



WHITE PAPER

Why AI Without Domain Expertise Falls Short

Turning AI into durable value through domain expertise

MAY 2026


By
Chitra Baskar

The AI arms race moves up the stack

The AI arms race in asset management has moved decisively beyond experimentation. Foundation model providers are no longer content to sit beneath the application layer as underlying infrastructure. They are making deliberate moves up the stack: into workflow definition, into operational influence, and ultimately into decision-making authority within financial institutions.

Anthropic's recent product direction illustrates the shift clearly. Claude for Financial Services launched with pre-built capabilities for DCF modelling, comparable company analysis, and earnings research, integrated directly with data providers such as S&P Global, FactSet, and Morningstar. Cowork plugins now package finance, legal, and compliance workflows directly into the platform. Opus 4.6 was benchmarked explicitly on financial research tasks. These are not incremental enhancements. They reflect an ambition to shape how work gets done across the industry.

That ambition is understandable. But the assumption that platform capability alone translates into operational transformation is where expectations need to be recalibrated.



Early progress is not the same as full transformation.

AI is already delivering measurable value in parts of the industry, but its impact remains uneven and contained. Many firms are using large language models to accelerate earnings call analysis, draft first-pass LP communications, or assist with DDQ responses and regulatory filings. These use cases generate real efficiency gains and are relatively straightforward to deploy.

What they do not yet represent is a fundamental redesign of the fund operating model. The distance between a compelling demonstration and a production-grade system embedded into day-to-day operations remains significant. Most funds continue to run AI pilots alongside existing processes rather than replace them. The narrative that entire operational functions are being compressed into minutes obscures this reality. Where AI is working well today, it is doing so in tightly scoped contexts. The most effective adopters are those approaching the technology with discipline, examining not only what it can do in ideal conditions, but how it behaves when precision, accountability, and edge-case handling matter most.

Platform strength is not operational understanding



Anthropic has built a capable and thoughtfully engineered platform. Claude for Financial Services is a credible offering for research-oriented workflows. The integrations with Snowflake, Databricks, and major market data providers are robust. The pre-built tools for coverage initiation, due diligence data packs, and comparable company analysis deliver genuine productivity gains in the right settings. When Anthropic says it can compress hours of analyst work into minutes, that is not hyperbole for a well-scoped research task.

But there is a meaningful difference between a platform that can analyse public filings and one that understands how a specific fund actually operates. A model can generate a DCF. It cannot know that Fund III's waterfall includes a non-standard catch-up provision negotiated via side letter with an anchor LP. It can draft an LP update. It cannot reconcile conflicting reporting obligations imposed by different institutional investors. It can parse an ISDA Master Agreement. It cannot recognise that a counterparty's CSA was amended last quarter in a way that materially alters eligible collateral for a particular netting set.

These are not shortcomings of model quality. They are the inevitable consequence of serving an industry defined by structural variation across funds, strategies, jurisdictions, and counterparties.



What generic AI can do



Generate a DCF on public filings



Draft a first-pass LP update



Parse an ISDA Master Agreement



What it cannot do

1

Know that Fund III's waterfall includes a non-standard catch-up provision negotiated via side letter


2

Reconcile conflicting reporting obligations imposed by different institutional investors

3

Recognise that a counterparty's CSA was amended last quarter, altering eligible collateral for a netting set

Generic software has never solved domain-specific problems AI will not be the exception.



Enterprise software history makes this pattern familiar. ERP platforms promised standardised operations across industries, only to encounter the irreducible realities of specialisation. Pharmaceutical supply chains, with cold-chain logistics and regulatory serialisation, differed fundamentally from oil and gas upstream operations. Manufacturing scheduling bore little resemblance to retail demand planning.

ERP systems delivered value only when deep industry expertise was layered on top, when practitioners adapted technology to reflect how work actually happened. Even then, many implementations failed to meet their original objectives.

The constraint was never computing power. It was context.

AI is following the same trajectory, with an important nuance. Not every process should be AI-native. There is a useful distinction, articulated well in a recent article published by IBM on AI-native design, between systems where AI is a core architectural component and systems where deterministic logic remains the right tool for the job. Math is math. NAV calculations, waterfall distributions, and fee computations demand precision and auditability above all else. In these domains, AI's role is not to replace deterministic logic but to orchestrate the workflow around it: data retrieval, exception handling, approvals, and the explanatory narrative that makes outputs actionable.

Not every process should be AI-native



Math is Math
NAV, waterfalls, fees
demand precision and
auditability



AI's Real Role
Orchestrate around deterministic logic: data
retrieval, exceptions, approvals, explanatory
narrative.

The organisations gaining real leverage are designing hybrid architectures. AI where judgment adds value. Deterministic systems where accuracy is non-negotiable. The connective tissue between the two is where the operational edge lives.

Domain expertise is not a layer. It is the foundation.



A generic large language model does not understand what it means to run a fund. It has no intrinsic grasp of waterfall mechanics, where a misapplied hurdle rate or catch-up provision can misallocate millions. It cannot experience the operational intensity of a month-end close, where fund accountants reconcile across multiple administrators, prime brokers, and custodians under time pressure that makes precision and speed equally critical.

Consider OTC derivatives workflows. A desk managing these positions faces a daily sequence that is deceptively complex: margin calls calculated against live portfolio valuations, collateral identified and mobilised from eligible pools that vary by counterparty agreement, and wire instructions organised and validated against standing settlement instructions, all before a daily deadline.

A generic AI tool might read a margin call notice. Without domain context, it cannot determine whether a disputed threshold applies, whether proposed collateral meets the relevant CSA, or whether settlement aligns with the custodian's cut-off windows. This is where agentic workflows show real promise, not as autonomous decision-makers, but as orchestration layers grounded in deeply encoded operational knowledge built by practitioners over decades.

The same principle applies across collateral management more broadly. Every counterparty agreement differs. Every fund structure introduces unique account hierarchies. Every jurisdiction adds regulatory nuance around segregation and rehypothecation. The more complex the document, the more essential domain expertise becomes to deriving any value from it.

This is the gap that no amount of model capability improvement will close on its own. The proprietary knowledge that defines how a firm operates, including its workflows, its exception-handling logic, and its institutional memory, lives nowhere on the public internet. It lives in the heads of operations teams, in annotated spreadsheets, in process documents that have not been updated since 2019. The firms seeing genuine operational uplift are grounding their AI in that institutional context, building carefully for the variation by client, by strategy, and by jurisdiction. That investment in specificity is what separates working deployments from expensive pilots.

Agentic workflows show real promise – not as autonomous decision-makers, but as orchestration layers grounded in deeply encoded operational knowledge.

The platform play has limits the platforms cannot solve



Model providers recognise that raw capability is not enough. Anthropic and OpenAI have both expanded into consulting and partnered with large systems integrators such as Deloitte, KPMG, PwC, and Infosys to embed firm-specific knowledge bases and operational logic directly into their models. Anthropic's own leadership has openly acknowledged that closing the gap between a model that works in a demo and one that works in a regulated industry requires domain expertise.

Yet a structural limitation remains. These teams excel at AI, cloud architecture, and data engineering. They are not, by default, experts in fund administration, derivatives operations, or the regulatory mechanics of complex cross-border fund structures.

Translating model capability into operational value requires practitioners who have lived inside these workflows, who understand why exceptions exist, and who know how processes fail under real-world pressure.

The data reality

There is also a data reality the platform model cannot easily address. The most valuable operational intelligence in alternative asset management, covering fund structures, LP terms, proprietary investment processes, and counterparty relationships, is also the most sensitive. Externalising this knowledge at scale introduces governance risk and erodes competitive advantage. Durable value will accrue to firms embedding this intelligence locally, within their own infrastructure and governance framework, guided by domain specialists who understand both the technology and the operational reality well enough to build what this industry actually requires.



Governance Risk

Externalising sensitive intelligence at scale introduces exposure



Competitive Erosion

Proprietary edge is diluted when shared with shared platforms



Local Embedding

Durable value accrues to firms keeping intelligence within their own infrastructure

The right tool, without the right hands, is just another tool

Three pillars of sustained AI value



PILLAR 01

Domain Expertise

Grounded in how the industry actually operates



PILLAR 02

Hybrid Architecture

AI for judgment, deterministic systems for precision



PILLAR 03

Rigorous Governance

Same discipline applied to investment risk

The firms that define the next decade in alternative asset management will not be those that adopted AI first. They will be the ones that implemented it thoughtfully: grounded in domain expertise, integrated alongside deterministic systems where precision demands it, and governed with the same rigour applied to investment risk.

Anthropic, OpenAI, and the broader AI ecosystem are building powerful platforms. Platforms are not solutions. The technology will continue to advance. Whether it delivers sustained operational value depends entirely on the expertise guiding its deployment.

"What Is AI Native?" IBM Think, February 2026. <https://www.ibm.com/think/topics/ai-native>