INSOFTEX

WHY AI IN FINTECH IS NO LONGER ABOUT MODELS

WHY AI IN FINTECH IS NO LONGER ABOUT MODELS

It's About System Readiness

Executive Summary

Fintech teams usually focus on picking the best AI model. But choosing the model is less important than many think. It seldom decides whether an AI project delivers real value or gets stuck in a long proof-of-concept stage.

The biggest challenge is deciding whether the system is ready. This means having good data pipelines, mature integrations, clear decision workflows, and the ability to deploy in regulated settings. These factors are better predictors of project success than model performance.

This paper introduces a four-layer framework for evaluating AI readiness. Also, it identifies common mistakes that slow progress, highlights where AI delivers the greatest value, and explains what is needed to move from pilot to production. According to McKinsey's 2025 State of AI report, only a third of organizations measure AI across their businesses. The gap between using AI and seeing results comes from system challenges, not model choice.

> The central argument is that AI in fintech does not fail because of the models, but because the surrounding system is not ready.

Problem & Market Reality

Most fintech organizations have started at least one AI initiative. According to McKinsey (2025), 88% report regular AI use in at least one business function, up from 78% in the previous year. However, enterprise-level impact remains restricted: only about one-third have started scaling AI across operations, and just 21% have radically redesigned workflows around it.

88%	~1/3	21%
of organizations use AI in at least one function	have begun scaling AI enterprise-wide	have redesigned workflows around AI
McKinsey, 2025	*McKinsey, 2025*	*McKinsey, 2025*

> The challenge is not access to AI. It is translating AI capability into reliable, auditable, production-grade systems inside a regulated industry.

There are three main reasons why this gap is especially big in fintech.

1. Regulatory Constraints Reshape What 'Working' Means

A model that works well in testing might not be usable in real operations if its decisions can't be explained, audited, or questioned. In fintech, explaining decisions is required for compliance, not just a nice-to-have. Credit decisions, fraud alerts, and KYC results must all be traceable. Because of this, many common AI tools can't be used, and extra engineering work is needed - something most pilot projects miss.

2. Financial Data Is Rarely Ready for AI

Risk data is kept in loan origination systems. Transaction histories are stored on different processing platforms. KYC documents are in tools that aren't built for machines. Thinking that AI can just connect to these systems and work right away is nearly always wrong. Combining, cleaning, and checking data usually accounts for 60 to 70 percent of the entire project.

3. Integration Is an Engineering Problem, Not a Data Science Problem

It's easy to get a model to produce a score. But making sure that score shows up in the right workflow, at the right time, for the right person, with a proper audit trail and linked to the right actions, is a much bigger engineering challenge. Most AI pilots are built by data science teams and then handed off to engineering without a clear integration plan. This often leads to a working model that nobody actually uses.

> Most AI projects fail not because of technology, but because the surrounding system - data, integration, compliance architecture - is not ready to support it.

Framework: The Four Layers of AI Readiness

One helpful way to assess whether a fintech system is ready for AI is to break it down into four layers. Each layer can fail in its own way and has its own requirements. An AI project will only be as strong as its weakest layer.

01	**INPUT** *Data Sources*	Risk records, transaction history, KYC documents, CRM data. AI cannot compensate for fragmented or inconsistent inputs - system readiness begins here. If this layer is weak, nothing downstream will work reliably.
02	**PROCESS** *Data Pipeline*	Cleaning, structuring, enriching, and validating data before it reaches any model. This is where most fintech pilot failures originate - not because the model is wrong, but because the data is not ready.
03	**DECISION** *AI / Logic Layer*	Models and business rules convert structured inputs into predictions, risk scores, or recommendations. This layer performs reliably only when the Input and Process are already sound.
04	**ACTION** *Application Layer*	Outputs delivered into real workflows: automated triggers, dashboards, audit trails, and compliance exports. Without this layer, AI produces insights - not operational value.

Diagnosing Your Current State

The real value of this model is its diagnostic support. Before choosing a use case, check each layer individually. The questions below help clarify this assessment.

Layer	Not Ready - Warning Signs	Ready - What It Looks Like
Input	Data scattered across 3+ systems with no unified schema	Key data sources consolidated and consistently structured
Process	No validation; analysts manually clean data before reviews	Automated pipeline with quality checks and version control
Decision	Black-box scoring; no explanation available per case	Every decision is traceable and auditable for compliance
Action	AI output goes to a report that no one acts on	Results trigger real workflows, alerts, or automated steps

The Sequencing Mistake - and the Correct Path

The most common way fintech AI projects fail is expected. Teams pick a model, build a proof-of-concept in isolation, and only later realize that the data and integration requirements make real deployment much harder than they anticipated. By then, the budget is gone, and patience has run out.

× Common Approach		
Step 1	Select a model	*High enthusiasm, wrong starting point*
Step 2	Build a PoC	*Works in a sandbox, not in production*
Step 3	Discover data issues	*Project stalls, budget consumed*
Step 4	Try to retrofit integration	*Too late - scope has exploded*

✓ Recommended Approach		
Step 1	Map & consolidate data sources	*Slow start, strong foundation*
Step 2	Build a validated pipeline	*Rule-based baseline first*
Step 3	Introduce the model as an enhancement	*Incremental, measurable improvement*
Step 4	Integrate into live workflow	*Real business value, auditable*

Teams that follow the right order start with data consolidation, then pipeline validation, then add the model, and finally, integrate with workflows. In such cases, they get to production faster and achieve more reliable results. The model is usually the least risky part and should come last, not first.

Where AI Creates Measurable Value in Fintech Operations

Not every fintech use case is a good fit for AI at every stage of system maturity. The matrix below shows the most common fintech AI applications based on three things: data complexity, integration effort, and compliance risk. This helps you decide which use case to start with, depending on your system's current state.

> Companies that succeed with AI do not start with the most sophisticated use case. They start with the one where system conditions for success already exist - or can be established within a realistic scope.

In practice, starting with internal operations and reporting is the safest approach. These use cases need less compliance work, the data is usually more organized, and integration is simpler. Teams that build a working AI system for internal use - even a basic one - learn the patterns that make bigger, more valuable projects easier and cheaper to launch later.

Use Case	Data Complexity	Integration Effort	Compliance Risk	Typical First Outcome
Credit Risk Scoring	High	Medium	High	50% faster decisions, -15% default rate
Transaction Fraud Detection	High	High	High	30–50% reduction in false positives
KYC / Onboarding Automation	Medium	Medium	High	60–80% of cases processed without manual review
Cashback & Loyalty Engine	Medium	Medium	Low	+20–30% repeat transaction rate
Internal Reporting & Ops	Low	Low	Low	Analyst hours freed for higher-value work

Use Case #1: Mobile Cashback Feature for E-commerce Platform

A fast-growing e-commerce platform requested to boost customer retention and encourage repeat purchases by introducing a cashback program in its existing mobile app. While the business already had significant traffic and transactions, it lacked a clear way to reward repeat customers.

The main challenge was not simply adding a cashback feature. It was about ensuring it fit smoothly into the existing payment flows, user accounts, and partner merchants, without disrupting daily operations.

Why the Standard Approach Did Not Work

The initial idea was to implement a standalone cashback module or use third-party loyalty solutions.

However, these approaches failed to address key constraints:

- No seamless integration with the existing payment and checkout flow
- Limited flexibility in defining cashback rules across different merchants and campaigns
- Lack of real-time balance updates and transaction tracking
- Poor alignment with mobile UX and performance requirements

The plug-and-play loyalty tool approach did not match the platform's architecture or future growth plans.

What Changed

The project started with restructuring the transaction and reward logic at the system level.

Phase one focused on building a unified cashback engine:

- Centralized transaction tracking across purchases
- Real-time calculation of cashback based on configurable rules
- Secure wallet system for storing and managing user rewards

Phase two integrated the system into the mobile app and partner ecosystem:

- Seamless cashback application during checkout
- Real-time balance updates and notifications
- Flexible campaign management for different merchants and promotions

The system was built to work as part of the core platform, not as an external add-on.

Result

+25-35%	Real-time	Flexible	Seamless
Increase in repeat purchases	Cashback calculation and tracking	Campaign management across merchants	Integration into mobile experience

What This Case Demonstrates

The value did not come from the cashback concept itself, since it is widely used.

The impact came from how it was implemented:

- deeply integrated into transaction flows
- built around real-time data
- aligned with both user experience and business logic

This is what turns a common feature into a measurable growth driver.

Use Case #2: Automated SaaS Lending Risk Assessment Platform

A SaaS lending platform handling thousands of loan applications needed to improve the speed and consistency of its credit risk assessment process.

The existing workflow relied heavily on manual review:

- Data collected from multiple sources (applications, financial records, external APIs)
- Analysts manually verifying and structuring inputs
- Decision-making is dependent on individual expertise

This created bottlenecks, inconsistencies, and limited scalability.

Why the Standard Approach Did Not Work

The company explored off-the-shelf risk scoring tools and generic AI models.

These solutions failed because:

- They did not integrate cleanly with existing data sources
- They required restructuring internal processes around the tool
- They lacked the transparency required for compliance and audit
- They could not reflect the company's internal risk logic

Assuming that a model alone would solve the problem overlooked how complex the data and workflows really were.

What Changed

The project first focused on building a reliable foundation for data and decision-making.

Phase one addressed the data layer:

- Integration of multiple data sources into a unified pipeline
- Automated data validation and normalization
- Creation of a consistent borrower profile structure

Phase two introduced the decision engine:

- Combination of statistical scoring and business rules
- Configurable logic aligned with internal risk policies
- Full audit trail for every decision

The system was developed to support both automation and human oversight.

Result

50%	30-40%	Up to 15%	100%
faster risk assessment time	Processing capacity	In the default rate performance	Auditability and traceability of decisions

What This Case Demonstrates

The breakthrough did not come from a complex AI model.

It came from:

- structuring the data correctly
- building a reliable processing pipeline
- embedding decision logic into real workflows

The same scoring logic that struggled in isolation became effective once the surrounding system was properly engineered.

How We Approach This at Insoftex

When we work with fintech teams, we always start by understanding the state of the four layers in your environment before suggesting any AI approach. Talking about the model comes after that.

In practice, our first step is to diagnose your situation. We look at where your data is stored, how your systems connect, your compliance and audit needs, and which important decisions are still made by hand. From this, we find one or two use cases where the system is already set up for success - or could be, within a reasonable scope - and estimate what a pilot would need in terms of data, engineering, and time.

Most fintech teams already have what they need for a successful AI project. The real question isn't whether to use AI, but which use case matches your system's current state and what needs to be fixed before starting development.

Ready to assess where AI fits in your system?

We offer a focused 30-minute working session - no pitch, no slide deck.

We map your data sources, identify the strongest first use case, and estimate what a realistic pilot looks like.

Book a 30-min AI Readiness Review

usa@insoftex.com | insoftex.com