AI in Asset Finance & Leasing

AI is widely prioritised—but has yet to materially change decision-making

The Implementation Paradox: High Experimentation, Low Impact

Prepared for senior leadership in the asset finance and equipment leasing industry

Christoph Auerbach, Invigors EMEA Ltd

April 2026

EXECUTIVE SUMMARY

The asset finance and equipment leasing industry is caught in an uncomfortable paradox. AI is everywhere in the boardroom — in strategy presentations, vendor pitches, and innovation budgets. Yet in the front office, in the credit committee, and across the asset management lifecycle, decisions and processes are performed much the same way they were a decade ago.

This paper argues that the implementation gap in asset finance is not primarily a technology problem. It is an organisational, cultural, and leadership problem — and the data bears this out. Across the industry, AI projects stall not because the algorithms fail, but because the organisations around them are not designed to learn, decide, or govern differently.

At the same time, a dangerous narrative is gaining ground: that AI can replace human judgment in research, credit assessment, and portfolio management. The evidence suggests the opposite. In a capital-intensive, relationship-driven, and regulation-sensitive industry, AI is a powerful amplifier of human capability, a huge productivity driver — but it is not a substitute for it.

This paper maps the implementation landscape, diagnoses the structural barriers, addresses the generational dynamics that complicate adoption, and offers a practical leadership agenda for closing the gap between AI ambition and AI value.

1. AI is everywhere — but changing almost nothing

1.1 Firms are investing in AI — but not yet using it to change decisions

Ask any leadership team in asset finance whether AI is a priority. The answer is almost always yes. Ask whether AI is meaningfully changing credit decisions, residual value management, or collections strategy. The answer is almost always 'not yet.'

This is the AI implementation paradox, and it is not unique to asset finance — but the industry exemplifies it. According to McKinsey's State of AI 2025 Survey, 88% of organisations now report regular AI use in at least one business function, up from 78% just a year prior. Yet a Gartner survey of finance leaders published in late 2025 found that AI adoption in the finance function has effectively plateaued, with 59% reporting use — the same figure as the year before. The gap between 'using AI somewhere' and 'AI changing outcomes' is wider than most executives are comfortable admitting.

In asset finance and equipment leasing, the paradox is sharper still. The industry handles long-dated, capital-intensive commitments in which information asymmetry, asset expertise, and portfolio judgement are genuine competitive advantages. It should be one of the most fertile environments for AI. Yet across credit underwriting, portfolio monitoring, residual value forecasting, and early default detection, the tools remain largely advisory at best, and decorative at worst.

1.2 What the numbers show

Key industry data points

73% of asset management executives say AI is critical to their organisation's future — but two-thirds report only modest ROI from AI investments, and 12% see no return at all. (ThoughtLab / Grant Thornton, Q3 2025)

Around 70% of failed AI initiatives stem from people- and process-related issues; only 10% from the algorithms themselves. (HES FinTech / industry survey, 2025)

Only 1 in 5 companies has a mature governance model for autonomous AI agents, even as agentic AI is expected to scale rapidly. (Deloitte State of AI in the Enterprise, 2025)

Half of asset management firms lack basic processes to clean, normalise, and tag internal data — limiting AI's effectiveness before it begins. (ThoughtLab, Q3 2025)

The picture is consistent: high aspiration, significant investment, modest measurable impact. The industry is paying full price for experimentation and receiving pilot-level returns. That is not a technology failure. It is a management failure.

Practitioner Perspective

In recent discussions with asset finance leaders across Europe, a consistent pattern emerges. Most can describe their AI roadmap with confidence — the use cases they are piloting, the innovation budget they have secured. Fewer can point to a single credit decision, a pricing call, or a collections action that was materially different because of AI last quarter. When pressed, the answers tend to converge on the same set of explanations: data is not ready, the business case needs further *validation, the right governance structure is still being defined. These are not unreasonable positions. But they have been the answers for three years running. The gap between the roadmap and the reality is the subject of this paper.*

2. Where AI Creates Value in Asset Finance

2.1 The operational use cases are clear

Before diagnosing the implementation gap, it is worth being precise about what AI genuinely offers the industry — and what it does not. Over-promising leads to disillusionment; under-estimating means ceding ground to competitors who invest more deliberately.

2.2 The main applications

Use Case (Examples)	**What AI Does Well**	**Value Unlocked**
Portfolio monitoring	Continuous signal detection across news, filings, macro indicators, sector data	Earlier risk identification; faster response to stress signals
Credit pre-screening	Pattern matching at scale across large application volumes; flagging anomalies	Reduced manual load; consistent first-pass quality
Residual value analytics	Regression and ML models trained on asset lifecycle, usage, channel outcome data	More defensible RV assumptions; reduced end-of-term surprises
Collections prioritisation	Early default prediction from behavioural and payment signals	Higher recovery rates; more targeted intervention
Document processing	Extraction, classification, summarisation from contracts, condition reports, PDFs	Significant time savings in origination and servicing
Fraud detection	Anomaly detection across application and behavioural data patterns	Materially lower fraud losses; real-time flagging
Market intelligence	Scanning, aggregating, and summarising public data at scale	Faster awareness of sector trends, OEM changes, regulatory developments

These are not theoretical capabilities. They are operational today in leading organisations. The question is not whether the technology works — it is whether the organisational conditions exist to deploy it at a level that changes decisions.

2.3 The limits are equally important

The industry narrative often conflates AI's capability to gather and summarise information with its ability to generate insight. They are not the same thing, and the difference is commercially significant.

Publicly available AI works predominantly with public, digitised, widely circulated data. In asset finance, a substantial proportion of the most valuable intelligence is proprietary, relationship-based, and informal — dealer incentive changes, OEM buy-back term adjustments, credit committee risk appetite shifts, informal remarketing dynamics. No model can see what it cannot access.

The second limitation is what might be termed the 'grey goo' problem. AI generates clean, well-structured, confident-sounding content at scale. In research terms, this creates a flood of generic market intelligence that is accurate at a high level but lacks the specificity to drive differentiated decisions. Everyone is reading the same AI-synthesised summaries. Strategic convergence is the inevitable result.

Third, AI is optimised for consensus and statistical frequency. The most commercially important signals in asset finance — subtle shifts in dealer behaviour, early stress in an asset class, informal changes in exit channel dynamics — are precisely those that do not show up in aggregated public data. AI smooths out the edge cases that later become trends.

> *"Information is not insight. AI can tell you what everyone is already talking about. The competitive advantage belongs to those who can interpret what is not being said — and that still requires human experience and industry judgment."*

3. Why Implementation Fails

3.1 The industry is designed for control, not iteration

The asset finance industry's implementation gap is not random. It reflects a set of structural characteristics that systematically impede the conditions AI requires to deliver value. Understanding these barriers precisely is the prerequisite for addressing them.

Asset finance businesses are, by design, optimised for risk control. Credit policies are codified. Decision authorities are defined. Processes are documented. This is entirely rational in an industry where single underwriting errors compound over multi-year contract terms.

AI, however, requires the opposite organisational characteristic: a structured learning loop. Models improve when decisions — good and bad — are captured, reviewed, and fed back into the system. In leasing, the feedback lag is inherent: default outcomes may not crystallise for 18–24 months. But the more fundamental problem is cultural. When a model flags a borderline credit and the credit committee approves it on relationship grounds, the override is rarely captured, attributed, or analysed. The model never learns. Six months later, when the outcome is poor, the organisation concludes that 'the AI wasn't accurate' — a self-fulfilling prophecy caused by the absence of feedback infrastructure.

Organisations built to avoid being wrong are poorly suited to the iterative, experimental, failure-tolerant process that AI capability requires.

3.2 Data is fragmented

Most equipment finance firms have more data than they realise. The challenge is not volume; it is architecture. Customer data sits in the core system. Asset specifications live in spreadsheets. Usage data, where it exists at all, resides with the OEM or asset manager. Remarketing results are stored in a separate platform. Condition reports are PDFs in a shared drive. Maintenance records are on the dealer's system.

AI does not need big data. It needs to be joined-up, decision-grade data: customer profile linked to asset profile linked to contract behaviour linked to exit outcome. The absence of this linkage is not an IT clean-up exercise. It is the single most fundamental constraint on AI's ability to support core decisions — pricing, structuring, residual value, and asset risk — at all.

The diagnostic question is simple: Can your organisation trace a single leased asset cleanly from origination through to remarketing, with usage, maintenance, and exit value in one coherent record? In most mid-market leasing businesses, the honest answer is no.

3.3 AI is treated as an IT project

One of the most consistent and underappreciated barriers to AI implementation is procedural rather than strategic: AI initiatives are routinely classified, governed, and managed as IT projects. This category error has consequences that compound at every stage of delivery.

The IT project model is well-suited to what it was designed for: defined scope, fixed requirements, staged delivery gates, RAID logs, steering committees, and a go-live date that marks the end of the engagement. It works for infrastructure upgrades, system migrations, and software rollouts.

It is structurally unsuited to AI, for a fundamental reason: AI is not a project. It is a capability that is uncertain at the outset, improves iteratively over time, and requires continuous review and adaptation to remain fit for purpose.

3.4 Decision ownership is unclear

When AI is governed as an IT project, the consequences are predictable. Business ownership migrates to the technology function, where accountability for commercial outcomes does not naturally sit. Approval processes designed for capital expenditure and system change are applied to model development, creating delays that are disproportionate to the risk being managed. Change management — the deliberate process of preparing users, decision-makers, and middle management to work differently — is treated as a communications task at go-live rather than a programme strand that runs from day one. Stakeholder management defaults to the IT governance map: architecture review boards, information security sign-off, data protection impact assessments. These are necessary, but they are not sufficient. The credit director whose team will act on the model's output, the collections manager who will adjust their workflow, and the asset manager who will challenge the RV forecast are the stakeholders who determine whether the AI delivers value — and they are rarely in the room until it is too late to change the design.

The change management failure is particularly costly. AI does not just change tools; it changes how decisions are made, who has visibility of information, and where accountability sits. In asset finance — where decision culture is entrenched and experienced professionals have earned their judgment through years of market exposure — this is not a minor adjustment. A credit analyst who has built their authority on personal judgement does not automatically trust a model that arrives pre-packaged from an external vendor with an IT project plan. Without a structured change programme that addresses the 'why' as much as the 'how', resistance is not irrational — it is the rational response to inadequate engagement.

The organisations that make AI work treat it from the outset as a business transformation programme with a technology component — not the reverse. Governance sits with the business. The steering group is led by a C-level business sponsor, not the CIO. Change management is funded as a first-class workstream. Stakeholder engagement begins at problem definition, not at user acceptance testing. And the programme does not close at go-live — it transitions into an ongoing operational rhythm of model review, performance monitoring, and iterative improvement.

AI initiatives in asset finance frequently stall at the threshold between piloting and production for a specific and consistent reason: no single person owns the decision that the AI is intended to support.

The dynamic plays out repeatedly. An AI model is developed to tighten pricing on certain asset classes based on risk signals. Sales resists because it threatens dealer relationships. Credit is unwilling to overrule without more validation. Risk wants further testing. Technology is focused on infrastructure. The model output is 'taken into consideration' and subsequently ignored. The pilot is deemed inconclusive. A new one begins.

Successful implementations are characterised by explicit decision ownership: credit owns the approval logic; risk defines the boundaries and monitoring thresholds; the business owns the outcomes; technology owns the reliability and performance of the models. When that accountability structure is absent, AI stays in pilot mode indefinitely — an expensive way to feel innovative without being so.

Real-Life Example: The Pilot That Never Ended

A mid-size equipment finance business invested in an AI-powered credit pre-screening tool in 2023. The tool was procured, integrated into the origination system, and demonstrated strong predictive accuracy in testing. Two years later, the credit committee was still manually reviewing every application above a modest threshold — using the AI output as a reference point, not as a decision input.

The diagnosis was straightforward. The project had been owned by IT. The head of credit had been consulted at requirements stage but had no formal role in the steering group. No one had defined what ‘acted upon’ meant for a credit recommendation — at what confidence level, for which ticket sizes, with what override documentation. When the model flagged a borderline application differently from the credit manager’s view, the manager’s judgment prevailed, unrecorded. The model never learned. The investment stalled at the pilot threshold.

The AI had not failed. The governance around it had.

3.5 Buying a tool is not the same as adopting capability

A significant proportion of the industry’s AI investment has been absorbed not into capability but into software licenses. The assumption is that buying an AI tool constitutes AI adoption. In practice, what is purchased is typically a sophisticated dashboard that layers on top of unchanged decision processes.

Credit committees continue to rely on static scorecards. Asset managers distrust model forecasts and revert to historical rules of thumb. Collections teams stick to trigger-based outreach. The AI tool provides output that is visible but not acted upon — an expensive way to confirm what the experienced professionals already believed.

The organisations generating measurable value from AI do not start with a tool. They start with a specific decision that is currently made poorly or at unnecessary cost, identify the data required to support that decision more effectively, and deploy a targeted model with clear accountability and a defined KPI. Trust is built through demonstrated results in a bounded context before being extended to adjacent decisions.

3.6 Regulatory is real, but often overstated

Regulatory uncertainty is frequently cited as a primary barrier to AI adoption in asset finance. The concern is legitimate in certain contexts — model explainability requirements under the EU AI Act, fair lending obligations in credit decisioning, and data privacy constraints all represent genuine compliance considerations.

However, a candid assessment of the industry suggests that regulation is often invoked as a barrier before it blocks anything. The EU AI Act's high-risk categories do cover automated credit decisions, and explainability requirements are real. But most high-value AI applications in asset finance — portfolio monitoring, collections prioritisation, document processing, market intelligence, maintenance prediction — are not in the high-risk category. Regulatory complexity is a genuine constraint at the margins. It is not an explanation for the industry-wide implementation gap.

4. The human factor in adoption

4.1 Generational dynamics shape adoption

Leadership literature on AI adoption rarely confronts the generational dynamics directly. This is a mistake.

In asset finance — an industry with a significant proportion of senior professionals who built their expertise before digital infrastructure existed, alongside a younger colleague who have never not had access to AI tools. That creates tension, but also opportunity.

4.2 Scepticism can be an asset

The data on generational AI adoption are unambiguous. Research from the London School of Economics found that 83% of Gen Z and 73% of Millennials use AI at work, compared to 60% of Gen X and 52% of Baby Boomers. Among younger workers, AI is increasingly the first port of call for information, analysis, and decision support. For a meaningful proportion of senior professionals in asset finance, it remains a tool viewed with scepticism about accuracy and appropriate scope.

This gap creates operational friction. Younger analysts produce AI-assisted research, summaries, and analysis at speeds that experienced professionals find either impressive or concerning, depending on their disposition. Senior leaders instinctively interrogate the provenance of conclusions that arrived faster than they should have. In many organisations, this dynamic suppresses AI's potential: the outputs are generated but not trusted, and the trust gap is never systematically addressed.

4.3 The Under-appreciated Value of Experience

Here a counterintuitive but important point deserves direct articulation: the scepticism that experienced professionals bring to AI is not a liability. It is a capability.
Older professionals — who came of age reading research reports in full, building sector relationships over decades, and developing analytical judgment through trial and error — are better equipped to audit AI outputs. Research published in early 2026 confirmed that productivity and judgement in knowledge-intensive roles peak decades after initial qualification. Experienced professionals in asset finance are, by construction, better positioned to identify when an AI-generated residual value forecast is plausible and when it is dangerously wrong, when a credit model is missing a structural risk, or when a market intelligence summary has smoothed over a signal that matters.

The organisation that deploys AI 'natively' — with younger staff generating outputs and senior staff treating them as a starting point rather than a conclusion — is structurally closer to the optimal model than one where scepticism simply blocks adoption. The goal is not to resolve the generational divide but to design workflows that harness both capabilities.

4.4 Designing for the Full Workforce: The AI Powered Organisation

The most effective AI implementations in financial services are those where the roles are explicit: AI handles scanning, monitoring, summarising, and pattern detection; humans handle sensemaking, prioritisation, interpretation, and decision framing. This is not a temporary transitional arrangement — it is the sustainable operating model.

Organisations that try to automate both sides produce shallow intelligence. The AI-generated insight is technically accurate but contextually unusable — it does not account for the firm's specific asset mix, portfolio vintage, geographic exposure, or dealer relationship dynamics. The experienced professional who could contextualise it has been bypassed or sidelined. The result is well-informed paralysis.

The AI – Human Division of Labour in Asset Finance
AI is best positioned for:
Scanning thousands of data sources for early risk signals | Summarising long documents and market reports | Detecting patterns across large portfolios | Flagging anomalies in application and behavioural data
Humans are indispensable for: Interpreting signals in the context of a specific portfolio | Accessing relationship-based and informal intelligence | Challenging comfortable AI-generated narratives | Making and owning capital allocation decisions | Sensing what is not in the data

5. The Research Illusion

5.1 AI enabled research is powerful, but limited

One of the most commercially consequential misconceptions currently circulating in the industry is this: with AI, we can now produce market intelligence ourselves. The implication is that high-quality, decision-relevant research can be generated on-demand, at low cost, without specialist knowledge or relationships.
There is enough truth in this to make it dangerous.

AI is genuinely transformative in its ability to aggregate, synthesise, and surface publicly available information at speed. Competitive intelligence at a high level, regulatory tracking, OEM announcement monitoring, macro trend identification — these tasks that once required significant analyst time are now executable in minutes.
The problem is that this is precisely the research that everyone in the industry is also producing. When market intelligence is derived from the same public sources, synthesised by similar models, and presented in similarly structured outputs, strategic differentiation disappears. Every firm sees the same trends, draws similar conclusions, and positions itself in the same direction. This is not intelligence. It is informed conformity.

5.2 What AI cannot reach

Meaningful intelligence in asset finance has three components that AI structurally cannot access. The first is proprietary and relationship-based data. Dealer incentive changes, OEM buy-back term adjustments, informal shifts in credit appetite at competitor firms, remarketing channel dynamics, and condition assessment practices that vary by asset manager are not in any public data set. They exist in conversations, in relationships, and in the accumulated experiential pattern-recognition of practitioners who have operated in the market across multiple cycles.

The second is what is not being said. AI is fundamentally retrospective — it analyses what has been published, filed, and digitised. The most important risk signals in asset finance often precede publication by months. An experienced practitioner reading the room at an OEM conference or observing changed behaviour in a key dealer will identify emerging stress before it registers in any data source that AI can reach.

The third is contextual relevance. A well-structured AI-generated summary of residual value pressure in a particular asset class may be entirely accurate at the market level. Whether it applies to a specific firm's portfolio — given its vintage profile, geographic concentration, dealer mix, and remarketing relationships — requires judgment that no model can substitute for.

5.3 False clarity is the risk

Perhaps the most under-appreciated risk associated with AI-generated research is not misinformation but false clarity. AI produces outputs that are clean, well-structured, and confident in tone. In a senior leadership context, this creates a cognitive trap: the output feels decision-ready, the analysis appears comprehensive, and the recommendations seem logical. The discomfort of uncertainty — which is often the appropriate response to genuinely ambiguous market conditions — is replaced by the comfort of apparent resolution.

Organisations that act on AI-generated market intelligence without subjecting it to experienced human interpretation risk the worst of both worlds: the speed of automated analysis combined with the confidence of expert judgment, without actually having expert judgment in the loop.

6. The Strategic Stakes: Why This Is a Leadership Issue

6.1 The implementation gap has strategic consequences

The implementation gap in AI is not a technical challenge awaiting a software solution. It is a leadership challenge that requires deliberate organisational choices. The firms that close the gap will compound their competitive advantage over those that remain in perpetual pilot mode.

6.2 Three Competitive Risks of the Status Quo

Risk	Consequence for Asset Finance
Strategic convergence	When all firms use the same AI-synthesised market intelligence, differentiation erodes. Pricing converges. Portfolio positioning aligns. The firms that access non-AI-reachable intelligence — through relationships, expertise, and physical market presence — gain disproportionate advantage.
Delayed risk recognition	AI confirms trends after they are visible in public data. Experienced professionals with sector relationships detect signals earlier. Organisations that over-automate research and under-invest in experienced practitioners will systematically lag on risk identification.
Decision paralysis	High AI output volume without structured human interpretation produces information overload. The organisation knows more but decides less confidently. Speed of action — in credit, in collections, in asset management — declines despite higher information investment.

6.2 What Separates Leaders from Laggards

The organisations generating material value from AI in asset finance are distinguishable by several consistent characteristics. They do not start with technology; they start with a specific decision and work backwards to the data and tools required. They invest in data infrastructure as a strategic asset, not an IT maintenance task. They designate clear ownership of AI-driven decisions at the business level, not the technology level. They build feedback loops into the process architecture so that models improve over time. And they treat experienced practitioners not as obstacles to AI adoption but as the critical layer of judgment without which AI outputs are commercially inert.

Crucially, these organisations also govern their AI deliberately. As agentic AI — models that take autonomous sequential actions — begins to enter enterprise workflows, the governance question becomes urgent. Deloitte’s 2025 enterprise AI survey found that only one in five companies has a mature governance model for autonomous AI agents. In an industry characterised by long-dated commitments and regulatory accountability, the organisations that build governance capability now will be better positioned to scale AI safely as its autonomy increases.

7. A Leadership Agenda for Closing the Gap

7.1 Start with one decision

The following agenda is not a technology roadmap. It is an organisational and leadership framework. The tools are available. The question is whether the conditions exist to use them.

The most reliable path to demonstrable AI value is to identify a single high-cost, high-frequency decision that is currently made with insufficient data or inconsistent quality, and to build a focused AI capability around that decision. In asset finance, natural candidates include early default detection in the SME portfolio, residual value refresh for specific high-volume asset classes, or collections prioritisation based on early behavioural signals.

Define the decision clearly. Identify the data required. Assign business ownership. Set a specific KPI. Deliver a result. Trust is built through demonstrated performance in a bounded context, not through platform-level investments that promise transformation but deliver dashboards.

Real-Life Example: Starting Small, Scaling Deliberately

A manufacturer-affiliated captive finance operation in Central Europe, supporting dealer networks across four countries, faced chronic inconsistency in its early-stage collections process: intervention timing varied by team, recovery rates differed significantly by region, and experienced collectors were spending material time on accounts that would self-cure without intervention.

Rather than pursuing a full AI transformation programme, the head of collections defined a single question: which accounts in the first 30 days of delinquency are genuinely at risk, versus those that will cure without intervention? A focused model was built using three years of payment history, asset type, dealer profile, and customer segment data. It was piloted in one country for 90 days. Recovery rates in that region improved by 18%. Collector time on self-cure accounts fell by 30%.

The model was then extended to a second country, refined, and extended again. Two years later, the same data and governance architecture underpins early-stage credit risk as well — not because the organisation set out to build a broad AI capability, but because a single successful decision built the trust and infrastructure to support the next one. Scope was a feature, not a limitation.

7.2 Build decision-grade data infrastructure

Data readiness is the single most under-invested capability in the industry. The investment case is straightforward: without joined-up data linking customer, asset, contract, behaviour, and outcome, AI cannot support core decisions. The cost of data remediation is a fixed investment. The cost of continued data fragmentation is an indefinite constraint on competitive capability.

The organisational intervention required is not primarily technical. It is governance: defining who owns data quality for specific decisions, establishing accountability for data completeness and linkage, and treating the data asset as a boardroom priority rather than an IT function.

7.3 Assign ownership and govern properly

Every AI-influenced decision in the organisation requires a named business owner who is accountable for the outcome, a risk owner who defines the boundaries and monitoring requirements, and a technical owner who is responsible for model performance and reliability. Without this triangular structure, AI outputs will be produced, reviewed, and disregarded.
The hardest conversation — and the most important one — is about who owns the outcome when AI is wrong. Organisations that avoid this conversation will never move AI from advisory to operational.

7.4 Govern the programme as business transformation, not an IT delivery

The governance model applied to an AI initiative determines, more than any other single factor, whether it reaches production or stalls in perpetual pilot. The default — classifying AI as an IT project and applying the standard delivery framework — is one of the most reliable ways to guarantee the latter.

AI programmes require a fundamentally different governance architecture. The programme sponsor should be a C-level business executive, not the CIO or CTO. Steering group membership should be weighted towards the business functions whose decisions the AI will influence — credit, risk, asset management, collections — with technology represented but not leading. Programme reporting should be framed around decision quality and commercial outcome, not delivery milestones and system uptime.

Change management must be treated as a first-class programme workstream, funded and resourced from day one — not retrofitted at go-live as a communications exercise. In asset finance, where experienced professionals have built their authority on personal judgment, the introduction of model-driven decision support is a significant organisational change. It requires structured engagement with the people whose working practices will change, honest conversation about what the model can and cannot do, and a transition path that builds trust through demonstrated performance rather than mandating adoption by decree.

Stakeholder management equally requires a different map. The standard IT governance stakeholders — architecture review, information security, data protection — are necessary but insufficient. The credit director who will act on the model's output, the collections manager who will restructure their team's workflow, and the asset manager who will challenge the residual value forecast are the stakeholders who determine whether the AI delivers value. Engaging them at user acceptance testing is too late. They need to be in the room at problem definition, involved in data design, and given genuine influence over how the model's outputs are surfaced and used.
Finally, the programme should not close at go-live. AI capability is not a system that is installed and maintained — it is a capability that is built, reviewed, and improved continuously. The governance structure that delivers the initial deployment should transition directly into an operational rhythm: regular model performance reviews, scheduled retraining cycles, defined escalation paths when model confidence is low, and a standing forum where business, risk, and technical ownership meet to assess whether the AI is still fit for the decisions it is supporting.

7.5 Design the Human-AI Workflow Deliberately

The division of labour between AI and experienced practitioners should be explicit, documented, and reviewed. AI handles scanning, monitoring, summarising, and first-pass pattern detection. Experienced professionals handle interpretation, contextualisation, exception judgement, and decision accountability.

This is not a transitional arrangement pending further AI capability. It is the appropriate permanent model for decisions that are capital-intensive, long-dated, and relationship-sensitive. Designing workflows that make this division explicit — rather than leaving it to emerge from informal practice — is a leadership responsibility.

7.6 Invest in Learning Loops, Not Just Models

A model that does not receive feedback does not improve. In asset finance, building the feedback infrastructure requires deliberate process design: capturing when model recommendations are overridden and why; tracking outcomes against model predictions over the relevant horizon; reviewing model performance at defined intervals with business, risk, and technical ownership present.

This investment is more organisational than technical. The technology for feedback capture is straightforward. The discipline of using it, and the cultural willingness to acknowledge when a model was right and a human override was wrong, is what most organisations lack.

7.7 Govern AI as a Strategic Asset

AI governance in asset finance must move from compliance theatre to genuine risk management. This means defining where AI operates in an advisory capacity and where it has decision influence; establishing model monitoring and periodic review as standard; maintaining human accountability for all consequential decisions; and building the organisational capability to audit AI outputs before they are acted upon.

As AI capabilities advance — and agentic AI in particular begins to automate sequences of decisions rather than single outputs — the organisations with governance infrastructure in place will be able to scale safely. Those without it will face a more difficult regulatory and reputational environment as AI's footprint expands.

8. Looking Ahead: The Next 24 Months

The competitive landscape in asset finance will be shaped materially by AI over the next two to three years — not by the technology itself, but by the quality of organisational adoption. Several dynamics deserve specific leadership attention.

8.1 The gap will widen

Organisations that have invested in data infrastructure and built genuine AI decision capability are already compounding their advantage. McKinsey data indicates that organisations with mature AI practices see labour productivity grow nearly five times faster than average. In asset finance terms, this translates to faster, more accurate credit decisions, more defensible residual value assumptions, earlier risk detection, and lower cost-to-serve in servicing and collections. The gap between AI-capable and AI-aspiring firms will widen faster than most leadership teams currently anticipate.

8.2 Agentic AI: Approaching with eyes open

Agentic AI — models capable of taking autonomous, sequential actions across multiple steps with limited human intervention — is moving from experimentation to early deployment in financial services. Its potential applications in asset finance are significant: automated covenant monitoring, proactive collections outreach, dynamic repricing triggers, and maintenance scheduling in fleet contexts.

The governance imperative is equally significant. An agent that takes a series of wrong actions in a portfolio management context does not produce a single correctable error — it produces a compounding sequence. The organisations that approach agentic AI with clear accountability structures, defined intervention thresholds, and audit trails will capture the capability without the risk. Those that do not will discover the hard way why human oversight is not optional.

8.3 Talent will matter more, not less

As AI handles more of the analytical and monitoring workload, the nature of required human expertise in asset finance will shift. The premium will move from data gathering and processing capability to interpretive judgment, domain expertise, and the ability to interrogate AI outputs critically.

This has direct implications for talent strategy. The industry should be cautious about interpreting AI efficiency gains as a mandate to reduce experienced headcount. The experienced practitioner who can identify when an AI-generated RV forecast is wrong, when a credit model is missing a structural risk, or when a market intelligence summary has smoothed over an important signal, is not replaceable by the tool that produced the output. The human in the loop is not a legacy cost. It is the differentiator.

9. The Invigors Perspective: From Experimentation to Operational Impact

9.1 Where organisations get stuck

Invigors works with asset finance and leasing organisations at the intersection of commercial strategy, operational capability, and organisational performance. We see the implementation gap described in this paper not as an abstract industry problem, but as a specific challenge that plays out in recognisable ways across leadership teams, governance structures, and deployment programmes.

Our perspective is built on a consistent observation: the organisations that make AI work in asset finance are not necessarily the most technically sophisticated or the best resourced. They are the ones that treat AI as a business transformation challenge — and bring the same commercial discipline to AI deployment that they bring to credit, pricing, and portfolio management decisions.

Most of the asset finance organisations we engage with are not at the start of their AI journey. They have run pilots. They have procured tools. They have AI referenced in their strategy. What they have not yet done is make AI change a decision in production — consistently, measurably, and at scale.

The sticking points are predictable: strategy without a specific decision to anchor it; data that is not joined up in the way AI requires; governance that has defaulted to IT project management; change programmes that were not funded; and ownership gaps that mean model outputs are noted but not acted upon. These are not technology problems. They are the organisational problems this paper has described — and they are solvable with the right combination of industry expertise and implementation rigour.

9.2 How Invigors Supports the Journey

We work with clients at three points where the implementation gap is most felt.

The first is strategy and prioritisation. We help leadership teams identify which decisions — in credit, asset management, collections, or portfolio monitoring — are most likely to yield measurable value, given their specific data architecture, organisational structure, and competitive context. This diagnostic work is the foundation for everything that follows. Without it, AI investment is undirected.

The second is programme design and governance. We help organisation's structure AI initiatives as business transformation programmes — with C-level business sponsorship, decision-oriented success metrics, change management as a funded workstream, and the right stakeholder architecture from day one. This is where most implementations either accelerate or stall, and where the combination of deep asset finance expertise and change management capability matters most.

The third is human-AI workflow design. Deploying a model is not the same as changing how decisions are made. We work with the credit teams, asset managers, and collections leaders who will use AI outputs in their daily work, designing the processes, escalation paths, and governance rhythms that make adoption durable rather than mandated.

9.3 What engagement looks like

Engagements with Invigors typically begin with a diagnostic: an honest assessment of where an organisation sits on the journey from experimentation to operational impact, what the most material decision-level opportunities are, and what the specific barriers need to be addressed before AI can deliver value. This is a practical, commercially grounded conversation — not a technology assessment.

From there, engagements are scoped to the specific need: strategy and prioritisation, programme design, change management, or ongoing advisory as the programme moves from pilot to production. We work alongside management teams, not as a substitute for them. The fundamental question Invigors helps organisations answer is not 'What AI should we buy?' It is 'What decisions do we need to make differently, and what does it take to get there?'

10. Conclusion: The Question Is Leadership, Not Technology

The asset finance and equipment leasing industry has spent several years experimenting with AI. The technology has been validated. The use cases are established. The tools are available, capable, and increasingly accessible. The gap is not in the AI.

The gap is in the organisational conditions required to deploy it: clean, joined-up data; clear decision ownership; feedback loops that allow models to improve; workflows that combine AI capability with experienced human judgment; and governance that keeps accountability with people, not algorithms.

These are leadership problems. They require leadership solutions: deliberate choices about organisational design, explicit accountability structures, cultural willingness to learn from AI-assisted failures, and the strategic confidence to move from pilots to production.

The organisations that close the implementation gap will do so not by buying better tools but by making better decisions about how AI fits into their decision architecture — and by refusing to allow the comfort of experimentation to substitute for the discipline of deployment.

Three questions for senior leaders:

1. Which specific decisions in our organisation are meaningfully different because of AI today — not just informed by it, but different?
2. Who in this organisation is accountable for the outcome when an AI-influenced decision is wrong?
3. Can we trace a single leased asset from origination to remarketing in a single, clean, joined-up data record?

If the answers are uncomfortable, the problem is not the AI. The work begins at the top.

Where Invigors helps

- Identify the decisions that matter
- Build decision-grade data
- Turn AI from pilot to production

Sources & References

McKinsey & Company, State of AI 2025 Survey • Gartner, Finance AI Adoption Survey (Nov 2025) • Deloitte, State of AI in the Enterprise 2026 • ThoughtLab / Grant Thornton, The AI-Powered Investment Firm (Q3 2025) • HES FinTech, AI in Lending: Implementation Challenges (2025) • London School of Economics, Generational AI Adoption at Work (2025) • The Register, Study Confirms Experience Beats Youthful Enthusiasm (Feb 2026) • Financial Stability Board, Monitoring AI Adoption and Vulnerabilities in the Financial Sector (Oct 2025) • OECD, AI Adoption by SMEs (Dec 2025) • BCG, AI Efficiency Gains in Financial Services (2024) • World Economic Forum, AI Talent Scarcity in Financial Services (2024)