Why Financial Services
Isn’t Getting Value From AI

Medan Gabbay,

CEO of Quod Financial

04 June 2026

Most firms have spent the last two years putting AI into their business and not getting much back. Productivity gains are marginal, token costs keep climbing, and the workflows people thought they could automate still need humans at every step.

The issue is architectural, not technological. The wrong architecture won’t just cost you productivity. It will cost you the business.

Human process, human speed

Cars were invented to solve a hygiene problem. Cities in the late 1800s were drowning in horse manure and the odour was a public health crisis. Speed wasn’t really the point. The point was getting rid of horses.

But once cars existed, the second-order effects were enormous. Suburbia, the highway system, drive-through retail, the entire shape of post-war American life. None of that was imagined when the car was being designed to replace the horse.

Imagine if cars had been built to behave like horses. Stop every ten miles to rest. Refuel with hay. Need a stable. Most of what made cars transformative would have been thrown away.

That is exactly what most firms are doing with AI. They are taking a human process, slotting AI into one or two steps, and running the whole thing at human speed inside human-shaped workflows. They are using a car like a horse.

AI as assistance saves time on specific tasks. It does not change the business. The productivity gains people expected aren’t showing up because the architecture is wrong.

Cost and the token economy

The economics of frontier AI are not what they appear. Major model providers are quietly compressing what you get per prompt, because compute economics are unsustainable at current usage patterns. Costs are rising even when prices look flat.

Firms running everything through frontier models are about to feel that. The ones that designed for it from the start will be fine. The ones that didn’t will face a choice between rising bills and ripping out their architecture.

A multi-model abstraction layer isn’t optional, it’s a cost survival mechanism. Frontier models for the hardest reasoning, lower cost or open source models for everything else, with the architecture routing between them based on the task. Anything else is a slow-motion disenfranchisement from the technology you’ve built your business on.

Probabilistic versus deterministic

Run the same workflow ten times and you get five, six, seven different results. That’s fine for ideation. It’s not fine for execution, risk, or the analysis feeding into those decisions.

You can’t run a trading desk on a model that maybe gets it right.

The fix isn’t a better model. The fix is architecture. The model can stay exploratory and non-

deterministic in how it reaches its answer. What matters is that the output passed downstream is deterministic, the workflow is bounded, and the action that gets taken is reproducible.

Flexible versus efficient

AI is flexible and self-programmable. It can adapt, reason over ambiguity, generate code, navigate problems it has never seen before. That's the Brain.

It's also slow, expensive and computationally inefficient compared to purpose-built software. Asking AI to handle a high-frequency action is like asking a CEO to staple invoices. They actually are not only a waste, but inefficient and error (hallucination) prone.

The right design uses each for what it's good at. AI handles the parts that need flexibility. The software handles the parts that need efficiency and performance. Doing 1M transactions per minute. That's the Muscle.

Most vendors are doing the opposite, throwing AI at problems where deterministic software would be faster, cheaper and safer.

AI is Logic. Software is Action.

This is the architectural principle the rest of this piece is built on. Agents are good at reasoning, navigating ambiguity, and making human-like logical jumps. Trading software is fast, deterministic, and contains years of accumulated know-how about how markets, venues, regulation and clients actually behave.

AI orchestrates. Software acts.

Smaller decisions for the AI, with tighter scope and predictable handoffs. The model reasons over the problem and delegates the high-performance actions to the deterministic trading software. That changes the cost, the control and the reliability of the whole system. It is also the only architecture that lets you swap models in and out without rebuilding the business logic underneath.

Architecture is survival

Cars didn't just replace horses. They redrew cities, rewrote retail, reshaped how people work and live. Nobody designing the early Model T was thinking about the mall or the school run.

AI is the same kind of shift, the cognitive shift. The Urbanism you adopt now will determine what your business looks like in ten years, what it can do, what it can't, and whether it still exists. The new game is different. Features get built faster, because AI plus a flexible architecture delivers an exponential pace of capability. Internal builds that previously took quarters can now be delivered in weeks, end to end, agentically. The advantage compounds.

Buyers used to pick vendors by checklist. The new question is whether your platform lets you build things you haven't even imagined yet.

Where Quod sits

This isn't a roadmap for us, the infrastructure exists. Twenty years of mature trading software, normalised data, microservices around it, and the ability to introduce AI at any point in any workflow with full visibility and control.

Customers can build their own micro-apps on top. An options RFQ workflow, a custom analytics layer, anything they need. They can't bypass the determinism because everything routes through the trading software. The guardrails compliance sets still hold. The audit trail still works. The IP stays inside the perimeter.

Agentic is logic. Software is action. We've spent twenty years building the action layer. Now AI sits on top of it. That's the architecture financial services needs. The firms that get this right will look unrecognisable in ten years. The firms that don't won't be here.

+1 646 494 1121 | +44 20 7997 7020 | info@quodfinancial.com
www.quodfinancial.com