AI AND BANKING - A TOUGH RELATIONSHIP

Embracing AI in financial services has become a complex, often confusing love–hate relationship. Leaders are on board with leveraging AI because they believe it should be based on inside/outside conversations they are having or what they are reading. Meanwhile, further down the stack, many workers are resistant to participating in and adopting it. The sheer volume of information available to leaders and workers alike is mind-boggling, highlighting the triumphs and tribulations of AI adoption, as well as the trade-offs between corporate value and unemployment. Depending on their perspective, readers may be amazed by the opportunities and value, and/or see a doomsday event coming.

And there is a middle ground; a credit analyst is using ChatGPT on her phone to draft a loan memo. A compliance associate is running regulatory language through Claude to check his interpretation. A branch manager is summarizing customer complaints using a tool nobody in IT approved, and an operations analyst has gone further — she's built a reconciliation helper that runs continuously, catches discrepancies faster than the legacy system, which has no plan for replacement, and has quietly saved her team hours every week for months. They're not waiting for the platform decisions. They've already decided — and each person is working at a different level of sophistication, from simple prompts to something that looks a lot like the automated workflows that a steering committee may still be debating whether to fund.

Every consultancy and evangelist of AI is producing material around CXO Playbooks, What Boards should be thinking about, Enterprise AI Strategies, Critical Success Factors, Blueprints and Frameworks, etc., all of which exude paths to a promised land. With all this collateral and wisdom available, financial services still seem to have real challenges turning all these paths to nirvana into reality.

Is this struggle due to a lack of true understanding of AI and what is needed to implement it? Do I buy, build, or outsource a solution? Is my data/technology in a usable condition, and am I required to fix it or develop a data/infrastructure strategy? Should I expect immediate ROI or burn a lot of money on POC's with no value achieved? Do I need to evolve my whole company's culture, push for an AI-first view, and embrace a "change is good" mentality?

The challenges seem more like leadership vs. worker, knowledge vs. fear, reality vs. expectation. Let's explore.

BANKS ARE AI ADOPTERS

Financial services institutions, from global to community, are in varying degrees of AI adoption, and much of it has to do with middle management thinking.

The Federal Reserve Bank of St. Louis reported in early 2025 that generative AI adoption *[do you even know what "generative AI" is*] reached 54.6% among working-age adults in the US, with finance ranking among the highest-adopting industries. An EY-Parthenon 2025 survey found 77% of banks have already launched generative AI applications of some kind. These numbers describe an industry where individuals and small teams adopted quickly — often faster than the institutions that employed them.

A counterintuitive finding is that the "banking is behind" narrative was always a story about institutional capability, not individual willingness. The demand signal for AI inside banks is enormous. It arrived before the infrastructure, before the governance frameworks, before the measurement systems, and, in many cases, before leadership had a coherent position on what AI should be used for (performing tasks vs. adjusting workflows vs. producing decisions and creating outcomes). The workforce didn't wait for leadership decisions/permission, which is both encouraging and potentially chaotic for overall company strategy — encouraging because it means the adoption energy exists, chaotic because it reveals exactly how wide the institutional gap is.

A fair objection: someone using ChatGPT on a personal device to rephrase an email is a long way from institutional AI deployment that touches core banking processes or credit decisioning. This is true, but the distance between those two things is precisely the distance that matters. Individual readiness and institutional capability are supposed to move together. When they diverge sharply, the institution is the bottleneck — and the question shifts from "how do we get people to adopt AI" to "what's preventing the institution from absorbing what its people already want to do."

WHERE ENTHUSIASM GOES TO STALL

Bank leaders have probably seen this pattern: a team builds a proof of concept, it works, everyone is excited, and then it dies — not because the technology failed, but because everything around it wasn't ready.

MIT's research initiative on AI implementation, published in 2025, found that 95% of AI pilots deliver zero measurable P&L impact. That number is jarring, but what sits underneath it is more useful: roughly 80% of the work required to move from pilot to production is institutional — data engineering, governance frameworks, stakeholder alignment, workflow integration, and the list goes on. The model itself may be 20% of the problem. The other 80% is the organization's ability to receive what the model produces and act on it.

This is why pilots fail in banking at the rate they do, and why the pattern is so consistent across institutions of different sizes. The pilot succeeds in a controlled environment with clean data and motivated people. Then it meets the production environment — siloed data architectures where customer data lives in one system, transaction data in another, and neither talks to the risk platform. The compliance review process, designed for a different era of technology change, adds weeks, and the core systems it needs to integrate with run on decades-old infrastructure that nobody fully understands, because the people who built them retired years ago. The pilot didn't fail. The institution itself showed you exactly where it breaks.

Gartner predicted that 30% of generative AI projects would be abandoned after proof of concept by the end of 2025. That number is probably conservative, and the reason connects directly to measurement. Bank Director's 2025 survey of 141 directors and executives at US banks with less than $100 billion in assets found that 82% don't measure ROI on any

technology investment. Not just AI……. any technology project. POC's conducted without a measurement infrastructure produce activity, not evidence. You can't tell whether a pilot succeeded if you never defined what ultimate success would look like, and you can't build a case for scaling something you can't quantify. This isn't a failure of discipline — measuring avoided losses, regulatory risk reduction, revenue generation, and improvements in operational resilience requires frameworks that banking's existing cost-accounting structures weren't designed to capture. But that makes the gap more important to close, not less.

The pattern compounds: pilots launch without clear success criteria, produce results nobody can measure, and leadership—reasonably, given the information available to them—concludes that AI doesn't work for their institution… or that adopting AI is simply too hard. What actually happened is that the institution couldn't evaluate what it built, which is a different problem entirely and a much more solvable one. The pilot graveyard is telling you something specific: banks haven't built the institutional infrastructure — starting with measurement — to know whether anything works.

LEADERSHIP BUYS PRODUCTS WHEN IT NEEDS TO BUILD CAPABILITY

S&P Global's 2025 banking survey revealed a figure that captures the core tension: 91% of bank boards have approved generative AI programs, while only 26% have developed the capabilities to execute them. Ninety-one percent approval. Twenty-six percent capability. The gap is between what leadership thinks AI is, what it requires, and what it actually is and requires.

The pattern shows up across multiple data sources. IBM's 2025 Institute for Business Value study found that only 8% of banks are approaching AI strategically, while 78% remain tactical — buying point solutions for specific issues/solutions and running department-level experiments that never connect to a broader capability. The institutional posture is piecemeal because the mental model is piecemeal: AI is a product you buy and install rather than a capability you build across strategy, people, processes, and technology.

In 2024, Gartner articulated an outdated CFO dimension: "GenAI requires a higher tolerance for indirect, future financial investment criteria versus immediate ROI. Historically, many CFOs have not been comfortable with investing today for indirect value in the future." That discomfort was rational back then, given how banks have historically evaluated technology as a cost center with measurable returns on a predictable timeline.

AI doesn't fit that model. Its value is often indirect, its timeline uncertain, and its returns compound rather than arrive. A CFO trained to demand immediate ROI from technology spending will systematically underfund the kind of patient, iterative investment that AI capability requires.

Part of the confusion is that "AI" currently has three very different layers of value, and most bank leadership teams are buying one while expecting another. The first layer is performing a task, which is how most people in/out of the industry use AI: an employee prompts a model to produce a message, summarize a document, or research a topic and create tables, reports, or presentations. It's useful, and it's what most banks are actually running when they say they've "adopted AI" — but it doesn't change how the business operates. The second layer is automation: kickoff sections for operational activities such as producing a regulatory report, displaying exceptions, and advancing a simple process. This is

faster and more consistent, but still optimizes yesterday's structure. The third layer is decisioning, learning, and self-evolving, and proactively displaying risks and opportunities.

Decisioning, learning, and evolving systems are where a company's operating model changes significantly — feedback loops shorten, decisions improve, and value compounds. Most bank AI investments are task-oriented, with only partial processes automated, while leadership expects larger results than an outcome-driven system can produce. That gap between what's purchased and what's expected is the bolt-on trap in a single frame: you can buy task-level AI the way you buy software, but decision-driven AI requires building institutional capability that no vendor can sell you.

The consistent finding across IBM, Deloitte, and S&P Global's research confirms this at scale. The primary barriers to AI execution aren't technological and aren't even primarily regulatory. They're cultural. Organizational structures that can't support cross-functional initiatives, an institutional habit of treating technology as something you purchase rather than something you become, and a vendor ecosystem that reinforces the bolt-on framing because bolt-on solutions — task-level AI — are what vendors sell.

You can buy a fraud detection model, but you can't buy the data governance that makes it accurate, or the cross-functional workflows that route its outputs to the right people, or the measurement discipline that tells you whether it's actually reducing losses or generating more revenue. Those capabilities have to be built, and doing so requires a sustained organizational commitment that the bolt-on mindset specifically avoids.

THIRTY YEARS OF RATIONAL DEFERRALS, ALL DUE AT ONCE

The gap between individual readiness and institutional capability didn't evolve overnight. It was dug over the course of decades, and the decisions that created it were rational at the time.

Start with mergers and acquisitions. US banking has consolidated relentlessly since the 1980s, and almost every merger brought together incompatible systems, many of which were never fully integrated. The acquiring bank kept its core platform. The acquired bank's data was migrated — partially and imperfectly — on a timeline driven by cost rather than completeness. The IIF and EY's 2025 joint survey found that 96% of financial institutions cite noisy, untimely, or inaccurate data as their primary challenge in deploying AI. That number reflects forty years of M&A in which data integration was treated as a back-office problem to be dealt with later — and later arrived, carrying decades of deferred integration costs.

The core banking infrastructure tells the same story. A meaningful share of US retail bank transaction processing still runs on COBOL-based systems designed in the 1970s and 1980s. Bank Director survey found that 41% of banks with less than $100 billion manage business-line data in spreadsheets. The Kansas City Fed and CCG Catalyst have documented that the top three core providers — Fiserv, FIS, and Jack Henry — serve approximately 70% of US banks, a concentration that emerged because outsourcing core processing was the rational response when technology was a cost center and economies of scale favored specialists. That decision made sense for decades. It also meant banks traded control of their most critical infrastructure for cost efficiency and now face vendor dependency at exactly the moment when technological flexibility matters most.

Each of these decisions — deferred data integration, aging core systems, and outsourced infrastructure- was individually rational. Technology was a cost center. Banking was a relationship and compliance business. Investing heavily in technology beyond what regulators and operations required didn't make the short list of strategic priorities, and the executives who made those calls were optimizing correctly for the environment they operated in. But each deferral constrained future options and raised the cost of eventual action, and the deferrals compounded in ways that weren't visible until something demanded them all at once. Think of real scenarios, like a bank that deferred data integration at the turn the century post an acquisition, outsourced its core processing 5 years later, patched the gaps with spreadsheets over the next 10 years, and added a compliance layer that talks to none of the above, That bank faces a starting position where AI's requirements read like an indictment of every shortcut that seemed reasonable at the time.

AI didn't create this debt. AI made it visible and made the cost of carrying it forward untenable, because AI's value depends on exactly the things these deferrals neglected: unified data, flexible infrastructure, and the organizational alignment to change how work gets done. And the irony is the same cost-center mentality that made each individual deferral rational is now the biggest obstacle to the technology that could finally make banking's cost structure sustainable.

SAME TOOLS, DIFFERENT RESULTS — AND THE DIFFERENCE ISN'T TECHNOLOGY

Every bank in America can license the same large language models, hire from the same talent pools, and engage the same consultancies. The tools are a commodity. The results are not. Deloitte's financial services AI survey solicited insights from 542 leaders across the industry, divided firms into "pioneers" and "followers" based on AI maturity, and the gap between them is instructive because the difference was how they organized around what they bought.

The difference in measurement is the most telling. Pioneer firms designed their AI initiatives to measure outcomes from the start, tying every use case to a specific value driver — revenue impact, cost reduction, risk mitigation, and client satisfaction. They built the evaluation infrastructure that 82% of mid-tier banks lack. The result: 74% of pioneers expected greater than 10% ROI from AI, compared to 44% of followers. That's not optimism. Pioneer firms built the systems to know whether something was working, which meant they could iterate and compound, while followers deployed without measurement and couldn't distinguish success from failure.

The access difference is equally stark. Pioneer firms gave 40% or more of their workforce access to AI tools, compared to 19% at follower firms. Same technology, different institutional commitment. Pioneers treated AI as an organizational capability — something that required broad access, training, internal evangelism, and the willingness to learn from early failures rather than demanding immediate returns. They set aside a budget and were willing to learn from how it was spent, which is a fundamentally different posture than demanding a business case for every experiment. Followers treated AI as a specialized tool for a small team, which limited both the learning surface and the organizational buy-in needed to scale. When you give 19% of your workforce access, you're telling the other 81% that this isn't their problem — and then wondering why organizational adoption stalls.

Capital One offers one of the clearest cases of what sustained commitment looks like. Ranked second on the 2025 Evident AI Index, Capital One didn't arrive there through a big strategic bet. CEO Richard Fairbank has described the company as being in its "thirteenth year of an all-in technology transformation." That means Capital One started

deliberately building this capability in 2012, a full decade before generative AI entered public consciousness. Thirteen years is an outlier timeline, and most banks shouldn't need that long. What Capital One demonstrates is that alignment across strategy, technology, people, and process compounds over time, and by starting earlier, with commitment, creates future advantage. Tying AI to the company's Mission, Vision, and Strategy, as well as its future aspirations, was paramount in this endeavor.

S&P Global's 2025 analysis suggests the divergence window is three to five years. Institutions that achieve this kind of alignment in the near term will pull away from those that don't, and the gap will widen rather than close — because the operational learning, the data quality improvements, and the organizational muscle that come from running AI in production all compound. Changing later is more expensive than changing now. Much more expensive.

BANKING'S ALIGNMENT PROBLEM

The alignment problem in banking lies at the intersection of strategy, technology, people, and processes. The technology is available and largely commoditized. The demand signal is already there. Banking professionals are among the heaviest individual adopters of AI compared to almost any industry. What's missing is the institutional machinery to connect the two — the measurement infrastructure, the data foundations, the cross-functional collaboration, the patience to invest in indirect value that compounds over time, and the collective culture to adopt. Pioneer institutions didn't succeed because they found better technology. They succeeded because they aligned what they were solving for with what they built, who had access, and how they evolved their processes and mindset to absorb the tools' output. Institutions that addressed those dimensions together produced outcomes. Those that addressed them piecemeal produced pilots.

The vendor(s) banks select to assist in achieving their aspirations, albeit important, are more commoditized today and will matter less than whether banks have built the measurement infrastructure to know if the vendor's solution works. The budget banks approve will matter less than whether they give their workforce broad enough access to generate the organizational learning that compounds into capability. The conversation that matters is whether the institution — from how it sets direction to how work actually gets done — is aligned to absorb what AI makes possible.

The credit analyst is still drafting loan memos on her phone. The operations analyst's reconciliation helper is still catching discrepancies faster than the system it's not supposed to replace. The demand signal is already inside your building. Your people are already using these tools, already finding value, already working around the institutional constraints. The institution just has to be built to hear it.

This article was authored by John Stefani of OPCO Advisory with contributions from Oliver King of Mythmaker Labs. We would also like to thank the entire OPCO Advisory team for their valuable input.