Velocity with Containment

An AI Adoption Doctrine for Mid-Market Leaders

Katerina Andreeva

The Brained Inc., New York

Version 1.0 · June 2026

Table of Contents

Abstract

Mid-market companies have adopted AI almost universally and have, almost as universally, struggled to show what it was worth. This document explains the gap and offers a way across it. The argument is that the mid-market holds a genuine structural advantage in AI adoption: decision velocity, the ability to move from a decision to a live system while a larger rival is still scheduling the meeting. Most firms destroy that advantage by imitating the enterprise, importing its governance, its roadmaps, and its coordination overhead along with its slowness. The discipline that preserves the advantage is containment: holding that speed inside two deliberate guardrails, a clearly defined business goal and a realistic grasp of recurring cost.

The deeper claim is that the determinant of success is not technology but the human layer. Motivated, capable, involved people are what turn a working tool into a working capability, and AI-native people are the precondition for an AI-native company. The document gives C-level leaders of US mid-market firms a strategic frame for making the adoption decision well, before governance and before tooling, and leaves the execution to leaders who, once they are thinking clearly, are entirely capable of it. It is a doctrine, not a manual: it supplies the distinctions and the small number of load-bearing ideas that make a good decision possible, and ends each chapter on a single question the reader must answer for their own company. Its argument is organized around a set of original frameworks developed by the author, including the Motivation Taxonomy, the Imitation Trap, and Velocity with Containment.

Preface — Why This Document Exists

Almost every mid-market company in the country has now started using artificial intelligence. In a 2025 survey of mid-market decision-makers, ninety-one percent reported using generative AI, up from seventy-seven percent only a year earlier, a rate of uptake with few precedents in business technology.[1] By the ordinary measure of technology adoption, the question is settled: the mid-market did not hesitate, did not wait for permission, and did not fall behind.

And yet the same period produced a second, harder number. In MIT's 2025 study of AI in business, roughly ninety-five percent of AI pilots produced no measurable impact on the profit-and-loss statement.[2] A separate global survey found the same shape at larger scale: near-universal adoption set against fewer than four in ten organizations able to point to any effect on earnings.[3] Nearly everyone has adopted; almost no one can show what it was worth. That gap, between adoption and value, is the entire subject of this document. The question is no longer whether to use AI. It is why so much use produces so little, and what the companies on the other side of that gap understood that the rest did not.

This document concerns the adoption of generative AI and large-language-model-based systems, the class of technology that drove the 2024–2025 adoption wave, rather than predictive machine learning or narrow automation, which carry different cost structures, risk profiles, and human dynamics.

It helps to be honest about how most mid-market leaders arrive at this question, because the starting conditions shape everything that follows. Few come to AI from a calm strategic review. They come to it from indirect pressure: a competitor's announcement, a customer's offhand question, a board that has read the same headlines everyone else has read and now wants to know what the company is doing

---

[1] RSM US Middle Market AI Survey 2025 (RSM US in partnership with Big Village; 966 middle-market decision-makers across the US and Canada; fieldwork February–March 2025). The ninety-one percent figure, up from seventy-seven percent the prior year, is the mid-market-specific anchor used throughout this document in preference to enterprise-skewed sources.

[2] MIT NANDA, *The GenAI Divide: State of AI in Business 2025* (July 2025). The roughly ninety-five percent figure refers to AI pilots that produced no measurable profit-and-loss impact. The study's base is modest, on the order of 300 deployments, 52 interviews, and 153 surveys, so the figure is used here for the scale of the gap, not for decimal precision.

[3] McKinsey & Company, *The State of AI in 2025*. The survey reports roughly eighty-eight percent of organizations using AI in at least one function, against only about thirty-nine percent reporting a material effect on earnings before interest and taxes. The sample skews to larger enterprises and is cited here only to show that the adoption-without-value gap is not unique to the mid-market.

about them.[4] The pressure is real, but it is secondhand, and secondhand pressure tends to produce secondhand decisions: a tool bought to have an answer ready, a pilot launched to be able to say one exists.

Compounding this, most leaders by their own account were not ready for what followed. A majority describe themselves as only “somewhat prepared” and a further tenth as not prepared at all; sixty-two percent found adoption harder than they expected, and seventy percent had to reach outside the company for help to get anywhere.[5] This is not a story of incompetence. It is a story of capable people making a consequential decision under borrowed urgency, without a frame for making it well.

That frame is what is missing, and its absence is the gap this document is written to fill. There is no shortage of guidance on AI adoption. Almost all of it is calibrated to the enterprise, written for organizations with thousands of employees, dedicated transformation offices, and the capital to absorb a year of false starts. Imported into a company a fraction of that size, enterprise guidance does not merely fail to fit; it misleads, prescribing machinery that a mid-market company cannot afford and does not need.

The mid-market has been handed the enterprise’s playbook and told to run the enterprise’s plays, and the resulting mismatch is a large part of why so much adoption produces so little.

There is a second, quieter absence as well. Most experienced leaders already hold sound intuitions about why adoption works: that people have to want to use the tool, that speed is an asset, that a clear goal matters more than a clever system. What they rarely have is those intuitions connected, in one place, to the body of knowledge that explains and validates them. This document does both. It speaks to the mid-market on its own terms, and it grounds the practitioner’s intuition in the science that backs it.

One organizing choice separates this document from nearly everything else written on the subject, and it is worth stating plainly at the outset. Almost every mid-market AI guide begins with governance, with its policies, controls, committees, and apparatus of oversight. Governance matters, and this document takes it seriously in its place. But governance is not where adoption begins, and starting there is one of the surest ways to produce motion without value.

---

[4] BCG, *Split Decision: The BCG CEOs and Boards Survey* (May 2026), reports that sixty-one percent of CEOs say their boards are pushing AI transformation faster than the business is ready for, consistent with the pattern that leaders arrive at AI through indirect pressure. That sample skews to companies above $100 million in revenue and is cited here as corroboration of the direction of pressure, not as a mid-market measurement.

[5] RSM US Middle Market AI Survey 2025. Fifty-three percent of respondents described themselves as only “somewhat prepared” for AI adoption and ten percent as not prepared; sixty-two percent found adoption harder than expected; seventy percent required outside help.

Adoption begins upstream of governance, at a question most documents skip past on their way to the controls: what, honestly, is the company trying to achieve, and is that goal worth the cost and the disruption of pursuing it? A company that has answered that question well can govern a clear purpose. A company that has not will govern a vacuum, and no amount of oversight will turn an unexamined motive into a result. Starting before governance, at the honest definition of goals, is the central commitment of this document, and the thread that runs through every chapter that follows.

That thread has a name, and it is the one on the cover. The mid-market's structural advantage is velocity, the ability to move from decision to live system while a larger rival is still scheduling the meeting; the discipline that turns velocity into results rather than waste is containment, holding that speed inside deliberate bounds so that fast never means blind. Velocity with containment is the spine the rest of the paper hangs on, and each chapter is one facet of it. The idea takes its full shape in Chapter 3.

A word on what this is, and what it is not. This is written for the people who own the adoption decision: the C-level leaders of US mid-market companies, those with roughly $10 million to $1 billion in revenue, and secondarily the operational leaders who will carry the decision into the work.[6]

It is a doctrine, not a manual. It will not hand the reader a checklist, a software shortlist, or a day-by-day plan, because those things age within months and, more than that, because they answer the wrong question. The purpose here is to reorganize how a smart executive thinks about the adoption decision, supplying the frame, the distinctions, and the small number of load-bearing ideas that make a good decision possible, then leaving the execution to leaders who, once they are thinking clearly, are entirely capable of it.

The tone throughout is not one of alarm. The pressure that brought the reader here is real, but fear is a poor guide and a worse strategy; it produces exactly the imitative, anxious decisions that this document exists to prevent. What AI can already do for a mid-market company is substantial, and what is coming is more so. The reason to engage is not to avoid falling behind. It is that the opportunity, approached deliberately, is one of the largest available to a company of this size in a generation, best seized by those who decide on purpose rather than by reflex.

[6] Mid-market defined per the National Center for the Middle Market as companies with $10 million to $1 billion in annual revenue, segmented into Lower ($10M–$50M), Core ($50M–$500M), and Upper ($500M–$1B).

Chapter 1 — AI Adoption Is a Strategy Decision, Not a Technology Project

Of all the small choices a company makes when it first decides to take AI seriously, the most consequential is also the most invisible: where the decision lands. In most mid-market companies it lands by reflex, not by judgment. Someone forwards the question to the person who owns technology: the CIO, the head of IT, the most technical lieutenant available, on the unstated logic that AI is software, software is technology, and technology is their department. The routing feels so natural that no one experiences it as a decision at all.

It is, however, the first and most expensive decision the company will make, because it reclassifies the entire effort. Sent down that path, AI becomes a thing to be implemented, evaluated, procured, installed, and rolled out, when what the company faces is a change in how it creates value. That is the category error this chapter exists to correct, and almost everything that later goes wrong can be traced back to it.

The error is understandable, because AI arrives wearing the costume of ordinary software. It is bought from vendors, it runs on the existing stack, it has logins and license terms and an admin console. Every surface signal says system migration, and the company already knows how to handle a system migration: scope it, assign it to IT, manage it as a project with a start, a middle, and a finish.

But the resemblance is superficial. A new CRM changes where information is stored; it does not change what the work is. AI changes the work itself: which tasks people do, which they hand off, how fast the business can move, what it can now offer a customer, and where its costs and capabilities sit relative to its competitors. A technology you install leaves the business model intact and improves a process inside it; a technology that rewrites the business model is not a project but a strategic choice wearing a project's clothes. Treating the second like the first is not a small misallocation of responsibility; it is a misunderstanding of what kind of thing is happening.

That misunderstanding has a structural consequence, and it is the reason the decision has to originate at the C-level rather than be delegated to a function. AI adoption reshapes three things at once: the operating model, how work flows and who does it; the cost structure, what the business spends and on what; and the competitive position, what the company can offer and at what price relative to its rivals. Each of those belongs to a different part of the building, and no functional head has authority over all three.

The head of technology can stand up a tool, but cannot decide that the company will keep its people and expand their capacity rather than cut headcount; that is an ownership-level choice about the operating model. The head of a single line of business can change a workflow, but cannot commit the recurring spend or accept the competitive trade-offs the choice implies across the rest of the company.

When a decision reorganizes the operating model, the cost base, and the market position simultaneously, the only seat with the authority to own it is the one that already owns all three. Push it any lower and it does not get smaller; it gets fragmented, executed in pieces by people who can each see one face of it and none of whom can see, or own, the whole.

What the fragmentation looks like in practice shows up in the contrast between two companies that made the same move and ended in different places. *From practice:* two firms in the same industry adopted the same class of tool in the same year. One handed the question to its technology function as a purchase, asking it to compare the market, run a pilot, and buy the best product. The rollout was clean and the tool was sound, and a year later almost nothing had changed, because no one above the project had decided what the business would do differently now that the capability existed. The other treated the question as strategy before contacting a vendor, working out what it was trying to achieve, what that would cost, who would work differently, and how the result would change what it offered customers. By the time it chose software, the tool was the least important part of the decision.

The first company bought software. The second changed how it competed, and the two often buy nearly identical technology. The difference is not the purchase. It is who did the thinking, and at what altitude, before the purchase was made.

The deeper lesson sits underneath even the question of org-chart placement. It is possible to keep the decision nominally at the top, to have the CEO announce it, sponsor it, put their name on it, and still commit the original error, because the part that cannot be delegated is not the signature but the understanding. A leader who declares “we are adopting AI” and then routes the substance to someone else to figure out has not actually owned the decision; they have owned the announcement of it.

Owning the decision means the people at the top personally grasp what this technology can realistically do for their specific business and, just as important, what it cannot: where it creates genuine advantage and where it is merely expensive, what it changes about the work and what it leaves untouched. That understanding is not a technical literacy that can be acquired secondhand from a briefing deck; it is the strategic judgment the rest of the decision depends on, and it is the thing that gets lost the moment “AI” is filed under technology and sent downstairs.

A company whose leaders understand what they are choosing can locate the decision wherever it is most useful and still own it. A company whose leaders have delegated the understanding has lost the decision no matter whose name is on the memo.

The decision the reader makes: Who, in your company, actually owns this decision, not who sponsors it or signs off on it, but who holds it? And have they personally understood what AI can and cannot do for your business, or have they delegated that understanding along with the work, and kept only the announcement?

Chapter 2 — Why Companies Adopt: The Motivation Taxonomy

Once the decision is located at the right level, the next question is the one most leaders skip: why are we doing this at all? The honest answer is rarely a single clean sentence. Most mid-market companies arrive at AI carrying a blur of motives: a competitor's announcement, a board's curiosity, a vendor demo that stuck, a vague sense that the work could be faster. The blur feels like enthusiasm. It is the first failure point, because a motive you have not named is a motive you cannot resource, sequence, or measure. The purpose of this chapter is to convert that blur into a deliberate strategic choice.

There are five reasons a company adopts AI, and they are not interchangeable. Each one implies a different strategy, a different first move, and a different definition of success. Mistaking one for another is how a company spends a year executing well against the wrong goal.

1. **Cost reduction:** Protecting or lowering the cost base, usually by automating work that people currently do by hand. The metric is dollars removed, and the temptation is to reach for headcount, because labor is the largest visible line.
2. **Growth:** Expanding the capacity of the people you already have, so the business can do more without hiring in proportion. The metric is output per person, not cost per unit, a different number, pointing in a different direction.
3. **Quality, or new capability:** Doing the work better than before, or doing what simply was not possible before. This is measured in the quality of the result, not its cost or its volume, and it is where risk reduction and compliance usually belong; they are quality goals wearing a defensive coat.
4. **Defensive positioning:** Deliberately holding competitive ground when a rival's AI is already moving their prices, their turnaround times, or their service levels.
5. **AI-native transformation:** Not improving the existing business but repositioning what the company is, the most demanding driver of all.

The fourth driver deserves a closer look, because it is the one readers most often misread in themselves. Defensive positioning and FOMO produce the same purchase order and feel almost identical from the inside, yet they are opposites. FOMO is a trigger, a reaction to someone else's motion, with no goal and no budget behind it. Defensive positioning is a strategy, a calculated decision that a competitor's capability is genuinely eroding your position and must be answered.

The difference is whether you can state, in a sentence, what ground you are defending and what it is worth. If you cannot, you are not positioning defensively; you are following. *From practice:* the tell is simple. Ask the leader to describe the threat without naming the competitor; if it dissolves the moment the rival's name is removed, the motive was imitation, not strategy. Naming it is the teaching moment of the chapter. FOMO is not shameful; it is the one driver that masquerades as the others.

The cut-versus-grow fork

The first two drivers, cost and growth, look like neighbors. They are the place where two companies, doing the same thing technically, walk away into different futures.

Consider the fork directly. Automating a process creates slack. A company adopting under the cost banner converts that slack into a smaller payroll: it cuts. A company adopting under the growth banner keeps its people and points the freed capacity at more work, better work, or work it could not previously reach: it grows.

On the day of the decision, both companies deployed the same tool against the same task. A year later they are not comparable. The one that cut has a lower cost line and a smaller, less capable organization. The one that grew has the same people, now more productive, and a body of practical AI fluency that did not exist before.

From practice: the second company tends to compound while the first plateaus, the dynamic this paper develops in Chapter 5 as Compounding AI Assets. The point here is narrower and worth sitting with: cost and growth are not two flavors of the same initiative. They are a fork, and the reader is standing at it whether or not they have noticed.

The cost-blindness problem

The other reason naming the driver matters is that every one of them costs more than the reader thinks. AI is not free, and it is not a twenty-dollar subscription. The recurring cost scales with the number of initiatives, and it is dominated by things that never appear in a vendor's pricing page.

In a 2025 study of AI cost governance across 372 companies, roughly 85% of organizations misestimated their AI costs by more than 10%, and nearly a quarter were off by more than 50%, almost always in the same direction: too low. The largest cost drivers were not model usage or tokens, but data platforms and integration: the unglamorous plumbing that makes a tool work inside a real business.[7]

This is why the polished case studies circulating on social media are a poor foundation for a budget. They describe a result, not its true cost, and the result rarely transfers cleanly to a company with different data, different systems, and different people. A driver chosen without a realistic grasp of recurring cost is a driver chosen blind. This is not an argument against adopting; it is an argument for adopting with the second number in hand. The cost-blindness problem is the second of the two guardrails introduced in the next chapter as Velocity with Containment; the first is a clearly defined goal, the work of naming your driver.

---

[7] Benchmarkit and Mavvrik, *2025 State of AI Cost Governance* (372 companies), as reported by CIO.com, October 2025. The study is vendor-sponsored; it is cited here for the direction and scale of the estimation error, which is consistent and large, rather than for a precise figure.

Which is the discipline this chapter asks for. Not three drivers, ranked. One. A company can pursue secondary benefits, but it cannot pursue two primary goals, because the moment a real trade-off appears, and it always does, between cost and capability, between speed and quality, a company with two masters serves neither. The primary driver decides which metric you watch, which process you start with, whether you cut or grow, and what “success” will even mean a year from now. Everything downstream changes with it.

The decision the reader makes: Name your single primary driver. Not three. One. And then ask whether the rest of your plan is built for that driver, or for the blur you started with.

Chapter 3 — The Structural Edge: Velocity with Containment

A mid-market company has one structural advantage over the enterprise it competes against, and most leaders cannot name it. It is not capital, talent, or data; the enterprise has more of all three. It is decision velocity, the simple fact that a company of the mid-market's size can go from a decision to a live pilot in the time it takes an enterprise to schedule the meeting where the decision will be discussed.

Fewer people have to agree. Fewer layers have to sign. The distance between the person who sees the opportunity and the person who can act on it is short enough to cross in a conversation. This is the real edge, and it is the subject of this chapter: both why it exists and how reliably its owners throw it away.

The edge is more than intuition. In a four-year study of 318 firms, the companies whose leaders made strategic decisions fastest went on to grow faster and earn more, and decision speed itself accounted for much of the gap.[8] The firms studied were smaller companies in a high-change industry, which maps onto the mid-market's own situation, and the finding covers only the speed half of this chapter's argument. What it settles is enough: fast strategic decision-making is a measurable advantage in its own right, one that shows up later in growth and profit.

Velocity is an advantage only while it stays bounded. Unbounded, it is just speed in an unknown direction, and speed without a destination is how a company buys a year of motion and ends up nowhere. The discipline this chapter names is **Velocity with Containment**, the idea this document takes its title from: keeping the mid-market's native speed while holding it inside the two guardrails Chapter 2 already built, a clearly defined business goal and a realistic grasp of recurring cost.

The two guardrails are not bureaucracy; they are the only things that distinguish fast from reckless. A company that moves quickly toward a named goal, with the true cost in hand, is using its structural advantage.

A company that moves quickly because a competitor moved, no goal behind it, no budget under it, has converted its one edge into a liability. That failure mode has a familiar shape: the FOMO reflex, software bought because someone else bought it, discovered months later to have answered no question the business was actually asking.

---

[8] J. Robert Baum and Stefan Wally, "Strategic Decision Speed and Firm Performance," *Strategic Management Journal* 24, no. 11 (2003): 1107–1129. The study tracked 318 firms over four years and found that strategic decision speed predicts subsequent growth and profit and mediates the influence of environment and structure on performance.

The Imitation Trap

The more dangerous threat to the edge is subtler than FOMO, because it looks like maturity. A mid-market leader, wanting to "do AI properly," reaches for the enterprise's apparatus: the steering committee, the formal RFP, the eighteen-month roadmap with phases and gates, the governance board that meets monthly to review progress against the plan. Each of these feels responsible. Each is, in an enterprise, necessary. And each one, imported into a mid-market company, dismantles the only advantage that company had. This is the **Imitation Trap**: copying the coordination machinery of organizations a hundred times your size, and importing their slowness along with it, in the belief that you are importing their rigor.

The trap is seductive because the enterprise's machinery is not foolish. A ten-thousand-person company needs a steering committee, because without one, ten thousand people cannot stay aligned. The machinery is the enterprise's solution to a problem the mid-market does not have. When a four-hundred-person company adopts it, it pays the full cost of the solution while having none of the problem it solves, and it loses the speed that was its answer to that same problem all along. The leader who installs an eighteen-month AI roadmap has not become more serious. They have voluntarily adopted the enterprise's reaction time without the enterprise's resources to absorb it.

Why the edge exists: the economics underneath

It is worth understanding why the speed advantage is real, because the explanation also reveals exactly what the Imitation Trap destroys. The cleanest account comes from transaction cost economics, the line of thought that begins with Ronald Coase's question of why firms exist at all and that Oliver Williamson developed into a theory of how organizations choose their boundaries.[9]

The core insight is that every act of coordination inside a firm carries a cost, the cost of communicating, negotiating, aligning, approving, and monitoring, and these internal costs rise sharply with size and structure. A decision that one person can make alone is nearly free to coordinate. The same decision routed through six functions, three committees, and a sign-off chain is expensive long before anyone has spent a dollar on technology.

Seen this way, the mid-market's decision velocity is not a cultural trait or a matter of being "scrappy." It is an outcome of structurally lower internal coordination costs. The speed is the visible surface of an invisible economic fact: fewer parties must transact for a decision to become an action.

And this is what makes the Imitation Trap a genuine error rather than merely a missed opportunity. When a mid-market company adopts enterprise governance, it is not adding safety on top of speed. It is manufacturing the internal transaction costs it was previously free of, re-importing, by choice, the exact coordination overhead that makes

[9] Ronald Coase, "The Nature of the Firm" (1937); Oliver Williamson's subsequent development of transaction cost economics. The theory is invoked here for a single, specific claim: that coordination inside a firm is itself costly and rises with structure.

the enterprise slow. The roadmap and the steering committee do not sit alongside the advantage; they consume it. The firm pays enterprise coordination costs to reach an enterprise's reaction time, and calls the result discipline.

The figure, in its place

There is a number behind all this, and its place is here, as support, not as the headline. In MIT's 2025 study of AI in business, the top-performing mid-market firms moved from pilot to implementation in roughly ninety days, against nine months or more for large enterprises.[10] The figure is useful for scale: the gap is not ten or twenty percent, it is a different order of speed.

But the figure is not the argument. That the edge exists is unsurprising; everyone senses that smaller firms move faster. The claim worth holding is the counterintuitive one: that the edge is routinely forfeited, not by competitors and not by technology, but by the firm's own decision to imitate. The number tells you what is at stake. The Imitation Trap tells you how it is lost.

Where the edge holds, and where it fades

Honesty about the advantage requires admitting it is not uniform across the mid-market. The edge is strongest at the lower and core ends of the range, companies from roughly $10M to $500M in revenue, where the organization is still small enough that coordination genuinely is cheap and a decision really can cross the firm in a conversation.[11]

At the upper end, from $500M toward $1B, the symptoms of enterprise begin to appear on their own: more layers, more functions with standing to object, longer natural distances between decision and execution. This is the onset zone of the Imitation Trap, and the reason it is dangerous there is that imitation feels justified: the company is large enough that some enterprise machinery looks proportionate.

The discipline at the upper end is to add only the coordination the size actually demands, and not a step more, because every step beyond necessity is a voluntary surrender of the very thing that distinguishes the company from the larger rivals it is starting to resemble.

---

[10] MIT NANDA, *The GenAI Divide: State of AI in Business 2025* (July 2025). A methodological note is required: NANDA defines "enterprise" as companies above $100M in revenue, which overlaps with the upper portion of the mid-market as defined here (the National Center for the Middle Market's $10M–$1B range). The ninety-day-versus-nine-month contrast should therefore be read as smaller, faster firms against the largest enterprises, not as a clean line drawn at the mid-market's own upper boundary. The study's base is modest, roughly 300 deployments, 52 interviews, and 153 surveys, so the figure is best used for the order of magnitude of the gap, not for precision.

[11] Segment definitions follow the National Center for the Middle Market: Lower ($10M–$50M), Core ($50M–$500M), Upper ($500M–$1B).

The practical form of all this is a question about distance. Velocity with Containment is not a posture or a value; it is measurable in the number of human approvals that stand between a decision and its execution. Each approval is a transaction cost, and each one is either earned, a guardrail protecting the goal or the budget, or imported, a habit borrowed from an organization that needed it and you do not.

In practice, Velocity with Containment looks like this: a decision made by two people in a single meeting, a pilot scoped to one process and one team, a budget ceiling set before the vendor is contacted, and a defined date, thirty or sixty days out, at which the result is assessed against the goal named at the start. The guardrails are not a committee; they are a sentence describing the goal and a number describing the budget. Everything else is speed.

The decision the reader makes: How many approval steps stand between the people deciding and the people executing, and how many of those steps are protecting the business, rather than imitating one ten times your size?

Chapter 4 — People Before Pipelines: Adoption Is a Motivation Problem

Underneath who owns the decision, what drives it, and how fast the firm can move sits a layer most documents on AI adoption skip on their way to architecture and governance. It holds the simplest fact in the entire subject and the most ignored: a tool that no one wants to use returns nothing.

The model can be capable, the integration clean, the business case sound, and the initiative will still produce nothing if the people meant to work with it quietly decline. Technology is rarely the bottleneck in mid-market AI adoption. Motivated, capable, involved people are. Their presence or absence separates an initiative that compounds from one abandoned six weeks after launch.

It also decides whether the previous chapter's advantage amounts to anything. Velocity with Containment describes how fast a firm can move from a decision to a live system; the people in this chapter are who carry it there. Without a workforce that wants to use the tools, that speed produces a signed decision and no working system behind it.

The reason the human layer is decisive is that AI adoption, unlike most software, cannot be installed. A new ERP system can be mandated; people will use it because the old way is switched off and there is no alternative path to doing their job.

AI is different in kind. Its value comes not from compliance but from initiative: an employee who reaches for the tool, finds the unobvious use, brings their own judgment to where it helps and where it does not. That behavior cannot be ordered. It can only be motivated. Which means the entire question of adoption resolves into a question of motivation: what would make a capable person want to make this work?

The two true drivers

From practice: two things reliably predict whether a person becomes an early and genuine adopter, and neither is on most rollout plans. The first is curiosity, the disposition to poke at a new tool, to wonder what it can do, to try the strange use case nobody sanctioned. The second is intrinsic motivation, the drive to do one's own work better for reasons that belong to the person, not to the org chart.

Where these two are present, adoption tends to take care of itself; the person finds uses the rollout never imagined. Where they are absent, no amount of mandated training fills the gap, because the behavior that makes AI valuable is the behavior that cannot be compelled.

The strategic implication is uncomfortable for a planning culture: the most important input to adoption is a human disposition, and the leader's job is less to install a tool than to create the conditions under which curiosity and intrinsic motivation are allowed to operate.

Research on AI adoption points the same way. When researchers test what predicts whether people use an AI tool, the strongest signals are how useful the person believes it to be and what they call their AI mindset: the belief that the tool will extend their abilities rather than hollow them out.[12] A tool pushed out by mandate, with no work done on whether people believe it will help them, stalls no matter how capable it is. What moves usage is the person's conviction that the tool makes their own work better.

Why mandates fail: the science underneath

This is not merely a soft observation; there is a well-developed science that explains exactly why control-based adoption underperforms. Self-determination theory, developed over several decades by the psychologists Edward Deci and Richard Ryan, distinguishes motivation that is autonomous, arising from the person's own interest and values, from motivation that is controlled, driven by external reward, pressure, or mandate.[13]

Its central, repeatedly demonstrated finding is that autonomous motivation produces better engagement, persistence, and performance on complex tasks, while controlled motivation tends to produce compliance at best and resistance at worst. The theory identifies three psychological needs whose satisfaction is what makes motivation autonomous: autonomy, a sense of acting by one's own choice; competence, a sense of growing mastery rather than exposed inadequacy; and relatedness, a sense of doing this alongside others, not against them.

Read through this lens, the standard AI rollout is almost perfectly designed to fail. A top-down mandate to "use the new AI tools" attacks autonomy: it converts a potential interest into an order. A demo that shows how much faster the machine performs the employee's own work attacks competence: it lands as a threat, not an invitation. And an initiative imposed by leadership on a workforce that had no hand in it attacks relatedness: it sets the technology against the people rather than among them.

The same theory points directly at the remedy, and it reframes the most expensive line item in most adoption plans. The real purpose of training is not to instruct. People can learn a tool's features from documentation.

The purpose of training is to motivate: to satisfy the three needs the mandate violates: to give people room to find their own uses, to build genuine competence so the tool feels like growth rather than exposure, and to make adoption a shared undertaking

---

[12] Fabio Ibrahim, Johann-Christoph Münscher, Monika Daseking, and Nils Torge Telle, "The Technology Acceptance Model and Adopter Type Analysis in the Context of Artificial Intelligence," *Frontiers in Artificial Intelligence* 7 (2025): 1496518. In a sample of 1,007 respondents, perceived usefulness and an "AI mindset" were the strongest predictors of AI use. Cited for that empirical finding.

[13] Edward L. Deci and Richard M. Ryan, self-determination theory. The theory is invoked here for a specific, load-bearing claim, that controlled, mandate-driven motivation underperforms autonomous motivation on complex work, and that autonomy, competence, and relatedness are the conditions of the latter.

rather than something done to them. Training that only teaches features answers a question nobody was afraid of, while leaving the real one untouched.

Training as motivation, and AI-native people as the goal

Reframed this way, the aim of training is threefold: to motivate, to demonstrate a real and personal advantage, and to remove the sense of being left behind. The demonstrated advantage matters because abstract benefit persuades no one: a person adopts when they have felt the tool make their own day better, not when they have been told it improves a corporate metric.

The sense of being left behind matters because it is the quiet condition under which capable people disengage, and it is dissolved not by reassurance but by competence: the person who has genuinely learned to work with AI no longer fears it.

The deeper goal underneath all three is to produce AI-native people, individuals fluent enough that working with these systems is ordinary rather than exceptional. This is the precondition the rest of the strategy rests on. There is no such thing as an AI-native company whose people are not themselves AI-native; the company-level transformation is the aggregate of the personal one, and it cannot be bought ahead of it.

The second purpose of training: conscious use

Motivation gets people to reach for the tool; it does not make their reaching produce good work. A person who wants to use AI but has no judgment about where it helps will run everything through it, including the tasks it handles worst, and mistake the volume for progress. So training carries a second job alongside motivation, and that job serves the same competence need the theory above already named: the judgment to know when AI is the right instrument, when it is the wrong one, and how to check what it returns. This is competence, not control, and the difference decides whether the remedy works.

The failure mode is the mirror image of the mandate. A firm hands out broad access, hopes people will get used to it, sets no boundaries, and gets thoughtless use everywhere. *From practice:* a large company opened AI to everyone with no guardrails and watched staff route work through it where it added nothing. Two costs follow. One is money, as unbounded use burns spend on tasks that never needed it, which is the behavioral source of the cost blindness Chapters 2 and 3 already counted at the level of the budget. The other surfaces more slowly: people generate documents from AI and then generate the next round from those, until the origin of any claim has receded a layer at a time and no one can trace or verify it. Chapter 7 returns to why that compounds.

The instinct to fix this with a rulebook fails for the same reason the mandate failed. Forbid a tool people find useful and they do not stop using it; they move it where no one can see. A ban relocates the behavior and surrenders the oversight it was meant to protect, and the autonomy it violates and the visibility it loses are one event. That is why shadow use returns in the security chapter, where it is the exposure no policy ever recorded.

How far to formalize depends on the field. A compliance-heavy business already runs strict rules for other reasons, and AI slots into that existing structure. A lightly governed business that bolts prohibition onto a new tool pays twice, in motivation and in shadow use, for control it never gains. The lighter the external regime, the more the discipline has to live in the judgment of informed people rather than in a document, and informed means more than a nod in a session: a person uses AI well once they understand how it moves the business and why unchecked, self-referential output cannot be audited later. No rule substitutes for that understanding, which is why training, not policy, carries the weight.

Involvement is not consultation

The remedy for the relatedness problem has a concrete form: the people who will work with a system must participate in shaping it. This is not the same as a survey or a town hall, and it does not mean handing strategic decisions to the front line.

It means that the employees who will live with an AI system are well-informed participants in how it is chosen and shaped: consulted on where it helps, listened to on where it gets in the way, given real influence over the working design even if not over the strategy.

Involvement of this kind does two things at once: it produces a better-fitted system, because the people who do the work know the work, and it produces ownership, because a person defends what they helped build and resists what was done to them. The cost of skipping it is not only a worse tool but an unmotivated workforce, and the second cost is the larger one. Involvement does not require slowing the decision down; it requires starting it earlier: bringing the relevant people into the design phase rather than the announcement.

The CTO Ownership Trap

There is a particular way that leaders, having half-understood that AI is “technical,” hand the entire human problem to the wrong part of the organization. It begins reasonably: AI involves models and integration and data, which are the technical function’s domain, so the initiative is handed to the CTO or the head of IT to “implement,” and leadership steps back to await delivery. In a mid-market company this fails, and it fails for a structural reason.

The CTO Ownership Trap is the error of treating adoption as a system migration to be executed by the technical function, when adoption is in fact a change to the character of the work itself, and therefore belongs to whoever owns business strategy. The trap is not the involvement of the CTO, who is an essential partner in execution.

The trap is the act of delegation: a CEO who announces “we are adopting AI, go build it” has already lost, because the people making the decision must themselves understand what AI can realistically do, and the people who will work with it must be involved in shaping it, and neither of those is something a technical function can supply on the leadership’s behalf. Ownership of the decision stays with the business. The CTO is a partner in how, never the owner of whether and why.

The Trainer's Dilemma

The deepest obstacle to adoption is one that feature demos cannot touch, because it is not about features. AI differs from ordinary software in that it does not replace a tool the employee uses; it augments, and sometimes absorbs, the employee's own judgment, and even where a role is automated, a human orchestrator typically remains to direct the system. This creates a wholly rational fear in the people asked to adopt it: that in transferring their hard-won experience into an agentic system, they are training their own replacement.

This is **the Trainer's Dilemma**, and the first thing to understand about it is that the fear is not irrational and cannot be motivated away with optimism. A person who senses that diligent cooperation might erode their own standing is reading the situation accurately, and they know it. Pretending otherwise insults them and confirms their suspicion.

The rationality has research behind it. A 2025 study across six European countries found that workers grow more afraid of automation when they read the technology as substituting for their tasks rather than supporting them, and that a person's sense of control over their own situation lowers that fear.[14] The finding lines up with the dilemma: resistance tracks whether people expect to be replaced or strengthened, which makes it a rational response to how leadership designs the adoption.

The broad anxiety is well documented. In a Pew Research Center survey of US workers in early 2025, fifty-two percent reported being worried about AI's impact on the workplace, and thirty-two percent expected it to lead to fewer job opportunities for them personally, against only six percent who expected more.[15]

The most telling detail in that data is that the workers most exposed to AI were the most concerned: those who already used AI in their jobs were more likely than non-users to expect fewer opportunities, by forty-two percent to thirty percent. Familiarity, in other words, did not breed comfort. It sharpened the perception.

The specific mechanism named here, the sense of training one's replacement, is the author's framing rather than a finding in the survey; what the survey establishes is that the underlying unease is real, widespread, and concentrated precisely among the people whose cooperation adoption most depends on. (Pew, Feb 2025; the "training my replacement" mechanism is the author's observation.)

---

[14] Renata Włoch, Katarzyna Śledziewska, and Satia Rożynek, "Who's Afraid of Automation? Examining Determinants of Fear of Automation in Six European Countries," *Technology in Society* 81 (2025): 102782. The study found that perceived task substitution raises fear of automation and that a greater sense of personal control reduces it. It documents the general anxiety; the "training my replacement" mechanism remains the author's framing.

[15] Pew Research Center, survey of US workers, February 2025 (n = 5,273). The "training my replacement" mechanism described in this chapter is the author's framing, not a Pew finding; the survey is cited only for the broad anxiety and for the users-versus-non-users nuance.

Naming the dilemma is not a counsel of despair; it is the entry to the only durable answer. An adoption designed so that capable people become more valuable as they engage with AI, more productive, more skilled, more able to direct the systems rather than be displaced by them, dissolves the dilemma at its source, because the rational calculation now runs the other way.

The point is not to suppress the fear or to deny it, but to make it untrue, by building an adoption in which the people who lean in are the people who come out ahead. That is a design choice the leadership makes, or fails to make, and the workforce can tell which one was made.

The decision the reader makes: Does your adoption plan make your best people more valuable, or does it quietly ask them to train their own replacement? They can tell the difference, and they will act on it.

Chapter 5 — Readiness and Sequencing

This chapter is where the decision becomes a move. A leader who has named a single primary driver, understood the speed edge and the guardrails around it, and grasped that adoption stands or falls on motivated people now faces a different kind of question: not whether and not why, but what first, and on what footing.

The temptation at this stage is to treat the first move as a procurement event: choose a tool, buy seats, switch it on. That instinct is the same category error the whole document has been working against, arriving one level down. The first move is not a purchase. It is the opening position of a sequence, and the quality of the sequence is decided before anything is deployed: in how readiness is assessed, which process is chosen, and how the capability is sourced and run.

Readiness is joint, or it is not readiness

Most readiness assessments examine one half of the problem. They ask whether the data is clean, accessible, and structured enough for a model to use, a real and necessary question, and they stop there, as if the only barrier were technical. For a mid-market company, data readiness does not require a data lake or a dedicated data engineering team; it requires that the records the process depends on exist in a consistent, machine-readable format, are accessible without manual extraction, and are clean enough that a human would trust them to make a decision.

But there is a second readiness, and it is the one Chapter 4 was about: whether the organization is prepared, whether the people who will work with the system are motivated and involved, whether the competence to use it exists or can be built. Technical readiness and social readiness have to be assessed together, not in sequence.

Treating them sequentially, getting the data right first and the people later, produces the familiar failure in which a technically flawless system meets a workforce that never wanted it, and returns nothing. The two readinesses are not stages; they are two axes of the same judgment, and a first move should be chosen where both are plausibly high, not where one is high and the other has been deferred.

That mid-market organizations are not, by default, ready on either axis is well established for this segment. In RSM's 2025 survey of middle-market decision-makers, only fifty-three percent felt even "somewhat prepared" for AI adoption and ten percent felt not prepared at all; sixty-two percent found adoption harder than they had expected, and seventy percent needed outside help to do it.[16] The rollout difficulties they reported

[16] RSM US 2025 Middle Market AI Survey (RSM US with Big Village; 966 decision-makers across the US and Canada, February–March 2025). Figures cited: 53% felt "somewhat prepared" and 10% "not prepared"; 62% found adoption harder than expected; 70% needed outside help; reported rollout challenges clustered around data quality, privacy/security, and skills.

clustered around data quality, privacy and security, and skills, which is to say around exactly the two readinesses, technical and human, in roughly equal measure.

The lesson is not that the mid-market is incapable; it is that readiness is the normal gap, not the exception, and that a leader who assumes the firm is ready because the strategy is sound has skipped the one diagnosis that most determines whether the first move lands.

Don’t go in head-on

From practice: the single most reliable way to waste a first move is to deploy before either readiness is in place: to go in head-on, mandating a tool into a workforce that has not been prepared and onto data that has not been cleaned, on the theory that momentum will sort out the details. It does not.

A premature deployment burns the two things hardest to recover: the goodwill of the people, who learn from a bad first experience that AI is a nuisance imposed on them, and the credibility of the initiative, which now has to overcome its own track record before it can do anything else.

Readiness precedes deployment, not as a bureaucratic gate but as a matter of sequence, because the cost of going early is paid in the human capital the rest of the strategy depends on. The discipline here is the same Velocity with Containment named in Chapter 3: speed is the advantage, but speed deployed before readiness is not velocity, it is recklessness wearing velocity’s clothes.

Choose the process before the tool

Given readiness, the next question is where to point the first move, and here there is a criterion sharp enough to use directly. Some processes are settled: well-defined, repeatable, governed by stable rules, performed the same way each time. Others demand fresh creative judgment on every instance, where the right answer depends on context that changes and cannot be reduced to a rule. The settled, repeatable processes are the ones that can be largely automated; the judgment-heavy ones resist it. The first move should target the former.

This is not because the creative work is unimportant, but because the settled work is where current systems are strong, where the data to support them already exists as a byproduct of doing the work, and where success is legible enough to build confidence for the harder moves later. Choosing a judgment-saturated process for the opening play inverts the odds: it is the hardest thing these systems do, attempted before the organization has learned to work with them at all.

The process-selection criterion is, in effect, a way of choosing a first move the firm can win, and a visible early win is worth more to adoption than a larger but contested one.

Sourcing is a strategic choice, not a technical one

Once the process is chosen, the capability behind it has to come from somewhere, and the options are genuinely different strategic commitments rather than a single technical

decision to be delegated. A firm can buy a finished capability off the shelf, partner with a vendor or integrator who supplies it as a service, train its own people to build and run the capability internally, or build the system itself and own it outright.

These are not points on a convenience scale; they are different bets about where the firm wants its capability to live. Buying is fastest and shallowest: the capability is rented and leaves when the contract does. Partnering imports expertise the firm lacks but creates a dependency that has to be managed. Training builds durable internal capability but takes time and commitment before it pays. Building gives the most control and the deepest ownership at the highest cost and risk.

The error is to treat this as the CTO's call about implementation. It is a leadership decision about what kind of company this is becoming, whether AI capability is something the firm consumes or something it owns, and it has to be made with the primary driver from Chapter 2 firmly in view, because each sourcing path serves some drivers far better than others. Buying and partnering suit cost reduction and defensive positioning, where speed to value matters more than ownership; training and building suit growth and AI-native transformation, where the capability itself is the point and is worth the time it takes to own.

The deployment-model decision and the cost it hides

Sourcing has a close cousin that deserves its own attention, because it is where the cost-blindness problem of Chapter 2 most often does its damage. The deployment model, how the AI actually runs in the business, comes in several forms with sharply different cost and headcount consequences.

The lightest is an off-the-shelf assistant layered onto the stack the firm already runs, such as a copilot inside an existing productivity suite: minimal integration, fast to switch on, priced per seat. Heavier is an agent delivered as a service, where a vendor runs the system and the firm consumes its output. Heaviest is a custom-built, self-run agent the firm operates on its own infrastructure.

The instinct is to read these as a simple ladder from cheap to expensive, but the cost structure is not that simple, and the place it bites is scale. Per-seat subscriptions are cheap for a pilot and can become punishing once every employee has a seat; at that point the math can tip toward building, because owning the system looks cheaper than renting it for hundreds of people. What that comparison routinely omits, and what Chapter 2's cost-blindness finding predicts it will omit, is that a self-run system carries its own recurring weight: infrastructure cost, integration and data-platform cost, and people whose job is to maintain, monitor, and babysit the agents so they keep working.

Recall the verified pattern from Chapter 2: roughly eighty-five percent of organizations misestimate their AI costs by more than ten percent, nearly a quarter by more than half, almost always too low. The dominant drivers were data platforms and integration, not model usage.[17] The deployment model is precisely where that error gets made, because

[17] Benchmarkit and Mavvrik, "2025 State of AI Cost Governance" (372 companies), as reported by CIO.com, October 2025. The study is vendor-sponsored and is cited here

the seat price is visible and the run cost is not. The decision cannot be made on the sticker; it has to be made on the fully loaded monthly cost at the scale the firm intends to reach.

The instinct that a custom build is cheaper at scale hides a second assumption beneath the first: that building the system is the work. Building it is only the visible part. A self-run agent is constructed once but operated and improved without end, kept running as models, prices, and dependencies shift beneath it, corrected when it drifts, re-evaluated every time the technology it rests on moves on. *From practice:* the firms that come to regret building are rarely the ones that couldn't build it; they are the ones that costed the build and forgot the operating model, and found they had not finished a project but acquired a service they now had to staff. The distinction is not pedantic: a project has an end, and an operating model has a payroll. Reading a custom build as a one-time capital cost rather than a standing operational commitment is the same cost-blindness from Chapter 2 in a new disguise, and it is the first appearance of the idea Chapter 7 makes central, that adoption is never delivered and done; it is run.

Compounding AI Assets

Underneath the choice of process, sourcing, and deployment model sits a single test that orders all of them, and it is the framework this chapter is built to deliver. Adoption capital can be spent or it can be invested. It is spent when the first move is a one-off tool bolted on to trim a single cost line: a step-change that happens once and yields nothing further; the line is lower, and there the story ends.

It is invested when the move builds capability that accumulates: employees who each learn to build and orchestrate their own agents, process knowledge captured and reusable along the way, a workforce that is measurably more capable after the move than before it. This is **Compounding AI Assets**: the recognition that invested capability compounds while spent capability does not.

AI-native people are the precondition for an AI-native company, and every capable person lowers the cost of the next capability, so that an organization which invests rather than spends finds each successive move cheaper and faster than the last. The test applies to every decision in this chapter and is blunt enough to ask out loud of any proposed first move: does it compound, or does it merely spend? A move that only spends is not wrong; sometimes a cost line needs trimming. But a leader should know which one they are choosing, because a sequence of moves that only ever spend never builds the advantage the firm is supposedly adopting AI to gain.

The reason to assess social readiness jointly with technical readiness, to choose a process the firm can win, and to weigh sourcing as a question of owned-versus-rented capability is that all three are really the same question asked from different angles: is this first move building something that will make the next move easier, or is it a transaction that ends when it ends? The compounding test is what keeps the sequence a sequence rather than a series of disconnected purchases.

---

via CIO.com. Carried forward from Chapter 2, where this finding is introduced in full.

A practical note on measurement during the early phase: before financial ROI is visible, track leading indicators instead. The most reliable ones are behavioural: the percentage of the target team using the tool at least weekly, the number of use cases identified by employees rather than management, and the reduction in time spent on the specific task the system was deployed against. These are not proxies for success; they are the conditions under which financial return becomes possible, and they are the right questions to ask at thirty and sixty days.

The decision the reader makes: Which deployment model, and have you costed it monthly, at scale, including the people required to run it? And before that number is even worth computing: does this first move compound, or does it merely spend?

Chapter 6 — Security Is Not the Old Security

The first move carries one question that is almost always asked too late, if it is asked at all: whether the thing being deployed is safe in a sense the firm has never had to think about before. Choosing the process, the sourcing path, and the deployment model settles how the system will work. It does not settle whether the system can be trusted with the firm's data and its customers'.

Most mid-market AI documents either omit security entirely or fold it into a governance chapter near the end, as a box to be checked once the real decisions are made.

That placement is itself the error. Security is not a late-stage compliance step; it is a property of the first move, designed in or absent from the start, and it is the second place, alongside starting before governance, where this document parts company with the standard advice. The reason is simple and uncomfortable: the security a mid-market firm already has was built for a different kind of system, and it does not extend to this one.

The category shift

Classical IT security was built to defend against code. Viruses, malware, intrusion, the unauthorized program slipping past the perimeter: the threat was a hostile instruction set, and the defenses answered it in kind with firewalls, signatures, access controls, and patches. That entire apparatus assumes the attacker speaks in machine language and the defense can be machine-checked.

An AI system breaks the assumption, because it communicates in natural language and can be manipulated through that language by anyone who can type a sentence. The attack surface is no longer only the code; it is the conversation. A model that reads text to do its job can be made to misbehave by text that was written to make it misbehave, and there is no signature for a persuasive paragraph. This is why existing IT security does not necessarily cover the new exposure. The IT team is not deficient; the category of threat is different from the one their tools were designed to catch.

The first thing leadership has to do is not assume coverage. The right posture is to ask, explicitly and early, whether the protections already in place reach the language layer at all, and to treat a confident "of course we're covered" with the same suspicion as any other unexamined assumption.

The new attack surface

The new failure modes are not exotic edge cases; they are the normal ways a language-driven system gets abused, and they already have names and established taxonomies in the security community.[18] The first is prompt injection: instructions

[18] Established, publicly maintained taxonomies for these attack categories exist, notably the security community's catalogues of large-language-model risks (e.g., OWASP) and the relevant federal AI risk guidance (NIST). They are referenced here at the category

smuggled into the text a model reads, telling it to ignore its real task and do something else instead. It is best understood as the evolved analog of phishing: where phishing manipulates a person through a crafted message, prompt injection manipulates the system through one, and the same data the model is meant to summarize or act on can carry the payload.

The second is the jailbreak by pretext: a conversation constructed to talk the model out of its own guardrails, by framing, role-play, or incremental steps that each seem reasonable until the boundary has been crossed. The third is data exfiltration through the model itself: using the system as the channel by which information it can reach is drawn out and exposed, not by breaking into a database but by asking the model, in the right way, for what it already has access to.

None of these requires breaching the network in the old sense. They are abuses of the system working as designed, through the very interface that makes it useful, and a firm that has never catalogued them has no way of knowing whether it is exposed.

The naive-trust problem

The exposure is widened by a second-order effect: tools and agents are multiplying far faster than anyone's ability to scrutinize them. A capable employee can adopt a new AI product in an afternoon, and the marketplace produces them faster than that. The naive assumption, and it is naive, is that any product polished enough to look professional, and plausible enough to fit the workflow, must also ship with adequate protection built in. There is no basis for that belief.

A vendor's security is a question to be verified, not a default to be assumed, and the gap between "this tool clearly works" and "this tool is safe to give our data to" is exactly the gap that gets skipped when adoption moves quickly and informally. This is not an argument for slowing down; it is an argument for knowing what to ask. The discipline is to treat every new tool and agent as something that has to earn trust on the security question specifically, separate from whether it does its job well, because doing the job well and handling data safely are independent properties, and the visible one is no evidence of the hidden one.

The complement to vendor scrutiny is a short, plain-language acceptable use policy, not a legal document but a one-page statement of what employees may and may not feed into AI tools, because the most common data exposure in a mid-market firm is not a sophisticated attack but an employee pasting sensitive client information into a public model without realising the implications.

From practice: within weeks of ChatGPT's release, and long before any vendor offered an enterprise version with real controls, people were using it for real work in companies that had no policy on it at all. Few thought hard about what they pasted in. The eventual response, firm after firm, was to ban it, and the ban came too late to matter. People were already used to the tool, so the use moved out of sight: a company name stripped

---

level; specific control names and version numbers are deliberately omitted rather than approximated.

from the text, that stripping mistaken for safety, the same material going into the same public model with no one watching. The useful question was never whether people would use it. It was whether they used it with judgment. A ban produces concealment instead, which is why prohibition without a workable alternative does not close this exposure; it hides it.

The data decisions that belong to leadership

Underneath the technical vectors sit a handful of choices that look technical but are not, and that leadership cannot delegate to IT alone without abdicating something strategic. The first: does the model train on the firm's data, and is that acceptable? Feeding proprietary information into a system that learns from it can mean surrendering a piece of the firm's advantage in exchange for convenience, and whether that trade is acceptable is a business judgment about competitive position, not a configuration setting.

The second: where does the AI actually run, on a third-party cloud the firm does not control, or on infrastructure it does? That choice determines who else can reach the data in transit and at rest, and it carries cost and capability consequences that connect directly to the deployment-model decision of Chapter 5.

The third, and the one that grows most pressing as agents take on real work: what data and what identity does an agent actually have access to? An agent granted broad permissions to be useful is also an agent that can be manipulated into using those permissions against the firm; the scope of its access is the ceiling on the damage a successful injection can do.

Each of these is a strategic exposure dressed as a technical detail. Leadership does not have to make the configuration, but it does have to own the decision, because the cost of getting it wrong lands on the business, not on the IT function that quietly chose a default.

Security as the hard edge of Containment

These threads pull together into a single principle, and it is the same one the document has been building toward from Chapter 3. Velocity with Containment named two guardrails on the firm's speed: a clearly defined goal and a realistic grasp of cost. Security is the hard edge of that same Containment: the boundary that lets the firm keep moving fast without the speed turning into exposure.

Built in from the first move, security is what makes velocity safe to exercise; bolted on after the fact, it becomes the thing that forces the firm to slow down, retrofit, and unwind decisions it has already shipped. The mid-market's structural advantage is the ability to move quickly from decision to live system, and that advantage is only worth having if the systems it produces can be trusted with the firm's data and its customers'.

Containment is not the brake on velocity. It is the condition that allows velocity to be used at all. Security, designed in early, is where that condition is met or missed.

That is why this belongs at the first move and not at the end. A security posture decided after deployment is deciding it too late; the exposures were created the moment the system went live with its access, its data, and its place in the workflow.

The firms that handle this well are not the ones with the largest security budgets but the ones that asked the language-layer question before they deployed, verified rather than assumed their vendors, and scoped their agents' access deliberately. None of that requires enterprise machinery. It requires leadership to recognize that the security it already owns was built for a different threat, and to check, early, explicitly, and without taking the comfortable answer, whether it reaches the new one.

The decision the reader makes: Does your current IT security cover language-layer attacks, such as prompt injection, jailbreak, and exfiltration through the model itself? Yes or no. And if you do not know the answer, the answer is no.

Chapter 7 — Adoption Is a Relationship, Not a Project

Every chapter so far has been about getting the first move right: framing the decision as strategy, naming the motive, protecting the speed advantage, putting people before pipelines, sequencing the readiness, and building security in from the start. All of it points at a single moment: the system goes live. And almost every mid-market AI document treats that moment as the finish line. The project is delivered, the tool is in production, the initiative is closed, and attention moves on.

That instinct is the last and most expensive category error in the sequence. A live AI system is not a delivered artifact; it is the beginning of a relationship the firm now has to manage for as long as it intends to get value from it. The project framing has a natural end. The relationship framing does not. The firms that keep getting value are the ones that understood, from the start, that there was never going to be a point at which they were done.

The False Negative Window

Before the relationship can be managed well, it has to survive its own early days, and the way most firms measure those early days quietly kills good initiatives. An honest AI initiative, one aimed at a real change in how work gets done rather than a cosmetic win, tends to show little measurable return at first. The data has to be prepared, the people have to climb the curve, the process has to be reshaped around the new capability, and none of that lands on a quarterly P&L on schedule.

So when a firm judges the initiative on a ninety-day financial snapshot, it is very likely to be looking through a window in which even a genuinely good initiative looks like a failure. *From practice:* the initiatives that get cancelled at the three-month mark are frequently the honest ones, because the honest ones front-load the cost and back-load the value, while the cosmetic ones do the reverse and photograph better early.

The discipline is to judge on time-to-value logic rather than a snapshot: to ask whether the initiative is on the trajectory it was expected to be on, not whether it has already paid for itself. This is the **False Negative Window**: the early stretch in which the measurement instrument is too blunt to tell a slow-maturing success from a real failure, and in which the firm is most likely to throw away something that was working.

Economists have a name for the underlying pattern: the Productivity J-Curve. Brynjolfsson, Rock, and Syverson showed that when a general-purpose technology arrives, measured productivity falls at first, because the firm pours effort into things the books never capture, such as new processes, new skills, and reorganized work, and the payoff appears only once those intangible assets are built.[19] AI follows the same shape.

---

[19] Erik Brynjolfsson, Daniel Rock, and Chad Syverson, “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies,” *American Economic Journal: Macroeconomics* 13, no. 1 (2021): 333–372. The J-Curve explains the economic logic beneath the False Negative Window; the two are complementary, and the False

The early fall is the firm building the capability, and the return arrives once that capability begins to carry the work.

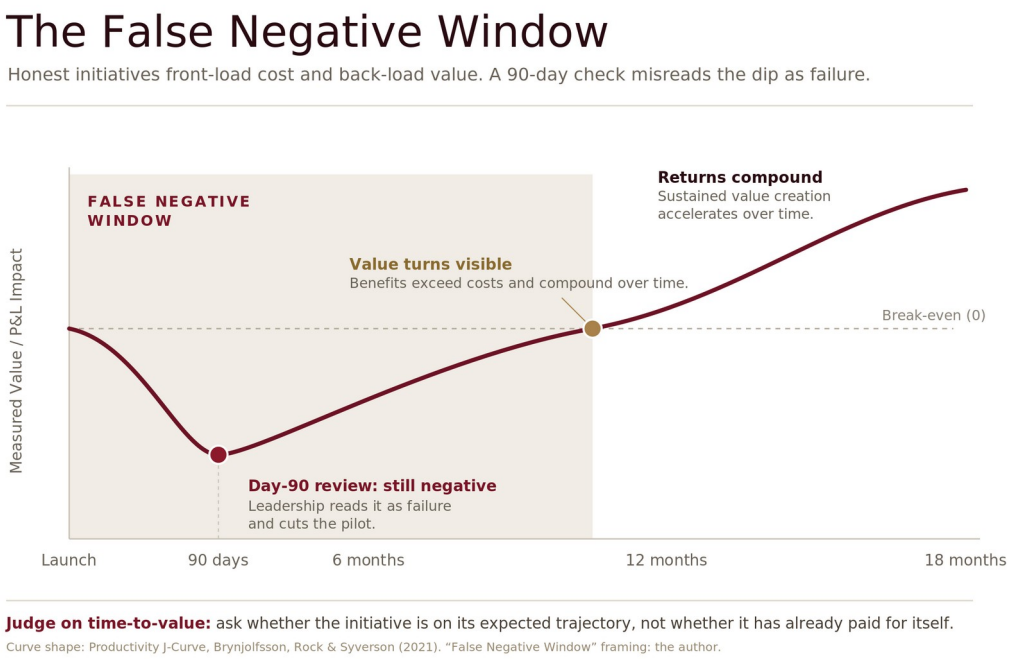

Figure 1. The False Negative Window. Measured value dips after launch while the firm builds capability, sits below the line at the ninety-day review, where good initiatives get cancelled, then crosses back above it as the capability begins to carry the work.

From practice: a mid-market healthcare provider put an AI tool into its patient-intake process, expecting lower administrative headcount inside the first quarter. At the ninety-day review the savings had not arrived: staff were still learning where to trust the system and where to correct it. Leadership read the initiative as a failure and ended the contract. The capability was weeks from compounding. They cut the organizational learning that had almost finished forming, mistaking a slow start for a permanent flaw.

Naming the window is most of the defense against it. A leader who knows the early read is unreliable will ask a different question at ninety days, not "what has it returned" but "is it where we expected it to be on the way to returning," and that single substitution saves more good initiatives than any measurement framework.

Adoption is not set-and-forget

Past the window, the deeper mistake is to imagine that a system which is working will keep working on its own. The most useful mental model for what a firm has acquired is not a piece of software but a person it has hired. You onboard a new hire, you train

---

Negative Window remains the author's framing of the management error the lag produces.

them, you review their work, you watch their performance over time, you tell them where they need to grow, and, when the role and the person stop fitting, you make a change.

An AI system asks for exactly the same attention. It has to be onboarded into the real workflow rather than dropped beside it; it has to be trained on the firm's actual material and corrected when it drifts; its output has to be reviewed rather than trusted by default; and its performance has to be watched, because the same system can degrade as the data around it shifts, as the people using it change how they use it, or as the underlying model is updated by a vendor without anyone at the firm being asked.

Reviewing the output is not a phase the system graduates from once it has proved itself. A model is not a calculator that returns the same answer to the same input, so three clean results do not guarantee the fourth, and the check has to hold every time by design rather than as a probation the tool serves and then leaves. The reason to keep checking grows over time instead of fading. As staff produce work with AI and then feed that work back as the basis for the next round, the origin of any figure or claim recedes a layer at a time, and a number no one verified at the start becomes close to untraceable once it sits three generations deep. A firm that relaxes its checking the moment the system looks reliable withdraws the discipline exactly as the material the system produces becomes hardest to audit.

And the analogy holds at the end as well: when the technology a system is built on stops advancing, when a better model, a better tool, or a better approach has arrived, the right move is to switch, the same way you would help a person move on when the role has outgrown them.

None of this is failure management. It is the ordinary, ongoing tending that any productive relationship requires, and a firm that budgeted for a one-time delivery has budgeted for the wrong thing. The cost of an AI capability is not the cost of standing it up; it is the cost of keeping it good, which connects directly to the recurring-cost guardrail of Chapter 3 and the deployment-model arithmetic of Chapter 5. A system no one is tending is not stable. It is drifting, and the only question is whether the firm notices before its customers do.

The strategy itself has to move

The relationship has to be managed at the level of the individual system, and it also has to be managed at the level of strategy, because the ground the strategy stands on is moving faster than the planning cycle that is supposed to govern it. Models are released, deprecated, and repriced on timescales of months. Capabilities that justified a build decision last year are commodity features this year. Pricing that made a per-seat model cheaper than building can invert without warning. Regulation that did not exist when the firm deployed arrives and changes what is permissible.

An annual strategy cycle assumes the world holds still between reviews, and in this domain it does not. The answer is not to plan more often in a panic; it is to decide, in advance, what would make the strategy wrong, and to revisit it when those things happen rather than when the calendar says to. The strategy is expected to move, and

the firm defines the triggers that move it instead of letting the planning calendar decide when to look.

A major model release in the firm's domain is a trigger. A pricing change that crosses the build-versus-buy line is a trigger. A regulatory shift, a security disclosure in a tool the firm depends on, a competitor's capability that resets customer expectations: each is a trigger that should pull the strategy back onto the table, regardless of where the firm is in its planning year.

Drift that happens by accident is how a firm ends up running last year's strategy on this year's technology without having decided to. Drift that happens by design is how it stays current without living in a state of permanent re-planning.

Continuous re-evaluation is the capability

There is a theory that names precisely what this requires, and it is the one place in this document where it earns its keep. Dynamic Capabilities, the framework David Teece developed to explain why some firms sustain advantage in fast-moving environments while others lose it, holds that durable advantage does not come from any particular resource or position, which gets competed away, but from a higher-order ability to keep adjusting them.

Teece breaks that ability into three movements: sensing changes in the environment, seizing the opportunities they open, and reconfiguring the firm's assets and routines to match.[20]

The non-obvious claim, applied to AI, is sharp and worth stating plainly: the capability the firm is building is not the AI system. It is the practice of continuously re-evaluating the AI system. The model will be superseded; the tool will be replaced; the specific deployment will be reconfigured more than once. What persists, and what compounds, is the firm's developed ability to sense when the ground has shifted, to act on it quickly, and to reconfigure without trauma, which is the same decision-velocity advantage of Chapter 3, now turned inward on the firm's own AI estate.

For a mid-market company, this does not require a dedicated strategy function or a formal review board; it requires one named person whose job includes watching the AI landscape in the firm's domain, and a standing agreement that certain events, a major model release, a pricing change, a competitor's capability shift, pull the strategy back onto the table within two weeks, not at the next annual planning cycle.

This reframes the most common request leaders make of their advisors. A firm that asks for a static three-year AI strategy is not asking for foresight; it is making a structural bet that the environment will hold still for three years, in the one domain where that is least likely to be true. The valuable thing is not the multi-year plan. It is the

---

[20] David J. Teece, Gary Pisano, and Amy Shuen, "Dynamic Capabilities and Strategic Management," *Strategic Management Journal* (1997); and David J. Teece, "Explicating Dynamic Capabilities," *Strategic Management Journal* (2007), which develops the sensing / seizing / reconfiguring articulation.

standing capability to keep the plan honest, and that capability is built, like any other, by exercising it, which is why the relationship framing is not a soft metaphor but the actual mechanism by which the advantage is sustained.

This is also where the whole sequence closes back on itself. The mid-market's structural edge was decision velocity, and the discipline that protects it was Containment. Adoption-as-a-relationship is simply that same edge applied over time rather than at a single moment: the firm that can move quickly from decision to live system is also the firm that can move quickly to re-decide when the system stops being the right one. The advantage was never the first deployment. It was the ability to keep authoring the next one faster than the environment can obsolete the last. A firm that treats adoption as a project it can finish has, without realizing it, set that advantage down.

The decision the reader makes: What are your defined triggers to revisit the strategy, the specific events that would tell you the current plan has gone stale, and who, by name, is watching for them? If the answer is "we review it annually," you do not yet have a strategy for a moving environment; you have a calendar.

Conclusion — From Imitation to Authored Advantage

This paper opened on a gap: nearly every mid-market company has adopted AI, and almost none can show what it was worth. The reason is now visible. Companies surrendered the value the moment they reached for someone else's answer: the enterprise's playbook, a competitor's tool, the urgency in a headline. They stopped short of authoring their own.

Everything here argues one shift, from imitation to authorship. Your structural edge is real, but you do not own it; you practice it. Decide on purpose and you keep it. Copy, and it is gone.

Those decisions are not seven topics. They form one act, each choice setting the terms of the next. You start by owning the decision: do the people who run the business understand what AI can and cannot do, or did they hand that understanding to IT and call it strategy? Owning it lets you name one honest driver, and the driver fixes your metric, your sequence, and what success means. The driver defines what you contain, and how much of your speed survives before it turns to waste. Containment lands on people: does your plan make capable people more valuable, or ask them to train their replacement? Motivated people make readiness real, and readiness decides where your first move lands and what it costs to run at scale. You build security into that move or you lose it. Then nothing is finished, because the ground keeps moving. You name the events that would make the strategy stale, and the person who watches for them.

Those decisions are the document. You will not find a checklist here. You leave with seven questions about your own company, ones only you can answer. The work now is to answer them yourself, in order, before you buy anything. Do that and you have already left the ninety-five percent.

None of this runs on fear. The pressure that brought you here is real, but let it drive and you make the anxious, imitative choice this paper exists to prevent. Move for the size of the prize. What AI does for a company your size today is substantial, and what arrives next is larger. You are close to your own work and free of the machinery the giants cannot put down, so you can seize it before they do.

A closing word, in my voice. I wrote this because the mid-market keeps getting advice built for companies it does not resemble, and you pay for that mismatch in value that never shows up. Use the frame, argue with it, bend it to fit. It is yours now.

The advantage was never in adopting AI. It is in authoring what you make of it.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the author used Anthropic’s Claude to draft and edit the text; all ideas, research, and source selection are the author’s own, and the author takes full responsibility for the content.

References

Every source used in this document, listed once. Order: mid-market and survey evidence, then the academic anchors, then standards. Each entry notes how it is used and any caveat carried in the text.

Mid-market and survey evidence

RSM US, in partnership with Big Village. *Middle Market AI Survey 2025*. 966 middle-market decision-makers across the US and Canada; fieldwork February–March 2025. Primary mid-market anchor: 91% of mid-market firms use generative AI (up from 77% the prior year); 53% feel only “somewhat prepared,” 10% not prepared; 62% found adoption harder than expected; 70% needed outside help.

MIT NANDA. *The GenAI Divide: State of AI in Business 2025*. July 2025. Used for two figures: roughly 95% of AI pilots show no measurable P&L impact (Preface); and a pilot-to-implementation speed gap of about 90 days for top mid-market firms versus nine months or more for large enterprises (Chapter 3). NANDA defines “enterprise” as firms above $100M revenue, which overlaps the upper mid-market; a reconciling note appears in Chapter 3. Base: ~300 deployments, 52 interviews, 153 surveys; used for scale, not precision.

Benchmarkit and Mavvrik. *2025 State of AI Cost Governance*. 372 companies; reported by CIO.com, October 2025. ~85% of organizations misestimate AI costs by more than 10%, nearly a quarter by more than 50%, almost always too low; top cost drivers are data platforms and integration, not model tokens (Chapter 2). Vendor-sponsored; cited via CIO.com for the direction and scale of the error.

Pew Research Center. *Survey of US workers on AI in the workplace*. February 2025 (n = 5,273). 52% worried about AI’s workplace impact; 32% expect fewer job opportunities for themselves, 6% more; AI users more likely than non-users to expect fewer opportunities (42% vs. 30%) (Chapter 4). Cited for the broad anxiety; the “training my replacement” mechanism is the author’s framing.

Boston Consulting Group. *Split Decision: The BCG CEOs and Boards Survey*. May 2026. 61% of CEOs say their boards are pushing AI transformation faster than the business is ready for (Preface). Sample skews to firms above $100M revenue; cited as corroboration of the direction of pressure, not a mid-market measurement.

McKinsey & Company. *The State of AI in 2025*. Roughly 88% of organizations report using AI in at least one function, while only about 39% report a material effect on EBIT (Preface). Sample skews to larger enterprises; cited only to show the adoption-without-value gap is not unique to the mid-market.

Academic anchors

Baum, J. Robert, and Stefan Wally. “Strategic Decision Speed and Firm Performance.” *Strategic Management Journal* 24, no. 11 (2003): 1107–1129. Four-year study of 318

firms; strategic decision speed predicts subsequent growth and profit and mediates the effect of environment and structure on performance (Chapter 3). Supports the decision-velocity component of the argument.

Coase, Ronald H. “The Nature of the Firm” (1937); and Oliver E. Williamson’s subsequent development of transaction cost economics. Basis for the coordination-cost argument that the mid-market’s speed is an outcome of structurally lower internal transaction costs (Chapter 3).

Deci, Edward L., and Richard M. Ryan. Self-determination theory. Basis for the claim that controlled, mandate-driven motivation underperforms autonomous motivation on complex work, and that autonomy, competence, and relatedness are the conditions of the latter (Chapter 4).

Ibrahim, Fabio, Johann-Christoph Münscher, Monika Daseking, and Nils Torge Telle. “The Technology Acceptance Model and Adopter Type Analysis in the Context of Artificial Intelligence.” *Frontiers in Artificial Intelligence* 7 (2025): 1496518. Sample of 1,007; perceived usefulness and “AI mindset” are the strongest predictors of AI use (Chapter 4). Cited for that empirical finding.

Włoch, Renata, Katarzyna Śledziewska, and Satia Rożynek. “Who’s Afraid of Automation? Examining Determinants of Fear of Automation in Six European Countries.” *Technology in Society* 81 (2025): 102782. Perceived task substitution raises fear of automation; a greater sense of personal control reduces it (Chapter 4). Documents the general anxiety underlying the Trainer’s Dilemma.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. “The Productivity J-Curve: How Intangibles Complement General Purpose Technologies.” *American Economic Journal: Macroeconomics* 13, no. 1 (2021): 333–372. Explains why general-purpose technologies depress measured productivity early, while firms build unmeasured intangible assets, with returns appearing later (Chapter 7). Complements, and does not replace, the False Negative Window.

Teece, David J., Gary Pisano, and Amy Shuen. “Dynamic Capabilities and Strategic Management.” *Strategic Management Journal* (1997); and David J. Teece, “Explicating Dynamic Capabilities,” *Strategic Management Journal* (2007). Sensing, seizing, reconfiguring as the durable capability; basis for the continuous-re-evaluation argument (Chapter 7).

Standards and definitions

National Center for the Middle Market (NCMM). Mid-market definition used throughout: $10M–$1B revenue; segments Lower ($10M–$50M), Core ($50M–$500M), Upper ($500M–$1B).

OWASP and NIST. Established taxonomies for prompt injection, jailbreak, and data exfiltration through models; referenced at the category level (Chapter 6). Specific control names and version numbers are not asserted.

Author & Disclaimer

About the Author

Katerina Andreeva is an AI strategy adviser and the founder of The Brained Inc. She studied Statistics and Economics at the Ludwig Maximilian University of Munich and spent a decade working across data analysis, data science, data engineering, and machine learning, most recently at IBM as a Customer Success Manager, Data & AI Architect. She has also taught database fundamentals at the Baden-Württemberg Cooperative State University (DHBW).

Disclaimer

Passages marked “from practice” draw on the author’s advisory work; identifying details are anonymized.