable PDF. It's the concept from this chapter working in production. The consistency, auditability, and control that Protocol Engineering promises---demonstrated.

Access the Protocol Vault at [flattenthej.com](http://flattenthej.com).

## Chapter 7 Takeaways

**Foundations determine outcomes.** 63% of firms lack unified data infrastructure. That gap prevents transformation regardless of model sophistication. The CFTC fined JPMorgan \$200 million for misreporting data fields---data quality failures carry real consequences at machine speed.

**The Four-Stage Information Flow.** Sources → Processing (Agent Chain) → AI Touch Points → Output Delivery. If Stage 2 is broken, Stage 3 fails catastrophically. 60% of firms are implementing manual workarounds because their infrastructure can't support automation.

**The Five-Stage Agent Chain.** Triage → Analysis → Decision → Data Change → Audit. Replaces manual data reconciliation with automated, auditable data integrity at scale. Humans handle the 10% that requires judgment; the chain handles the 90% that's routine.

**Your AI doesn't know when it's wrong.** Princeton researchers found that agents' self-reported confidence is no better than random chance. That's why you need Janitor Agents as an independent quality layer with externally computed confidence scores---not reliance on the model's own self-assessment.

**Protocol Engineering is empirically validated.** AI agents handle technical failures gracefully but break when instructions are rephrased with the same meaning. Locking input structures eliminates the variability that degrades performance. Users fill forms; systems execute protocols. Consistency, auditability, and control---mapped to Squad Model tiers.

**Foundation Value equals Utility divided by Cost.** Evaluate infrastructure investments by how many AI use cases they unblock per unit of organizational energy. Unified data layers almost always win. Custom solutions for narrow problems almost always lose. Build for the next phase of the J-curve, not the final state.

**Only infrastructure-first organizations escape the Kill Zone with their momentum intact.** Only 17% of firms achieved full straight-through processing (InvestOps 2026). The other 83% built AI on broken foundations. Fix the foundation first.

*Next: Chapter 8---Driving Adoption: The Radar Speed Sign and The Frozen Middle*

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for principles on foundations determining outcomes and how infrastructure work compounds over time even when invisible.

**Carnegie, Andrew.** Address to students at Curry Commercial College, Pittsburgh, PA, June 23, 1885. Later published in Andrew Carnegie, *The Empire of Business* (New York: Doubleday, Page & Co., 1902). Epigraph source.

**WBR Insights and SimCorp.** "The AI Race in 2026: Building Foundations for a Lasting Advantage. Perspectives from 200 Global Buyside Operations Leaders." InvestOps Report 2026. Referenced for unified data layer gaps (63%), data quality challenges (62%), manual reconciliation rates (40%), data fragmentation (43%), manual workarounds (60%), vendor priorities (57% stability, 54% innovation), regulatory/compliance automation priority (48%), and straight-through processing rates (17%). Hugues Chabanis quoted on AI trained on fragmented data. April Wilcox (CalSTRS) quoted on foundation of reliable predictive analytics.

**Bain & Company.** Generative AI Survey, December 2024 (N=199). Referenced for quality/accuracy concerns slowing GenAI adoption (39%), lack of in-house expertise (42%), and data security concerns (44%).

**Commodity Futures Trading Commission (CFTC).** Order filing, May 2024. JPMorgan fined \$200 million for misreporting critical data fields. Referenced via Insight Partners, "Financial Services Has a Data Problem," December 2024.

**Securities and Exchange Commission (SEC).** Enforcement actions totaling nearly \$5 billion in 2023, referenced for regulatory consequence of data quality failures.

**Rabanser, Stephan, Sayash Kapoor, and Arvind Narayanan.** "On the Reliability of LLM-Based Agents." Princeton University, arXiv:2602.16666, February 2026. Referenced for self-reported confidence carrying no signal (models unable to distinguish correct from incorrect predictions), instruction rephrasing fragility (agents break when same-meaning instructions are worded differently), and consistency scores ranging 30--75% across 14 models and 500 benchmark runs.

**Luria, David.** The Failure of I.T. Project Management: Why It's Broken and How to Fix It. Self-published, 2022. Referenced for the principle that systems built on bad foundations fail in production.

“What you do speaks so loudly that I cannot hear what you say.”

—Ralph Waldo Emerson, “Social Aims,” *Letters and Social Aims* (1876)

## Chapter 8

# Driving Adoption

### The Radar Speed Sign and The Frozen Middle

In *Same as Ever*, Morgan Housel writes about the power of social proof. People don't change their behavior because you explained why they should. They change because they saw someone else do it and get rewarded for it.

A toddler learns to walk not by reading an instruction manual but by watching other toddlers take wobbly steps. Investors pile into tech stocks not because they ran the DCF models but because their neighbor made 40% last year. Smokers quit not when their doctor lectures them but when their friends quit first.

This isn't weakness. It's how humans are wired. We're social creatures who scan our environment constantly for signals about what's working and what isn't. We trust our own eyes more than we trust your PowerPoint.

Emerson saw this 150 years ago. What you do speaks louder than what you say. No amount of executive messaging about “embracing innovation” will move the needle if employees can't see anyone around them actually innovating. And no mandate will stick if the people closest to the work—the middle managers, the squad leaders, the domain experts—aren't visibly using the tools themselves.

This chapter is about two problems that kill enterprise AI adoption after the technology works: invisibility and resistance. The first is solved by radical transparency. The second is solved by a five-step protocol that turns your biggest resisters into your most powerful advocates.

## The Visibility Gap

Your company spent millions on AI tools. You mandated training. You had the CEO record a video about “embracing innovation.” Three months later, the usage reports land on your desk and the numbers are ugly.

OpenAI’s 2025 State of Enterprise AI report—based on real usage data from over a million business customers and a survey of 9,000 workers across nearly 100 enterprises—found a widening gap between leaders and laggards that should alarm every executive reading this. Frontier workers are sending six times more messages than median users. Frontier firms are sending twice as many messages per seat as median enterprises. The capability gap isn’t about who bought the tools. It’s about who’s actually using them.

The Wharton School’s 2025 AI adoption study, conducted with GBK Collective across three consecutive years, paints the same picture from a different angle. Eighty-five percent of employees remain stuck at stages two and three of AI adoption—basic prompting and light assistance. Fewer than 10% reach stage four, where AI becomes a genuine collaborative partner. Meanwhile, 42% of companies abandoned most AI initiatives during the year, up from 17% in 2024.

Writer’s 2025 enterprise survey of 1,600 knowledge workers found that only 28% of employees know how to use their company’s AI applications. Not “don’t have access.” Don’t know how. And 35% are paying out of pocket for AI tools because the enterprise offerings aren’t meeting their needs.

The problem isn’t training. It’s visibility. Nobody knows what “good” looks like because the data is locked in IT dashboards that employees never see.

## **Radical Transparency: The Radar Speed Sign**

In the early 2000s, hospitals were losing the hand-hygiene war. Infection rates were climbing despite years of training, signage, and policy memos. Then a handful of systems tried something different: they installed real-time compliance scoreboards outside every ward. Not buried in quarterly reports. Not discussed behind closed doors. Visible to every nurse, physician, and visitor who walked past. Compliance rates at participating institutions jumped from the low

40s to nearly 90% within months. Nobody changed the rules. Nobody added incentives. They just made the data impossible to ignore.

The same principle operates on residential streets. Those radar speed signs don't issue tickets. They don't change the speed limit. They just show you your speed compared to the limit. And people slow down.

This scales to enterprise AI adoption. When you show people their behavior compared to their peers, they change without being told to change.

Most companies track AI usage in secret. IT runs reports. Leadership reviews them behind closed doors. Nobody tells the people who aren't using the tools that they're falling behind. The data exists. It's just invisible to the people who need to see it.

The fix is radical transparency. Publish the usage data. Make it visible by squad, by department, by team. Not daily active users—that's a vanity metric that tells you nothing about value. Track the metrics that matter: tokens per complex task, iteration rates per session, protocol usage versus blank-box prompting, and Gamma-to-Beta-to-Alpha certification progression within each squad.

Make it comparative. Make it impossible to ignore. And watch what happens.

## **The Monday Leaderboard**

Radical transparency requires an operating cadence, not a one-time dashboard. Here's how to build it.

Every Monday morning, the Business AI Disruptor from Chapter 3—or a designated squad lead—publishes a leaderboard to a channel visible to all squads. The leaderboard contains five fields for each squad: current Gamma/Beta/Alpha certification mix (how many members at each tier), tokens per complex task (are people using AI for substantive work or just email cleanup), protocol adoption rate (percentage of AI interactions using structured protocols versus blank-box prompting), cycle-time reduction on the squad's primary workflow, and one featured peer win—a specific example of a squad member who used AI to deliver a measurable result that week.

The featured peer win is the most important field on the leaderboard. It puts a name and a story next to the numbers. It's the difference between "Squad 7 reduced cycle time 18%" and "Maria in Squad 7 used Protocol L-01 to cut contract review from four hours to 45 minutes on the Henderson deal." The first is a statistic. The second is a story people repeat at lunch.

What do leaders do with the leaderboard? Three things. First, celebrate the top-performing squads visibly—mention them in leadership meetings, copy their squad lead on recognition emails. Second, investigate any squad that drops two weeks in a row. Not punitively. Diagnostically. Is the squad lead engaged? Does the squad have a blocker—a data access issue, a process bottleneck, a missing protocol? Third, rotate the featured peer win into a monthly company-wide showcase. This creates the flywheel: visibility drives adoption, adoption creates wins, wins create visibility.

## **What Visibility Looks Like in Practice**

Qualcomm's AI rollout is one of the best-documented examples of how transparency and peer pressure drive enterprise adoption without mandates.

The company deployed an enterprise AI platform initially with its marketing and communications team. During the pilot, 100% of users wanted to adopt the platform full-time—an unheard-of retention rate that gave the team confidence to scale. The onboarding approach was deliberate: cohort-based training, office hours, use case development workshops. Within months, there were over 100 active users.

Then the critical thing happened. Word spread. Other departments saw what marketing was producing and how fast they were moving. Senior Manager of Marketing Brent Summers described the dynamic: "WRITER helped us be wildly successful. So much so that other teams had fear of missing out and asked to get access as well."

That's the Radar Speed Sign in action. Nobody mandated that legal, product, analytics, sales, learning and development, HR, and internal communications adopt AI tools. They saw marketing saving 2,400 hours per month across all users. They saw 85% weekly active usage rates—60% using the tools multiple times per week. They saw

25 vetted use cases and 70 defined workflows generating visible, measurable output.

VP of Marketing Carmen True put it plainly: “We use it across all the different disciplines and it powers the whole collective suite of assets and deliverables that we create.”

The adoption engine wasn’t executive memos. It was visible success. One team demonstrated results. Other teams noticed. Fear of missing out did the rest.

This is Housel’s principle made operational. People copy what works. They copy faster when the results are visible. And they copy fastest when the people succeeding are their peers—not some abstract case study from a company they’ve never visited, but the team down the hall that’s suddenly delivering faster than everyone else.

## **Squads as Adoption Engines**

Chapter 4 established that squads of roughly 15 people—small enough for everyone to know each other, large enough for meaningful peer comparison—create the social container where transparency works. The Gamma-to-Beta-to-Alpha certification tiers are visible within each squad, generating organic pull without mandates.

This chapter applies that structure to adoption execution. The question isn’t whether squads drive behavior change—the research is clear that they do. The question is what, specifically, to make visible and how to translate visibility into action.

Squad-level leaderboards matter more than company-wide metrics. A company-wide dashboard showing “65% adoption rate” is abstract. A squad dashboard showing that 11 of your 14 teammates have completed Beta certification and you’re one of three people still at Gamma is visceral. It’s personal. It’s the difference between reading a national obesity statistic and stepping on a scale in front of your running club.

OpenAI’s enterprise data confirms this dynamic at scale. Usage of structured workflows—projects, custom GPTs, repeatable processes—increased 19 times year-to-date. Average reasoning token consumption per organization increased approximately 320 times in 12 months. The firms seeing these numbers aren’t the ones sending

more emails about AI strategy. They're the ones making usage visible, making progress measurable, and letting social dynamics do the heavy lifting.

## **The Frozen Middle**

You published the squad leaderboards. Usage is climbing. Gamma-to-Beta progression is accelerating. Then you hit a wall.

The wall has a name. Actually, it has about 40 names—the Senior Vice Presidents and Directors who've spent 15 to 20 years building expertise that AI now threatens to commoditize. This is The Frozen Middle—the management layer where AI initiatives go to die.

Writer's 2025 survey captured the scale of this problem: 42% of C-suite executives report that the process of adopting AI is “tearing their company apart.” Seventy-two percent say their company has faced at least one major challenge on the AI adoption journey. The challenges aren't technical. They're political: power struggles, conflicts, organizational silos, and—the survey's word, not mine—sabotage.

Middle management resistance isn't irrational. These are competent people who built careers on being the expert, the gatekeeper, the person who knows how things work. AI threatens that identity. Their power comes from accumulated knowledge—knowing which vendor to call, which process shortcut works, which client needs special handling. When AI can surface those answers in seconds, the value of 20 years of institutional knowledge feels like it's evaporating.

The InvestOps 2026 data tells you where this resistance shows up operationally: 26% of firms cite resistance to change among employees as a top AI challenge. But the more revealing number is what firms did when T+1 settlement arrived—65% expanded operational teams instead of accelerating automation. That decision doesn't happen at the C-suite level. It happens in middle management, where the people closest to operations choose the familiar path of adding headcount over the uncertain path of trusting AI to handle the load.

## **What Resistance Looks Like Up Close**

At a mid-sized energy company in the Permian Basin, the CTO rolled out an AI platform for geological analysis in early 2025. Forty-five senior managers were scheduled for onboarding. Thirty-eight found reasons to skip the first session. Seven showed up. Two of those spent the session on their phones.

Steve was one of the 38. Twenty-two years in petroleum engineering. Three patents. The guy other engineers called when a well log didn't make sense. Steve didn't say he was against AI. He said he was "too busy"—the polite version of "this threatens everything I've built."

The CTO made a smart move. Instead of mandating attendance, she asked Steve to review the AI's geological analysis for a formation the company was evaluating in the Delaware Basin. Not to use the tool. To grade it. To find what it got wrong.

Steve found three errors in the first report. He documented them thoroughly—formation depth misidentified, porosity estimate overstated based on incomplete log data, and a lateral spacing recommendation that ignored a known fault line. He delivered a five-page critique to the CTO with visible satisfaction.

Then the CTO showed him the second report—generated after the AI incorporated his corrections. The revised analysis identified a drilling pattern Steve hadn't considered: a correlation between two formation characteristics across 14 wells that would have taken his team weeks to map manually. Steve stared at it for a long time.

Within six weeks, Steve was running the AI geological analysis workflow for his entire division. Not because anyone mandated it. Because he realized the tool was better when his 22 years of domain expertise shaped its inputs. He wasn't being replaced. He was the reason the AI produced useful output instead of generic analysis. His institutional knowledge—the exact thing he feared losing—was what made the system work.

Three months later, 31 of the original 38 no-shows had completed onboarding. Not because the mandate changed. Because Steve told them to.

## **The Five-Step Unfreezing Protocol**

Steve's conversion wasn't an accident. The CTO followed a pattern—whether she knew it or not—that works across industries and functions. Here's the protocol, formalized.

**Step 1: Identify.** Find three to five middle managers who are already using AI successfully. They exist in every organization, even resistant ones. Look for the quiet adopters—the Director who built a protocol for her team without telling anyone, the SVP who uses AI for board prep but doesn't mention it in meetings. Document their results privately. You need proof before you need visibility.

**Step 2: Reposition.** Take the most respected resisters—the Steves—and reposition them as quality reviewers, not AI students. The framing matters enormously. “Complete this training module” triggers identity threat. “We need your domain expertise to evaluate whether this AI output meets our standards” triggers professional pride. Same activity. Opposite emotional register. Qualcomm's organic spread followed this pattern: Legal's Senior Director of Counsel saw what marketing was doing with trademark compliance workflows and recognized the application to her own domain. Nobody told her to adopt. She saw professional relevance.

**Step 3: Demonstrate.** Showcase the quiet adopters at visible forums—leadership meetings, department all-hands, quarterly reviews. Present business outcomes, not technology features. “Sarah's team reduced RFP turnaround from three weeks to four days” lands. “Our AI platform processes 40,000 tokens per complex query” doesn't. The audience is other middle managers, not the CTO. Frame every showcase as a peer success story, not AI evangelism.

**Step 4: Formalize.** Create an Alpha certification track within each squad that carries visible prestige—recognition at leadership meetings, first access to new capabilities, and explicit consideration for promotion discussions. The message isn't “learn AI or else.” It's “the people building automated workflows are the ones getting promoted.” Position Alpha certification as elevation, not compliance. Middle managers who combine 20 years of institutional knowledge with AI capability become your most valuable people—if they're willing to make the shift.

**Step 5: Accountability.** After six months of transparency, showcases, and prestige incentives, assign leadership of critical-path AI projects to managers with demonstrated adoption results. This isn't

punishment for holdouts. It's organizational design—matching leadership scope to demonstrated capability. AI-fluent managers lead high-visibility work. Managers who haven't engaged lead lower-priority initiatives. No demotion. No pay cut. A quiet organizational signal that AI fluency is now a prerequisite for leadership on high-stakes projects.

Most holdouts read the signal and enroll in certification within 90 days. The Wharton study found that among decision-makers lagging behind their peers, the constraints include tighter workplace usage restrictions and low trust. The unfreezing protocol addresses both—loosening restrictions through repositioning, building trust through demonstrated peer results.

## **The Gap Compounds**

OpenAI's enterprise data reveals something that should keep every lagging executive up at night: the gap between leaders and laggards isn't closing. It's widening. Frontier firms aren't just using AI more. They're using it for increasingly sophisticated tasks. They're building structured workflows, deploying custom applications, integrating AI into core business processes. The median firm is still experimenting.

NVIDIA's financial services survey tells the same story: 42% of firms are now using or assessing agentic AI. Twenty-one percent have deployed agents in production. These aren't the firms still fighting about whether middle managers should complete basic training. These are firms where Alpha-certified users within squads are building autonomous workflows that compound capacity week over week.

Every month you spend fighting the Frozen Middle is a month your competitors spend building capabilities you'll struggle to match. The J-curve from Chapter 2 is unforgiving enough. Adding six months of middle management resistance on top of the productivity dip turns a survivable trough into an extinction event.

The answer isn't more training slides. It isn't another CEO video about innovation. It's making success visible, making progress measurable at the squad level, and letting the most powerful force in organizational behavior—peer pressure—do the work that mandates never will. Publish the scoreboard. Elevate the converted skeptics. Tie

leadership scope to demonstrated AI fluency. That's the operating instruction.

## Chapter 8 Takeaways

**Visibility beats mandates.** People change behavior when they see peers succeeding, not when executives send emails. Publish squad-level usage data every Monday. Make progress comparative and impossible to ignore. Qualcomm’s 85% weekly active usage and organic cross-department expansion happened without mandates—through visible results and peer-driven FOMO.

**The Monday Leaderboard is your operating cadence.** Five fields per squad: certification mix, tokens per complex task, protocol adoption rate, cycle-time reduction, and one featured peer win. Celebrate top performers, investigate two-week declines, and rotate wins into company-wide showcases.

**The Frozen Middle is your biggest adoption risk.** Forty-two percent of C-suite executives say AI adoption is “tearing their company apart.” The resistance layer is middle management, where domain expertise meets identity threat. Mandates produce checkbox compliance. Prestige and peer pressure produce genuine transformation.

**Turn resisters into quality reviewers, then experts.** Reposition resistant domain experts as AI evaluators—grading the system’s output, not learning to use it. The shift from threat to pride happens when institutional knowledge becomes the ingredient that makes AI reliable. Steve’s 22 years of geological expertise made him the reason the system worked.

**Follow the five-step protocol: Identify, Reposition, Demonstrate, Formalize, Accountability.** Find quiet adopters, reframe resisters as reviewers, showcase peer wins, create prestige through Alpha certification, and assign critical-path AI leadership to managers with demonstrated results.

**The gap is widening now.** Frontier firms are sending six times more messages per user and building structured workflows that compound advantage. Every month spent fighting the Frozen Middle is a month competitors spend building capabilities you’ll struggle to match.

Next: Chapter 9—Making the Business Case: Capacity, Not Time—and the 90-Day Sprint

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for social proof as the primary driver of behavior change: people copy what works, not what's explained.

**Emerson, Ralph Waldo.** "Social Aims," in Letters and Social Aims (Boston: James R. Osgood and Company, 1876). Epigraph source.

**OpenAI.** The State of Enterprise AI: 2025 Report. Based on real-world usage data from enterprise customers and survey of 9,000 workers across nearly 100 enterprises. Referenced for: frontier workers sending 6x more messages than median users; frontier firms sending 2x messages per seat vs. median enterprises; weekly Enterprise messages increasing 8x year-over-year; structured workflow usage increasing 19x year-to-date; average reasoning token consumption increasing approximately 320x in 12 months; widening gap between leaders and laggards.

**Wharton Human-AI Research and GBK Collective.** "Accountable Acceleration: Gen AI Fast-Tracks Into the Enterprise." October 2025. Third annual study. Referenced for: 85% of employees stuck at stages 2–3 of AI adoption; fewer than 10% reaching stage 4; 42% of companies abandoning most AI initiatives in 2025 (up from 17% in 2024); lagging decision-makers constrained by tighter workplace restrictions and low trust.

**Writer and Workplace Intelligence.** "Generative AI Adoption in the Enterprise: 2025." Survey of 1,600 knowledge workers (800 C-suite, 800 employees). Referenced for: only 28% of employees knowing how to use AI applications; 35% paying out of pocket for AI tools; 42% of C-suite saying AI adoption is "tearing their company apart"; 72% facing major adoption challenges including power struggles, conflicts, silos, and sabotage.

**Writer.** "How Qualcomm Increases Employee Productivity and Satisfaction with Writer." Customer case study, October 2025. Referenced for: 100% pilot retention rate; 85% weekly active usage; 60% using tools multiple times per week; 2,400 hours saved per month; 25 vetted use cases; 70 defined workflows; organic expansion across legal, product, analytics, sales, L&D, HR, and internal communications; Brent Summers (Senior Manager, Marketing) and Carmen True (VP, Marketing) quoted on adoption dynamics and FOMO-driven expansion.

**NVIDIA.** State of AI in Financial Services: 2026 Trends. N=839 respondents. Referenced for: 42% using or assessing agentic AI; 21% deployed agents in production; continued acceleration among leading firms.

**WBR Insights and SimCorp.** “The AI Race in 2026.” InvestOps Report 2026. Referenced for: 26% citing employee resistance to change as top AI challenge; 65% expanding operational teams for T+1 rather than automating; 55% shifting top priority to competitive differentiation; portfolio management and trading innovation rising from 32% to 56%.

**Hand Hygiene Compliance Research.** Referenced for hospital scoreboard case. Multiple studies including WHO Clean Care is Safer Care initiative and Joint Commission hand hygiene monitoring programs have documented compliance increases from approximately 40% to 80–90% when real-time monitoring with visible feedback was introduced. Specific implementations vary by institution.

“Extraordinary claims require extraordinary evidence.”

—Carl Sagan, *Cosmos* (1980)

## Chapter 9

# Making the Business Case

Capacity, Not Time—and the 90-Day Sprint

Morgan Housel wrote about how a stock price is really just a number today wrapped in a story about tomorrow. The same principle governs every budget decision, every Board presentation, every enterprise initiative that lives or dies based on whether leadership believes in the outcome. The numbers matter—but only after the story makes them care.

Sagan demanded extraordinary evidence for extraordinary claims. In enterprise AI, the extraordinary claim is that this technology will transform your business. The extraordinary evidence isn't a 127-slide deck full of financial projections—it's a story that makes the CFO lean forward, backed by controls that address the risks keeping the Board up at night, proven in 90-day sprints that produce results faster than anyone expected.

Consider two ways to present the same AI implementation results. Data version: “AI implementation increased processing capacity by 30% while maintaining headcount, generating \$2.1M in incremental revenue over 18 months.” Story version: “The claims processing team went from telling customers ‘we’ll get to it in 3 weeks’ to ‘approved this afternoon’—without hiring a single person. The CEO stopped losing deals to faster competitors.”

Same facts. Different impact. The first version proves you're right. The second version makes people care. Trial lawyers don't lead with forensic evidence—they lead with narrative and circle back to facts. This chapter is about making the case that survives Month 6 when you're deep in the J-curve trough and the CFO is questioning every dollar.

But making the case isn't just about rhetoric. You need a specific frame that separates winners from losers in enterprise AI: the capacity

frame. You need a set of controls that address the five risks every CFO worries about before they'll sign off. And you need a repeatable execution engine—the 90-Day Sprint—that proves value on timelines that matter. This chapter gives you all three.

## **IBM's Capacity Expansion Playbook**

In May 2025, IBM CEO Arvind Krishna told the Wall Street Journal that the company had used AI agents to replace the work of several hundred human resources employees. Their internal platform, AskHR, now automates 94% of routine HR tasks—employment verification, internal job transfers, payroll queries, vacation requests, policy clarifications.

Here's what makes this a masterclass in business case framing: IBM's total employment actually went up. The company used the resources freed by HR automation to hire more programmers and salespeople—what Krishna called “critical thinking” roles that require human interaction and creativity. IBM employed 270,300 workers globally as of their 2024 annual report. Their AI division had grown into a \$6 billion business.

Krishna's framing was deliberate. He didn't walk into the boardroom and say “we're firing HR staff to save money.” He said “we're automating routine work so we can invest in growth.” Same math. Different story. One version triggers organizational resistance. The other gets budget approval.

This is the Capacity, Not Time principle in action. IBM didn't frame AI as a cost-cutting tool that eliminates jobs. They framed it as a capacity expansion tool that frees investment for higher-value work. The savings from automating rote HR processes didn't disappear into margin improvement. They were redeployed into programmers who build products and salespeople who generate revenue.

IBM's CHRO Nickle LaMoreaux described this explicitly: the company shares information about what skills are declining and what roles are growing, helping workers prepare rather than panic. Transparency turns potential resisters into partners. The narrative isn't “your job is being eliminated.” It's “your role is evolving, and here's the path.”

## The Pattern Repeats

IBM's capacity frame isn't an anomaly. The companies getting the biggest results from AI are all telling the same story—and it's never about saving time.

Spotify's engineering organization achieved a 90% reduction in time spent on complex code migrations. But Chief Architect Niklas Gustavsson didn't frame it as speed. He framed it as expanding who can build. Engineers could now push code changes from Slack on their phones using plain-language descriptions. People who never could have contributed to migrations before were suddenly doing it from a mobile device. That's not a time-savings story. That's a capability story.

Novo Nordisk's clinical study documentation process—the regulatory paperwork required before a drug can advance to the next phase of testing—went from ten-plus weeks to ten minutes. Director of Content Digitalization Louise Lind Skov framed the transformation not as automating the old process, but as fundamentally changing how medicines reach patients. When a pharmaceutical company can move drugs through documentation in minutes instead of months, the constraint shifts from paperwork to science. The organization can now advance more drugs, faster, with the same regulatory compliance team.

Salesforce and Slack's internal deployment of AI embedded in workflow tools reclaimed 97 minutes per employee per week. For a company of 5,000 knowledge workers, that's the equivalent of roughly 120 full-time employees of recaptured capacity—every week. EVP Rob Seaman framed this as the journey toward an “agentic enterprise”—not an enterprise that uses AI tools, but one where AI is woven into the fabric of how work actually happens. The 120 FTE-equivalents don't show up on a layoff list. They show up as faster response times, deeper customer engagement, and the ability to pursue opportunities that were previously impossible given bandwidth constraints.

The pattern across all four companies is consistent: dramatic efficiency metric, deliberately reframed as a capability story. Spotify expanded who can build. Novo Nordisk accelerated how medicines reach patients. Salesforce created 120 FTE-equivalents of new capacity. The math is identical to a cost-cutting story. The organizational response is completely different.

## Why the Capacity Frame Works

When you frame AI as capacity expansion rather than cost cutting, you unlock three strategic advantages that change the conversation entirely.

**You align AI with growth.** Instead of fighting over who gets fired, you're discussing how to capture market opportunities. The conversation shifts from "which positions do we eliminate" to "what new capabilities can we offer customers." Growth stories get approved. Cost-cutting stories get fought. The InvestOps 2026 survey of 200 global buyside leaders found that after two consecutive years prioritizing operational efficiency, 55% now prioritize competitive differentiation through innovation—the first time innovation has overtaken efficiency as the top strategic priority. The companies with excess capacity can pivot toward new priorities. The companies locked into cost-cutting commitments can't.

**You create flexibility for adaptation.** Capacity gives you options. When market conditions shift, you redeploy that capacity toward new priorities instead of being locked into headcount reduction targets. Portfolio management and trading innovation jumped from 32% to 56% in a single year in the InvestOps data. That kind of strategic pivot requires spare capacity. An organization that has already eliminated 120 positions can't suddenly pursue a new market opportunity that requires 80 people. An organization that redirected 120 FTE-equivalents of capacity into existing priorities can reallocate part of that capacity overnight.

**You build organizational support.** People who would resist a cost-cutting initiative will champion a growth initiative. When the claims processing team hears "we're going to eliminate 35 positions," they protect the old workflow and slow-roll the transition. When they hear "we're going to handle 53% more volume so the company can expand into three new states," they become your biggest advocates. This is why IBM's approach worked—and why Novo Nordisk's regulatory team embraced rather than fought the documentation transformation. The people doing the work saw a path to more meaningful work, not a pink slip.

The math is straightforward. Start with current state: how many transactions does your team process monthly, what's the average cycle time, where are the bottlenecks. Calculate the capacity ceiling—

maximum sustainable volume before quality degrades. Project AI-enabled capacity after implementing Protocol Engineering and workflow automation. Then frame the opportunity: “We can handle 30% more volume with the same team, creating room for strategic growth without proportional cost increases.” Compare that to: “We can process our current volume with 35 fewer people.” One version positions AI as an enabler of strategic objectives. The other positions it as a headcount optimization play. Which story survives Board scrutiny?

## What CFOs Actually Lose Sleep Over

You’ve reframed the ROI story. Now you need to address risk. The CFO cares about five specific things going wrong—and whether you’ve thought about them before they happen.

The FT/Infosys global survey of 300 board directors confirms how seriously governance leaders take this. Seventy-five percent expect AI to significantly improve strategic decision-making. Seventy-one percent say AI will enhance the timeliness and depth of information flows to the board. But fewer than half report having human-in-the-loop oversight for AI-generated insights, and only about a third conduct regular audits of AI models and data sources. The gap between enthusiasm and governance infrastructure is where risk lives—and where your business case needs to demonstrate you’ve built the controls your Board hasn’t gotten around to requiring yet.

Here are the five controls that get you from “this sounds risky” to “they’ve thought this through.”

**Budget Limits.** No AI agent initiates spending above a pre-set threshold without human approval. Hard-coded at the system level—not relying on AI to “follow guidelines.” Start deliberately low—think \$50, not \$500. When an agent attempts a transaction above the threshold, it generates an approval request. No approval within four hours? Transaction auto-cancels. You’d rather review fifty legitimate transactions than miss one catastrophic one. As the system proves itself, raise the threshold—but start where the CFO’s comfort zone is, not where yours is.

**Rate Limits.** System-level throttling that overrides agent instructions. Cap emails per hour, documents per day, API calls per

hour—at levels that prevent a marketing automation agent from blasting thousands of prospects with broken personalization before anyone notices. Rate limits are the AI equivalent of speed governors on company vehicles. They don't prevent you from getting where you need to go. They prevent you from causing a ten-car pileup on the way there.

**Human Circuit Breakers.** Big red button architecture: any employee can instantly halt AI operations with full audit trail. When triggered, all AI operations pause, all pending transactions cancel, and an incident report generates automatically. The key design principle: stopping should be easier than starting. If it takes three approvals to launch an AI workflow and one click to stop it, you've designed the system correctly.

**Data Privacy Controls.** Zero-trust data access: AI agents access only data required for their specific function. Customer service AI sees support tickets but not financial records. Finance AI sees transaction data but not HR records. Role-based access that maps to Squad Model tiers—Gamma users get sandbox data, Beta users get production data for their function, Alpha users get cross-functional access with full audit logging. Every data access is logged, timestamped, and traceable. When the regulator asks “who saw what, and when?” you have the answer in seconds, not weeks.

**Model Governance Protocols.** Automated performance monitoring that tracks accuracy, response time, and error rates against trailing 30-day averages. When accuracy drops or error rates spike, the system flags for human review. The InvestOps data reveals why this matters: only 31% of firms have achieved business-user-level data access. Most organizations still require IT support for complex queries. Building real-time visibility into AI performance from day one separates organizations that catch problems early from those that discover degradation six months too late.

When you walk into the CFO's office, lead with the risk controls. Not buried in appendix slides—on page 3 of your executive summary. The CFO's internal monologue shifts from “this sounds risky” to “they've thought this through.” That shift determines whether you get approval.

These controls aren't just governance theater for the Board. They're operational necessities. Princeton researchers studying AI

agent reliability found that agents handle expected failures—server crashes, timeout errors—gracefully. But they break in unexpected ways when conditions change: the same task, described with different words but identical meaning, produces different results. Budget limits, rate limits, and circuit breakers are what catch those failures before they reach production. They’re the infrastructure that makes AI capability safe to deploy.

## The 90-Day Sprint

You’ve got the story. You’ve got the risk controls. Now you need to prove it works—fast. The 90-Day Sprint isn’t a one-time project. It’s a repeatable framework for systematic AI adoption. Every 90 days, you audit, pilot, and protocol. Each cycle builds on the previous one. The organizations running sprint cycles are the ones moving from assessment to deployment while their competitors are still writing planning documents.

**Weeks 1–4: The Audit and J-Curve Presentation.** Map the target process by interviewing the people who actually do the work—not their managers. Document the workarounds, the manual fixes, the things people do that aren’t in the official process documentation. Every enterprise process has a shadow version—the way it actually works versus the way the process map says it works. Your audit needs to capture both, because the shadow version is where the real opportunity lives.

Calculate J-curve parameters using the framework from Chapter 2: depth and duration of the productivity dip, breakeven point, leading indicators. Build the 3-slide deck you’ll show the CFO in Month 4 when things look terrible. Slide 1: “Here’s where we said we’d be.” Slide 2: “Here’s where we actually are.” Slide 3: “Here are the leading indicators showing we’re still on track.”

Get this approved before you start the pilot. Get explicit commitment: protected 18-month budget regardless of Month 4 metrics. This is the insurance policy that keeps the Kill Zone from killing your initiative. If you skip this step—if you launch without pre-committing leadership to weathering the dip—you’re betting your career on the hope that nobody panics when productivity drops 15%. That’s not a bet. That’s a coin flip.

**Weeks 5–8: The Scrappy Pilot.** One team, one process, one metric. The Stepping Stone framework from Chapter 3 applies: high pain, visible results, room to fail, repeatable. Use Protocol Engineering to create locked-down workflows. Deploy with Squad Model tiers—Gamma users get sandbox access first, Beta users start using it with real data once certified.

Track J-curve metrics: process completion time, error rates, user satisfaction. But also track reliability—the four dimensions Princeton researchers identified as what matters for production AI. Consistency: does the system produce the same quality output for the same type of input across repeated runs? Robustness: does it hold up when conditions aren't perfect? Predictability: can the system tell you when it's uncertain versus confident? Safety: when it fails, are the failures contained and recoverable, or do they cascade?

When you stand in front of the Board at Month 4, deep in the J-curve trough, telling them that productivity is down but leading indicators are positive—consistency is at 85%, robustness testing shows graceful handling of edge cases, safety protocols have caught and contained three failures with zero customer impact—you're not just showing results. You're showing governance. That's the evidence that keeps the Board from panicking during the Kill Zone.

Capture the stories—the moments when someone realizes this actually helps. The analyst who used to spend three hours reconciling data and now spends twenty minutes. The compliance officer who caught a regulatory issue the old process would have missed. The customer who got a response in hours instead of weeks. Those stories become your political ammunition for the next phase.

**Weeks 9–12: The Protocol Rollout.** Refine the protocol based on pilot learnings. Design the certification track using the Squad Model framework. Launch the win story internally—not a generic “we're using AI now” announcement but a specific narrative with names, problems, solutions, and results. BCG's 2025 research found that employee-centric organizations are 7x more likely to be AI mature. Make early adopters visible and celebrated.

Then identify the next three Stepping Stone projects. Sprint 1 proved you could eliminate manual touchpoints in one workflow. Sprint 2 replicates that success in three more. Sprint 3 moves from workflow automation to end-to-end process transformation. Each

sprint builds on the infrastructure, the organizational muscle, and the political capital of the previous one.

The discipline of the sprint keeps you focused. Three months is long enough to prove value but short enough to maintain urgency. NVIDIA's data shows the top spending priority isn't chasing new use cases—it's optimizing AI workflows that already exist, cited by 41% of respondents. The organizations running 90-day sprint cycles are the ones turning pilots into operating systems.

## **From Sprint to Scale: The 18-Month Arc**

The 90-Day Sprint is the engine. But the business case that survives needs to show the Board the full journey—from first audit to organization-wide transformation. Here's how the sprints compound.

Months 1–3 are your first sprint. You're proving the concept with a single Stepping Stone project. The goal isn't transformation—it's proof of life. One team, one process, measurable results. This is where you earn the right to continue.

Months 4–6 are the Kill Zone. Productivity drops. Double-keying hits—people are doing the work the old way while also learning the new way, and workload increases roughly 1.5x before it decreases. This is where the Month 0 insurance policy pays off. You pre-committed leadership to an 18-month budget. You have the 3-slide deck ready. You're reporting leading indicators and reliability metrics, not just lagging productivity numbers. Your second sprint launches toward the end of this period, replicating the first sprint's success in three more workflows.

Months 7–9 are the inflection point. If you've built the infrastructure right—Protocol Engineering locking down inputs, Janitor Agents cleaning data, Squad Model certification ensuring quality—this is where the J-curve starts bending upward. The third sprint moves from Little-to-Little automation into Middle-to-Middle workflow integration. Individual task automation gives way to connected workflows where bots talk to bots through APIs.

Months 10–12 are the Victory Lap. Productivity has recovered and exceeded baseline. Capacity has expanded 20–30%. You have stories, metrics, reliability data, and organizational momentum. This is where you make the expansion ask—not for a new pilot, but for organization-

wide deployment. The Month 10 Don't Panic script is designed for exactly this moment: lead with the story, back it with the 30% capacity number, and build the expansion request into the narrative.

Months 13–18 are the compounding phase. Each sprint cycle now deploys faster because the infrastructure exists, the certification framework is established, and the organizational resistance has been replaced by organizational expectation. This is where AI isn't a project to be approved but an operating assumption to be leveraged.

## **The Case That Survives**

The business case that survives isn't the one with the most impressive projections. It's the one that tells a story leadership actually cares about, addresses risks they're actually worried about, and proves value on timelines that actually matter.

Frame AI as capacity expansion, not cost cutting. Show executives how you'll handle the five critical risks: budget controls, rate limits, circuit breakers, data privacy, and model governance. Lead with controls, not buried in appendices. Prove value in 90-day sprints that build momentum through visible wins. Report not just productivity metrics but reliability dimensions—consistency, robustness, predictability, safety—that demonstrate control.

NVIDIA's survey shows 83% of financial services firms now reporting positive ROI on AI investments, with 89% seeing impact on both revenue and costs. The InvestOps data shows the strategic shift from efficiency to innovation, with 55% of buy-side firms prioritizing competitive differentiation. The FT/Infosys board director survey shows 75% of directors expecting AI to improve strategic decision-making—but governance safeguards lagging far behind ambition.

The gap between where Boards want to go and how prepared they are to get there is your opportunity. The organizations that bridge that gap—with the right story, the right controls, and the right execution cadence—don't just implement AI. They build the kind of compounding advantage that competitors can't reverse-engineer from a press release.

Carl Sagan was right. Extraordinary claims require extraordinary evidence. But the evidence isn't a spreadsheet. It's a story backed by controls backed by sprint results backed by reliability data. That's the

business case that turns a technology experiment into a competitive moat.

Companies that frame AI as headcount reduction will get politics. Companies that frame it as capacity expansion, governed by real controls, and proven in short cycles before the organization panics—they get compounding advantage. The only question is which company you're running.

## **Digital Assets**

The Don't Panic tool at flattenthej.com provides three pre-written executive communication scripts—one for each critical J-curve milestone. Month 0 (“The Insurance Policy”) locks in the 18-month budget and gets the productivity dip on the record before deployment. Month 4 (“Hold the Line”) is the Kill Zone survival script: no apology, no hedging, leading indicators front and center. Month 10 (“The Victory Lap”) leads with the story, backs it with the 30% capacity number, and builds the expansion ask into the narrative. Each script includes talking points, anticipated objections, and the 3-slide J-curve update.

## Chapter 9 Takeaways

**Stories beat spreadsheets.** The CFO needs the numbers, but the story determines whether they care. IBM framed AI as “investing in growth,” not “firing HR staff.” Same math. Different outcome.

**Capacity, Not Time.** “Saved 2 hours” means nothing. “Process 30% more volume with the same team” means everything. The capacity frame aligns AI with growth, creates flexibility, and builds organizational support.

**Five risk controls get you approved.** Budget limits, rate limits, human circuit breakers, zero-trust data privacy, and automated model governance. Lead with controls on page 3 of the executive summary, not buried in appendices.

**Report reliability, not just productivity.** Track consistency, robustness, predictability, and safety during pilots. Reliability data is the evidence that keeps the Board from panicking during the Kill Zone.

**The 90-Day Sprint proves value fast.** Audit → Pilot → Protocol Rollout. Repeatable engine. Each sprint builds infrastructure, organizational muscle, and political capital for the next.

**The governance gap is your opportunity.** 75% of board directors expect AI to improve decision-making. Fewer than half have oversight in place. Build the controls they haven’t required yet.

\* \* \*

## Sources

**Housel, Morgan.** Same as Ever: A Guide to What Never Changes. Harriman House, 2023. Referenced for principles on narrative, stories vs. data, and how humans make decisions based on stories rather than spreadsheets.

**Sagan, Carl.** Cosmos. Random House, 1980. Epigraph source.

**IBM / Arvind Krishna.** Interview with the Wall Street Journal, May 6, 2025. Referenced for capacity expansion framing: AI replaced several hundred HR workers; AskHR automates 94% of routine HR tasks; total employment went up as savings were reinvested into programmers and

salespeople (“critical thinking” roles). IBM employed 270,300 workers globally (2024 annual report). AI division: \$6 billion business. CHRO Nickle LaMoreaux quoted via HR Brew (February 2026) on skills transparency and role evolution.

**Gustavsson, Niklas.** Chief Architect and VP of Engineering, Spotify. Quoted in Anthropic Enterprise Briefing, “The Briefing: Enterprise Agents,” February 2026. Referenced for 90% reduction in engineering time on complex code migrations and expansion of who can build software through plain-language code changes via Slack.

**Skov, Louise Lind.** Director of Content Digitalization, Novo Nordisk. Quoted in Anthropic Enterprise Briefing, “The Briefing: Enterprise Agents,” February 2026. Referenced for clinical study documentation reduction from 10+ weeks to 10 minutes and reframing as accelerated discovery-to-patient timelines.

**Seaman, Rob.** EVP and General Manager, Slack (Salesforce). Quoted in Anthropic Enterprise Briefing, “The Briefing: Enterprise Agents,” February 2026. Referenced for 97 minutes/week reclaimed per employee through AI embedded in workflow tools and “agentic enterprise” framing.

**Rabanser, Stephan, Sayash Kapoor, Peter Kirgis, Kangheng Liu, Saiteja Utpala, and Arvind Narayanan.** “Towards a Science of AI Agent Reliability.” arXiv:2602.16666, February 2026. Princeton researchers tested 14 models from OpenAI, Google, and Anthropic across 500 benchmark runs. Referenced for four reliability dimensions (consistency, robustness, predictability, safety) and the argument that agents handle expected failures gracefully but break when instructions are rephrased.

**NVIDIA.** State of AI in Financial Services: 2026 Trends. Survey Report. N=839 respondents across financial services, fielded August–September 2025. Referenced for positive ROI among adopters (83%), revenue and cost impact (89%), budget increases (44% expecting >10% growth), spending priorities (41% optimizing existing workflows).

**WBR Insights and SimCorp.** “The AI Race in 2026: Building Foundations for a Lasting Advantage.” InvestOps Report 2026. Survey of 200 global buy-side operations leaders. Referenced for strategic priority shift toward competitive differentiation (55%), portfolio management and trading innovation growth (32% to 56%), and data access gaps (31% business-user-level).

**FT Longitude and Infosys.** “Rethinking Board Insight: Why AI Is Becoming a Governance Essential.” Global Board Director Survey, 2025. N=300 board directors across 20 countries and 15 sectors. Referenced for board expectations for AI in decision-making (75%), information flow

enhancement (71%), human-in-the-loop oversight gaps (47%), and audit/validation gaps (36%).

**Boston Consulting Group.** “The AI Adoption Puzzle.” 2025. Referenced for employee-centric organizations being 7x more likely to be AI mature.

# Acknowledgments

First, I want to acknowledge the two friends to whom this book is dedicated: Anne Kahn and Bob Kaufman. They were both two of the sharpest, kindest, most intuitive people that I have ever met. I was lucky enough to count them as friends. Over two decades, their insights into human behavior, capability models, and change management shape many of the ideas in this book.

Second, I want to acknowledge my family, especially my beautiful wife, Anita, our grown children, Sarah and Dylan. All of my family. All of Anita's family. And all of our friends, which we are lucky to count in the hundreds.

Finally, I want to thank my friends, mentors, and team members at T. Rowe Price. They continue to challenge and encourage me to be a better version of myself. That includes: Vinit Agrawal, Mike Cotter, Katherine Crawford, Ed Dekema, Maria Elena-Drew, Ashley Hogan, Nicole Holman, LQ Huang, Matthew Keiser, Magnus Kjoeller, Heather McDonold, Brigitte Mitchell, Gabe Morris, Natasha Nicholas, Selina Pattraine, Ulla Pitha, Michael Ray, Suha Read, Chris Smith, Lindsey Taff, and Amanda Young.

# About the Author

I was barely back on my feet from Stage IVB cancer when AI arrived—not as a curiosity, but as an existential threat to the career, the expertise, and the professional identity I’d spent twenty years building. The same feeling I had in that hospital bed came back: the ground shifting under you, the rules changing faster than you can learn them, the dawning realization that what got you here will not get you there.

I survived the first crisis through extraordinary medicine and the people who refused to give up on me. I decided to face the second one the same way—not by waiting for someone else to figure it out, but by building the playbook myself.

This book is that playbook.

\* \* \*

David Luria is a PMP-certified program manager, Lean Six Sigma Black Belt, and certified change management practitioner with twenty years of enterprise transformation experience in financial services. He has led large-scale technology implementations, regulatory remediation programs, and organizational change initiatives across some of the most complex operating environments in the industry.

He is the author of *The Failure of IT Project Management* (2022) and *Flatten the AI J-Curve* (2026). He publishes under the Corso & Alexander imprint, named after his mother’s and his wife’s mother’s maiden names.

David lives in Frederick, Maryland, with his wife and two children.

\* \* \*

*For speaking inquiries, consulting engagements,  
or to access the digital tools from this book:*

**flattenthej.com**

david@flattenthej.com

## Appendix A

# Digital Asset Library

*Tools, Templates, and Executive Scripts*

Every framework in this book has a corresponding digital tool designed to move you from reading to doing. These are not summaries or checklists---they are working instruments built for the conference room, the boardroom, and the budget meeting. All assets are available at flattenthej.com.

Free tools are accessible immediately with no account required. The Protocol Vault is reserved for book purchasers and requires an access code printed on the inside back cover.

\* \* \*

### **Start Here --- Project Selection Wheel**

**Access:** Free, no sign-up required

**Location:** [flattenthej.com/start-here](http://flattenthej.com/start-here)

**Grounded in:** Chapter 3 --- Stepping Stone Projects

The hardest part of any AI initiative is not the technology---it is choosing where to start. The Start Here wheel is a facilitation tool designed to spark a 45-minute leadership conversation about which project deserves your first investment. Spin the wheel and it reveals two cards: one testing Utility (visibility and business impact) and one testing Cost (complexity and risk). The cards ask questions like "If we fix this, does a VP mention it at the next all-hands?" and "How many systems does this touch? Name them." Eight cards in each pool, randomly paired on every spin. The tool does not calculate a score. It forces the right conversation.

\* \* \*

## **Permission Slip --- Process Destruction Wheel**

**Access:** Free, no sign-up required

**Location:** [flattenthej.com/permission-slip](http://flattenthej.com/permission-slip)

**Grounded in:** Chapter 6 --- Process Destruction and the E-S-A Framework

Before you automate a process, you need permission to ask whether it should exist at all. The Permission Slip wheel uses the same spinning mechanic as Start Here but with a different purpose: challenging assumptions about existing workflows. Two card types appear on each spin. Existence cards ask questions like "Why does this process exist? Can anyone in this room answer that without guessing?" Redesign cards push further: "What would you have to stop believing about this process to redesign it completely?" Eight cards in each pool. This tool operationalizes the E-S-A Framework's first step---Eliminate---by making it safe to question what everyone else takes for granted.

\* \* \*

## **Don't Panic --- Executive Communication Scripts**

**Access:** Free with email registration

**Location:** [flattenthej.com/dont-panic](http://flattenthej.com/dont-panic)

**Grounded in:** Chapter 2 --- The J-Curve and the Kill Zone

Three downloadable Word documents containing board-ready scripts for the three critical moments of any AI implementation. These are not slide templates---they are the actual words you say when the room gets quiet and everyone is looking at you.

**Month 0: "The Insurance Policy."** The presentation you deliver before

deployment begins. Lock in 18-month budget protection, get the productivity dip on record, and establish the leading indicators you

will track through the Kill Zone. When Month 4 arrives and the CFO wants answers, you pull out this presentation and say: "As predicted."

**Month 4: "Hold the Line."** The board update you deliver when

everything looks wrong. You are in the Kill Zone. Do not apologize for the dip. Do not promise a faster recovery. Report the leading indicators, point back to Month 0, and hold the line. Includes a "What You Must Not Say" section---because what you do not say in this meeting matters more than what you do.

**Month 10: "The Victory Lap."** The board update you deliver when it is

time to claim the win and expand. Lead with a story---find a Sarah, someone whose daily work was visibly transformed. Then deliver the 30% capacity number. The expansion ask is built into the script.

\* \* \*

## **Protocol Vault --- AI-Powered Enterprise Protocols**

**Access:** Book purchasers only (access code on inside back cover)

**Location:** [flattenthej.com/protocol-vault](http://flattenthej.com/protocol-vault)

**Grounded in:** Chapter 7 --- Protocol Engineering

The Protocol Vault is Protocol Engineering in action. Ten pre-built protocols across four business categories---Sales, HR, Legal, and Vendor Management. Each protocol presents a structured form where you fill in business-specific fields. The system constructs and executes the prompt on the backend. You never see the prompt. You never write the prompt. You get a polished, enterprise-grade output tailored to your inputs. This is what replaces "prompt engineering" for the other 80% of the workforce.

## ***Sales Protocols***

**S-01 Prospect Research Brief.** Company overview, pain points, opening questions, and competitive positioning tailored to your target account.

**S-02 Proposal Executive Summary.** Client-ready one-pager with benefit-first framing and industry-aligned ROI language.

**S-05 QBR Agenda Builder.** Structured quarterly business review agenda with talking points, time allocations, and renewal framing.

## ***HR Protocols***

**H-01 Job Description Generator.** Compliant job description with bias-flag language review built in.

**H-03 Interview Scorecard.** Competency-based questions organized by interview round with legally compliant framing.

**H-06 Promotion Justification Memo.** One-page memo for HR and executive approval, structured around outcomes rather than tenure.

## ***Legal Protocols***

**L-01 Contract Review Summary.** Plain-language key terms, flagged non-standard clauses, and redline recommendations. Not a legal opinion.

**L-03 Compliance Incident Memo.** Regulatory citation lookup, remediation checklist, and escalation recommendation. Auto-routes to General Counsel if regulatory reportable.

**L-04 Vendor Due Diligence Summary.** Auto-flags GDPR, CCPA, and HIPAA applicability by data access level and vendor jurisdiction.

## ***Vendor Protocols***

**V-01 AI Vendor RFP Builder.** Complete, customized request for proposal for AI vendor evaluation. Tailored to

industry, use case, regulatory environment, and stated evaluation priorities. Includes weighted vendor scorecard.

\* \* \*

*All digital assets are maintained and updated at [flattenthej.com](http://flattenthej.com). If you encounter issues or have suggestions for new protocols, contact [david@flattenthej.com](mailto:david@flattenthej.com).*

## Appendix B

# Glossary of Terms

### *Key Concepts and Frameworks*

This glossary defines the core terminology, frameworks, and concepts used throughout Flatten the AI J-Curve. Terms are listed alphabetically. Cross-references point to related entries. For expanded discussion of any term, consult the chapter where it is introduced.

\* \* \*

**Alpha Squad** The most advanced tier of the Squad Model. Alpha squads deploy autonomous AI agents with human-on-the-loop oversight, handling end-to-end process transformation. Promotion from Beta requires demonstrated reliability and governance protocols. See Squad Model.

**Beta Squad** The middle tier of the Squad Model. Beta squads run approved AI workflows integrated into daily operations with human-in-the-loop oversight. Members have proven competency in sandbox environments and earned production access. See Squad Model.

**Bridge to AI** The third layer of each chapter's opening structure, connecting the historical epigraph and Housel Principle to the chapter's enterprise AI challenge. Part of the book's mandatory three-layer opening framework.

**Business AI Disruptor** A role embedded within business units---not IT---to own AI transformation at the operational level. The Disruptor translates between technical capability and business need, ensuring adoption is driven by the people closest to the work.

**Capacity, Not Time** The reframing of AI return on investment from time saved to capacity gained. Instead of "we saved two hours per analyst," the winning frame is "we can now process 30% more

volume without adding headcount." Stories beat spreadsheets. The target metric is a 30% capacity increase.

**Double-Keying** The period during AI adoption when employees perform their original work manually while simultaneously learning and supervising AI-assisted workflows. Workload increases approximately 1.5x before it decreases. Double-keying is the core mechanism that produces the J-curve's productivity dip.

**E-S-A Framework** Eliminate → Standardize → Automate. A three-step process destruction methodology applied in strict sequence. First, eliminate processes that should not exist. Then standardize what remains into consistent, repeatable workflows. Only then automate. Most organizations skip directly to automation---paving the cow path---and wonder why results disappoint.

**End-to-End Transformation** The final phase of the Little-to-Little → Middle-to-Middle → End-to-End migration path. Complete process transformation where AI handles entire workflows with human-on-the-loop oversight. The destination, not the starting point.

**Frozen Middle, The** The layer of middle management where AI initiatives go to die. Middle managers fear irrelevance---they built careers on expertise that AI threatens to commoditize. Always capitalized. Unfreezing the middle requires making adoption visible, safe, and career-enhancing rather than career-threatening.

**Gamma Squad** The entry tier of the Squad Model. Gamma squads operate in sandbox environments for learning and experimentation with no production access. The safe space where employees build AI fluency before graduating to Beta. See Squad Model.

**Housel Principle** The second layer of each chapter's opening structure. A two-to-three sentence summary of a relevant principle from Morgan Housel's Same as Ever, establishing the universal human pattern before connecting it to enterprise AI adoption.

**Human-in-the-Loop** An oversight model requiring manual human intervention for each AI-generated output before it proceeds. Every task is reviewed. Appropriate for high-risk or early-stage deployments but unsustainable at scale.

**Human-on-the-Loop** An oversight model where humans monitor AI operations and intervene only on exceptions or flagged anomalies. The AI operates autonomously within defined parameters; the human provides governance, not approval. The target state for mature enterprise AI deployment.

**Imagination Gap, The** The inability to envision what AI makes possible beyond current capabilities. Organizations trapped in the Imagination Gap pursue incremental efficiency gains---doing the same things slightly faster---rather than building entirely new business models or capabilities that only AI enables.

**J-Curve, The** The predictable productivity decline that hits around Month 4 of enterprise AI implementations, driven by double-keying, learning curves, and organizational friction. Productivity drops approximately 15% before recovering and exceeding baseline. The J-curve is not a failure---it is the cost of transformation. Understanding it prevents premature abandonment.

**Janitor Agents** Small, inexpensive AI models deployed to clean, tag, and structure data before a primary model processes it. Janitor agents handle metadata enrichment, format standardization, and quality checks---the unglamorous foundation work that determines whether enterprise AI produces reliable outputs. A core component of RAG 2.0 architecture.

**Kill Zone, The** Months 3--6 of enterprise AI adoption, where the J-curve productivity dip is deepest and executive patience is thinnest. The period when most CIOs get fired, budgets get cut, and initiatives get abandoned. Surviving the Kill Zone requires advance preparation: budget locked before deployment, leading indicators established, and stakeholder expectations set at Month 0.

**Little-to-Little** The first phase of the enterprise AI migration path. Individual productivity gains---one person using AI to do their own work faster or better. Necessary but insufficient. The trap is staying here and calling it transformation.

**Middle-to-Middle** The second phase of the enterprise AI migration path. Workflow integration where AI connects processes across teams and systems via APIs rather than chat interfaces. The phase where real organizational value begins to emerge.

**Protocol Engineering** The practice of building pre-structured, locked input frameworks---enterprise "Mad Libs"---that generate prompts on the backend without exposing prompt logic to end users. Users fill in business-specific fields; the system constructs and executes the prompt. Protocol Engineering replaces "prompt engineering," which is dead as an enterprise strategy.

**Squad Model, The** A tiered organizational framework for AI adoption built on Dunbar's research, organizing approximately 15-person squads across three certification levels: Gamma (sandbox/learning) → Beta (approved workflows) → Alpha (agent deployment). Progression is earned through demonstrated competency, not calendar time.

**Stepping Stone Projects** Fast-win AI implementations designed to produce visible results during the J-curve trough. Scoped to one team, one process, four weeks, with measurable outcomes. Stepping stone projects build organizational confidence and provide proof points that sustain executive support through the Kill Zone. Always lowercase "stone."

**Toy Mode** The state of trivial AI usage---email summarization, Slack message rewrites, basic chatbot interactions---that costs \ \$30 per seat per month and delivers negligible enterprise value. At 5,000 employees, Toy Mode represents \ \$1.8 million in annual waste. The chat interface is a dead end for 80% of enterprise work.

**Value Formula** Utility ÷ Cost. The fundamental equation for evaluating AI implementations. Utility measures the actual business impact delivered; cost includes not just licensing but adoption friction, training, double-keying overhead, and organizational disruption. A high-utility tool with crushing adoption costs still fails.

\* \* \*

*Digital tools referenced in this glossary---including the Protocol Vault, Start Here and Permission Slip spinning wheels, and Don't Panic executive scripts---are available at [flattenthej.com](http://flattenthej.com).*

# Index

*Note to formatter: Replace all "000" placeholders with actual page numbers after final pagination.*

## A

- Adoption curve. *See* J-Curve
- AI champions. *See* Business AI Disruptor
- AI governance, 000, 000
  - board-level oversight, 000
  - compliance requirements, 000
  - regulated industries, 000
- AI leaderboard dashboard, 000
- Allina Health, 000, 000
- API Mandate, 000, 000
  - board-level implications, 000
  - chat interface limitations, 000
  - enterprise integration, 000
  - human-on-the-loop transition, 000
- Autonomous agents. *See* Human-on-the-Loop

## B

- Beta tier. *See* Squad Model
- Board of directors, 000
  - governance role, 000
  - imagination gap and, 000
  - ROI expectations, 000
- Business AI Disruptor, 000, 000
  - embedding in business units, 000
  - job description template, 000
  - vs. AI Champion, 000

## C

- Capacity, Not Time, 000, 000
  - 30% capacity increase target, 000
  - ROI reframing, 000
  - vs. cost-cutting, 000
- Carnegie, Andrew, 000
- Change management, 000

- double-keying period, 000
- executive communication, 000
- Frozen Middle and, 000
- Chat interface limitations, 000
- Churchill, Winston, 000
- COiN (JPMorgan), 000
- Cost-cutting. *See* Capacity, Not Time

## **D**

- Data readiness, 000
  - Janitor Agents and, 000
  - scorecard, 000
- Digital assets, 000
  - 90-Day Gantt Chart, 000
  - Capacity Expansion ROI Calculator, 000
  - Existential Threat Audit Checklist, 000
  - Human-in-the-Loop Decision Matrix, 000
  - J-Curve Prediction Model, 000
  - Process Kill Sheet, 000
  - Squad Certification Design Kit, 000
  - Stepping Stone Selection Matrix, 000
- Double-Keying, 000, 000
  - definition, 000
  - managing through, 000
  - workload increase (1.5x), 000
- Drucker, Peter, 000
- Dunbar's research, 000

## **E**

- Earhart, Amelia, 000
- Einstein, Albert, 000
- Emerson, Ralph Waldo, 000
- End-to-End transformation, 000, 000
  - definition, 000
  - vs. Little-to-Little, 000
  - vs. Middle-to-Middle, 000
- E-S-A Framework, 000, 000

- Automate (step 3), 000
- Eliminate (step 1), 000
- Standardize (step 2), 000
- Existential Threat Audit, 000
  - eight critical questions, 000
  - regulated industries, 000

## **F**

- Financial services, 000
  - JPMorgan Chase, 000
  - Klarna, 000
  - regulatory environment, 000
- Frozen Middle, 000, 000
  - breaking through, 000
  - definition, 000
  - fear of irrelevance, 000
  - management layer resistance, 000

## **G**

- Gamma tier. *See* Squad Model

## **H**

- Healthcare industry, 000
  - Allina Health, 000
  - Optum partnership, 000
- Housel, Morgan, 000
  - Same as Ever principles, 000, 000, 000, 000, 000, 000, 000, 000, 000
- Human-in-the-Loop, 000, 000
  - decision matrix, 000
  - definition, 000
  - transition from, 000
- Human-on-the-Loop, 000, 000
  - autonomous agents, 000
  - definition, 000
  - exception handling, 000
  - transition to, 000

## **I**

Imagination Gap, 000, 000  
    board-level implications, 000  
    definition, 000  
    efficiency trap, 000  
    overcoming, 000  
Insurance industry, 000  
Integration. *See* API Mandate

## **J**

J-Curve, 000, 000  
    danger zone (Months 4–6), 000  
    definition, 000  
    Kill Zone and, 000  
    managing through, 000  
    Month 4 inflection point, 000  
    prediction model, 000  
    15% productivity dip, 000  
Janitor Agents, 000, 000  
    definition, 000  
    implementation guide, 000  
    RAG 2.0, 000  
JPMorgan Chase, 000  
    COiN system, 000  
    Proxy IQ, 000

## **K**

Kill Zone, 000, 000  
    CIO turnover during, 000  
    definition, 000  
    navigating, 000  
    vs. J-Curve (see also J-Curve), 000  
Klarna, 000

## **L**

Lee, Bruce, 000  
Lilli (McKinsey), 000  
Little-to-Little transformation, 000

- definition, 000
- limitations, 000
- upgrading from, 000

## **M**

- Machiavelli, Niccolò, 000
- McKinsey, 000
  - Lilli platform, 000
- Middle-to-Middle transformation, 000
  - definition, 000
  - workflow integration, 000
- Middle management. *See* Frozen Middle

## **O**

- Optum, 000
- Organizational change. *See* Change management

## **P**

- Pilot programs, 000
  - 30-day sprint, 000
  - stepping stone selection, 000
- Process Destruction, 000, 000
  - cow path metaphor, 000
  - E-S-A Framework, 000
  - litmus test questions, 000
  - only 12% of firms ask the right question, 000
- Process Kill Sheet, 000
- Productivity dip. *See* J-Curve; Kill Zone
- Protocol Engineering, 000, 000
  - definition, 000
  - enterprise Mad Libs, 000
  - locked protocols library (20 protocols), 000
  - vs. prompt engineering, 000
- Proxy IQ (JPMorgan), 000

## **R**

- RAG 2.0. *See* Janitor Agents

Regulated industries, 000  
    compliance challenges, 000  
    energy sector, 000  
    financial services, 000  
    healthcare, 000  
    insurance, 000

Return on investment. *See* Capacity, Not Time  
ROI. *See* Capacity, Not Time

## S

Sagan, Carl, 000  
Same as Ever. *See* Housel, Morgan  
Squad Model, 000, 000  
    Alpha tier (agent deployment), 000  
    Beta tier (approved workflows), 000  
    certification criteria, 000  
    Dunbar's research and, 000  
    Gamma tier (sandbox/learning), 000  
    15-person squads, 000  
    12-week rollout calendar, 000  
Stepping Stone Projects, 000, 000  
    30-day sprint, 000  
    definition, 000  
    High Value/High Pain matrix, 000  
    one team, one process, 000  
    selection criteria, 000

## T

Toy Mode, 000, 000  
    \$1.8M annual waste example, 000  
    \$30/seat/month cost, 000  
    definition, 000  
    escaping, 000  
    examples (email, Slack), 000

## V

Vendor evaluation, 000

RFP template, 000  
selection criteria, 000

## **W**

Workforce transformation, 000  
Business AI Disruptor role, 000  
capacity expansion, 000  
Frozen Middle, 000  
reskilling, 000