



AI TRANSFORMATION

Why Most Pilots Never Fly

Enterprises have folded thousands of AI pilots. Almost none of them ever climb. The reason is rarely the plane.

By **Guillaume Pajeot**, Insigniam Partner

Somewhere in the architecture of childhood sits a single sheet of paper and the unshakable conviction that it can be made to fly. The folds are a kind of liturgy, passed down on playgrounds and in the backs of classrooms. The lengthwise crease. The nose drawn to a point. The wings pressed flat beneath a thumbnail. Then comes the throw, that small and serious act of faith, the arm cocked back and the wrist snapping forward. And then, far more often than not, comes the disappointment. The plane stalls, tips its nose, and spirals into the floor a few feet from the hand that launched it. The folding was earnest. The launch was committed. The thing simply refused to climb.

It is a quiet accident of language that the word we hand to that hopeful little aircraft happens to be the same word the modern enterprise has chosen for its own hopeful little experiments. We call them pilots. We launch them with the same earnestness, and we get back roughly the same result. A brief and promising arc. A stall. A slow descent into the carpet of the quarterly review, where the program is logged as a learning and the next clean sheet of paper is reached for.

The arithmetic behind that descent is no longer a matter of anecdote. Researchers at MIT's NANDA initiative, studying the recent wave of enterprise generative AI spending, found that 95% of deployments produced no measurable impact on the profit-and-loss statement, this despite the tens of billions of dollars committed to them. The Stanford HAI AI Index, surveying the same landscape from a different altitude, reports that organizational adoption of AI has climbed to 88% and that corporate investment in the technology more than doubled in a single year. Those two findings do not argue with each other. They describe one room from opposite walls. Nearly everyone has thrown a plane. Almost no one is flying.

Distance of that kind is not crossed by buying a faster engine. It is crossed by getting the order of operations right, and the order turns out to matter more than any single tool inside it. Elixirr calls the sequence AI Value Sequencing, and its central, slightly uncomfortable claim is this: most organizations run it exactly backwards.



THE HARD TRUTH ABOUT AI TRANSFORMATION

Most AI programs are busy, and busy looks like progress. Pilots multiply, tools accumulate, dashboards fill, and none of it changes how a single decision gets made or who is accountable for the result. That is motion, not flight. A program moves the moment money is spent and climbs only when the work is rebuilt. Almost everyone funds the first and calls it the second.

AI ADOPTION 101



How to Close the AI Value Gap

The thing nobody wants to say out loud is that the technology mostly works. The models are fine, the integrations are fine, the demos are fine. What is broken is the conversation around them, which measures activity instead of value. Prompts run, pilots launched, seats deployed. None of it answers the only question the CFO asked. Did this change the P&L?

That is the value gap most programs sit in. Capability has raced ahead of business model. The AI is in the building, but the decisions, roles, and operating model it is meant to change have not moved. Hours get saved and never become capacity redeployed. The cost line holds, the revenue line holds, and nothing was ever set up to move.

Last quarter a CEO told me his company had forty-seven use cases. I asked which one had a named owner accountable for a specific outcome. He thought, then said zero. Forty-seven use cases is not an AI strategy. It is an AI hobby with a budget.

What separates the few who scale from the many who stall is sequencing. They settle the outcome, the owner, and the trade-offs before a single agent ships, then redesign the operating model around it. Most companies do it backwards: buy the platform, run a hackathon, declare a strategy, then spend two years retrofitting value out of code never anchored to an outcome.

Agentic AI only sharpens the point. The agent does the task and the person supervises outcomes. That is a redesign of the job, not a productivity gain.

AI adoption is not a technology problem with a culture wrapper. It is an op model problem with a technology trigger. The companies that scaled did not get smarter at AI. They got serious about redesign and made the hard choices early. Until that line is crossed, every new pilot is just movement that never shows up where it matters. **IQ**

—Adam Hofmann, Elixirr Partner

WHAT GROUNDS AI FROM SOARING

• • •



71%

PROVE IT OR LOSE IT
CIOs say AI budgets will be frozen or cut without proof of value within two years.¹



6%

CORE TRUST
Only six percent of firms fully trust AI agents to run core business processes.²



0%

THE ADVANTAGE MIRAGE
Across 800 public firms, automation potential showed no link to profit; the gains get competed away.³



MOMENT OF TRUTH

Ask yourself:
Has your AI changed *how* the work actually gets done, or only how fast?

The Launch

The trouble usually begins, as trouble often does, with a misreading of what is being attempted. An organization decides it must “do something about AI.” A platform is purchased. A hackathon is convened. Use cases are nominated from every function with a pulse, and a leaderboard of pilots is born. Only later, once the agents are half-built and the invoices have arrived, does anyone get around to the question that should have come first. What outcome were we trying to produce, and how would we know if we produced it?

That inversion is visible in the data, if you look past the adoption headline to what the adoption actually consists of. The U.S. Census Bureau, through its *Business Trends and Outlook Survey*, has begun asking American firms not whether they are excited about AI but whether they used it, in an actual business function, in the prior two weeks. As of the spring of 2026, the answer was about one company in five. More telling still, among the businesses that had adopted AI at all, well over half were applying it in three or fewer functions.

The Federal Reserve, tracking the same question through its own surveys, lands in the same neighborhood. The picture is not one of broad transformation. It is one of narrow experimentation, distributed widely and connected to very little.

A handful of disconnected pilots scattered across a handful of departments is not a strategy. It is a collection of hopes with a budget line. And a hope, however well funded, has no opinion about what the business is trying to win. This is the first and most expensive misunderstanding about enterprise AI: that the technology is the project. The technology is never the project. The outcome is the project, and the technology is merely the most recent thing that makes the outcome reachable.

Strategy, in this telling, is not the slide that names twelve priorities. It is

the harder act of naming the one or two outcomes worth reorganizing around, and then having the discipline to let the others wait. A pilot launched without that clarity is a paper plane thrown without a thought for where it is meant to land. It will travel exactly as far as its initial momentum carries it, and then it will stop.

You Cannot Throw It Harder

Watch a frustrated child whose plane keeps nose-diving and you will see the universal first instinct. Throw it harder. Put more arm into it, more shoulder, more want. It never works. A plane that is folded to stall will stall at any speed. The defect is in the design, not the effort, and no amount of additional force will argue a bad airframe into the air.

Enterprises reach for the same instinct and call it scale. The pilot showed promise, the thinking goes, so we will simply do more of it, faster, with a bigger budget and a steering committee. But speed applied to an unchanged operating model does not produce transformation. It produces a faster version of the thing that was not working, which is usually just an earlier



Design the Flight, Not the Throw

Everyone launches pilots, but no one builds the aircraft: fewer than 1-in-10 companies run real AI agents, says Stanford HAI, 2026 AI Index Report.⁴

arrival at the same disappointment. This is the second misunderstanding, and it is the costly one, because it feels like progress while it happens.

Here is the claim that the evidence keeps insisting upon, however much leaders would prefer a tidier one: *AI does not deliver value. Operating model change delivers value. AI is simply the reason you finally have permission to do it.*

The phrase belongs to Elixirr, and it reorders everything. It means that the second move in the sequence, after the outcome is named, is not to deploy technology. It is to redesign the work itself: how decisions get made, who is accountable for what, where human judgment belongs and where it does not. AI becomes the occasion for that redesign rather than a substitute for it. Skip this step and you are throwing the same plane harder.

The clearest illustration of the principle in practice comes from a banking client that decided, sensibly, to stop optimizing its old product process and to redesign it. Rather than bolt AI onto the existing lifecycle, the firm rebuilt the lifecycle as an AI-native, agent-led model, with intelligent agents handling

opportunity identification, business case creation, development, quality assurance, and benefits tracking. Security, risk, and compliance were built into that workflow from the outset rather than added at the end. The point was not to make the old assembly line quicker. The point was a different assembly line.

The numbers the redesign was built to produce are the kind only a redesign can produce. Across product, technology, and operational teams, the new model was engineered to unlock 50% productivity over ten years. Best-case time to market, the stretch it takes to carry an idea from conception to launch, was set to compress from roughly six months to a window of two to six weeks. Long-term technology run costs were put on a path to fall by as much as 18%. None of that arrives by throwing the old plane harder. It arrives by folding a new one.

There is an instructive irony in who tends to understand this first. Consulting firms built on the traditional pyramid, with armies of junior staff billed by the hour, have a structural reason to flinch from a technology that compresses the very hours they sell. Elixirr, which is built on a different shape, does not carry that conflict. “Elixirr was made for AI,” says Dr. Emiko Caerlewy-Smith, a partner at the firm. The business model, she explains, was “intentionally created to be nimble, to be agile, and to have this cylindrical shape so we had as many people adding value at senior levels as we do at the more junior levels of our organization.” The firm has long, in her words, “focused on outcomes-based pricing versus time and materials pricing,” which means that when AI makes the work faster, the incentive points in the same direction as the client’s. “It’s in everyone’s interests, ourselves and our clients, to deliver outcomes, to deliver value quickly,” she says, “and AI just supports us to be able to do that.” A firm that profits from speed has no reason to dread it.

“Organizations are deploying AI tools extremely quickly, often without applying the same level of security assurance they would normally apply to other enterprise software.”

—Joe Hubback
Elixirr Partner & CISO

CLIENTS SOUND OFF



When AI pilots fail to leave the ground, here’s what clients say about their stalled investments.



Our AI adoption is too slow.



Our AI investments aren’t delivering the results we want.



How do we get our AI initiatives adopted?



We flat out don’t know how to use AI.



The Airframe

Only now, with the outcome named and the operating model redesigned around it, does the technology earn its turn. Third in the sequence, not first. And here the conversation finally becomes the one most organizations wanted to have at the very beginning, about models and agents and infrastructure, except that it now has somewhere to land.

The instinct in many companies is to acquire a fleet of individual AI tools, one per problem, and to hope they cohere. Adam Hofmann, a partner at Elixirr, hears the ambition stated plainly. “I want to have all these AIs that are working inside of my business,” clients tell him. His reply is a set of questions they have rarely asked themselves. “Okay, how are those AIs going to get work done?” he says. “Can they talk to each other? Can they be orchestrated to combine things together?” The answer, very often, is a shrug. A drawer full of clever tools that cannot coordinate is not an intelligent organization. It is a drawer.

The firm’s own house offers a worked example of the alternative. Elixirr, in Mr. Hofmann’s account, “structured all of our internal data, think proposals, past projects, deliverables, SOWs, into this advanced data layer that we can build AI on top of to drive workflows.” A consultant who once needed days or weeks to produce a proposal can now hold a conversation with that system and, as he puts it, “get proposals out instead of taking days or in some cases weeks to now hours.” The legal team runs contracts and statements of work through a similar process, with guardrails and security wrapped around it, turning a slow series of reviews into a fast and governed one. None of this is magic. It is the unglamorous work of giving the airframe enough structure to actually catch the air.

The same disciplined building shows up inside client environments, and two examples mark the range of it. A leading U.S. financial services group

had been spending as much as thirteen hours of employee time to produce a single marketing asset, grinding through thirty-nine manual steps for every email. In a six-week proof of concept, Elixirr built a campaign agent on large language models, calibrated to the firm’s content, its customer personas, and its brand standards. The agent now assembles an entire campaign in thirty-seven minutes, less time than it once took to write a single email, a pace twenty-one times faster than before, holding to brand guidelines at better than 85% compliance and personalizing to the individual recipient. Speed, here, was the by-product of a process that had been rethought rather than rushed.

The second example trades velocity for sheer scale. A professional services organization needed to classify, rename, and migrate more than a quarter of a million documents trapped in inconsistent systems across two hundred and forty-nine clients, all of it under regulatory scrutiny. Elixirr built a classification engine with correction mechanisms designed in from the start, reached better than 90% accuracy, and established one hundred and forty-nine standardized folder structures, automating work that would otherwise have consumed thousands of manual hours while leaving an audit trail behind it. The system also surfaced something no one had asked it to find and the firm badly needed to know: regulatory documents that were missing entirely, absences no one had thought to go looking for.

What unites the three is not the cleverness of the models. It is that in each case the technology arrived third, after someone had decided what outcome mattered and how the work would be rebuilt to produce it. The plane flew because the plane was designed to fly, not because it was thrown by a better arm.

Instruments and Controls

A plane that climbs introduces a problem the stalling plane never had.

CASE STUDY

When Translation Became Oversight

A leading pharmacy services provider processes thousands of prescriptions a day, and a stubborn share of them arrive in a form no system can fill: free-text directions, written the way clinicians actually write, that must be converted into standardized Sig Codes before a prescription can be dispensed. The work is exacting, it is repetitive, and it had only ever been done by hand. A technician could clear about thirty prescriptions an hour, each one a small opportunity for error. The only obvious way to handle more volume was to hire more hands, which is to say there was no sustainable way to handle more volume at all.

The firm could have asked for a faster version of that manual line. It asked for a different one. Rather than bolt a tool onto the existing process, the team built an AI system that reads the full, untidy range of natural-language instructions and renders them into standardized codes, then wired it directly into the order management system so it lived inside the work rather than beside it. A feedback loop let technicians review and correct what the model produced, so the system grew more reliable the longer it ran and, more importantly, earned the trust of the people who depended on it. The technician was not removed from the process. The technician was moved up it.

What the redesign produced was not merely speed. It was more than half a million dollars in annual savings, a workforce redirected from translation to quality control, faster turnaround, and for the first time a process the firm could scale without simply scaling its headcount alongside it.

The Hard Truth

Automation no one can correct is automation no one will trust, and untrusted systems don’t scale — they get worked around. The value was never faster translation. It was moving people from doing the work to governing it. Speed any vendor can sell; trust has to be designed in.

Now it can go somewhere, which means it can go somewhere wrong. Altitude is the moment controls start to matter. This is the fourth and final move in the sequence, the one most often treated as paperwork to be completed after the interesting work is done, and it is the move that determines whether a working system is one you can actually trust at scale.

Governance, in the AI context, is not a brake pedal pressed against innovation. It is the instrument panel that lets you fly the thing in weather. The distinction shows up plainly in the cases already described, where security, risk, and compliance were not appended at the end but designed into the workflow from the first stroke, and where a document system earned its keep precisely because it could produce an audit trail and catch a regulatory gap. Trust was not a constraint on those outcomes. Trust was the outcome that made the others usable.

That posture has an advocate. Joe Hubback, a partner at Elixirr who advises clients on the security and trust dimension of these programs, reduces what a board ought to ask to two fundamental questions. The first: “How do we know it works?” The second: “How do we know we can control it?” A board that cannot answer the first has bought a rumor. A board that cannot answer the second has bought a risk it does not understand.

Beneath the questions sits a principle he calls “risk transparency,” knowing precisely “where they are accepting risk and what their fallback options are if something goes wrong.”

Much of the AI ecosystem, he cautions, still “runs on trust between buyers and vendors,” which is exactly the condition that governance, designed in rather than bolted on, exists to retire. The questions sound basic, which is why they are so often skipped in the rush to ship, and why the programs that endure are the ones that treat them as design requirements rather than afterthoughts. Get the instruments right and a working



About the Author
Guillaume Pajeot,
Insigniam Partner

Guillaume Pajeot has consulted for clients in most Western and Central European countries, as well as in Algeria, Australia, China, Egypt, South Africa, Turkey, and the Commonwealth of Independent States. Among his clients are firms in the fast-moving consumer goods, pharmaceutical, film, automotive, and chemical industries, as well as agricultural cooperatives. He has trained internal consultants and implemented numerous innovation projects. Mr. Pajeot possesses expertise in breakthrough projects, culture transformations, tackling the challenge of digital transformation, growth and new marketing strategies, and the transformation of labor-management working conditions. He has experience creating breakthrough strategies and consulting on the execution of those strategies, with boards and executive committees, through management to shop-floor teams. Mr. Pajeot holds an MSc in Chemical Engineering from the École Nationale Supérieure de Chimie de Lille, an MSc in Environmental Pollution Control Management from Heriot-Watt University, and an MBA from La Sorbonne.

system stops being an experiment you have to keep relaunching. It becomes something the organization can fly in any weather, and keep flying. The pilot is no longer the point. The capability is.

The Climb

Return, for a moment, to the child and the paper plane. The remarkable thing about that small failure on the classroom floor is how completely it can be solved, and how rarely the solution is effort. The plane that finally flies is not thrown harder than the one that stalled. It is folded differently. The weight is moved forward, the wings are trued, the design is corrected, and then a perfectly ordinary throw carries it the length of the room. The transformation feels like magic and is in fact just sequence: get the design right, and the flight takes care of itself.

The enterprise version is no different, though it is more expensive to learn. The 95% of AI deployments that go nowhere are not failing for want of better models. The models are extraordinary and getting more so by the month. They are failing because they were thrown before they were designed, scaled before they were rebuilt, deployed before anyone decided what winning meant or how the work would have to change to achieve it. S&P Global has found that 42% of companies abandon the majority of their AI initiatives before they ever reach production. Those are not crashes. They are planes that were never folded to fly in the first place.

The organizations that break the pattern do something that looks, from the outside, almost anticlimactic. They name the outcome first. They redesign the work second. They deploy the technology third. They build the controls fourth. They run the sequence in the order that actually produces flight, rather than the order that produces motion. It is not the glamorous path. It is merely the one that leaves the ground.

Anyone can throw a plane. The discipline is in building one that climbs. **IQ**

For over 35 years, executives at the world's largest and best companies have relied on Insigniam's unique consulting to produce critical, unexpected outcomes, utilizing proprietary methodologies that marry breakthrough performance and innovation.

WE PROMISE

that your people will
think newly,
act differently
and deliver
unprecedented results.®

For more information,
visit www.insigniam.com