



CIONET

DISCUSSION SUMMARY

AGENTIC AI IN FINANCIAL SERVICES

/thoughtworks

Agentic AI in financial services

Your next transformation engine or your biggest compliance headache?

Agentic AI is a rapidly evolving technology with profound implications for banking, financial services, and insurance. Today, a number of financial services organisations are beginning to evaluate and experiment with these autonomous systems capable of planning, reasoning, and executing complex tasks with minimal human intervention.

To better understand industry sentiment – and to assess the barriers and opportunities on the path to adoption – CIONET, in partnership with Thoughtworks, brought together a dozen IT leaders for an in-depth discussion in late November. These are the key takeaways:

What is Agentic, anyway?

Definitions often differ when new technologies emerge. This appears the case with Agentic AI. Some stretch the definition so they can rebadge existing solutions as something new and shiny ('Agentic-washing', is emerging as common coinage). Others are legitimately working out where the boundary is between generative AI (GenAI) and Agentic, and between Agentic and existing AI assistants.

The consensus around that table is that Agentic AI is "a more outcomes-focused system". In the words of one attendee, organisations are using existing AI to do their jobs as they currently do them, just more efficiently. Agentic, meanwhile, "let's us reimagine how to do those jobs". If the process needs changing, this technology can help. Another, alluding to its autonomous capability, said simply of Agentic AI: "It will work it out."

Others have pointed out that Agentic AI differs from conventional AI assistants in three key ways. First, Agentic AI allows for more complex workflows. Second, it requires less supervision. Third, it can be embedded across the enterprise, end-to-end.

Agentic AI adoption? It's early

While financial services firms are actively deploying GenAI-based applications in large numbers, the same is not true of Agentic AI. Not yet, at least. While some are beyond the proof of concept stage – and a smaller group are using the technology internally – most are in an earlier stage of experimentation.

That is certainly the conclusion based on the straw poll of senior IT leaders at this event. Asked to rate their progress on scale of zero to 10 (where zero is yet to begin and 10 is where the technology is fully embedded), no firm rated their progress higher than four, many much lower.

To illustrate contrasting adoption patterns, one voice around the table offered his experience talking to two organisations. The first has rebuilt its complaints process end-to-end using Agentic AI. It features dozens of guardrails, a handful of large language models (LLMs), and an additional LLM to review and audit the other models. In contrast, the second firm has employed a sizeable number of highly-trained post-graduates but has, so far, made little progress. "They've got 150 PhDs but business see little value."

Agentic AI use cases

Despite this, attendees shared a small handful of Agentic use cases that have progressed beyond proof of concept. One major bank has embarked on an enterprise architecture use case. The prevalence of legacy systems has left the firm with a lack of documentation which, in turn, leaves it lacking oversight. By applying "six agents and lots of guardrails" to the problem, it is back-filling documentation, including architecture diagrams. It is also ingesting policies and requirements. In doing so, the bank can assess what is – and what is not – compliant. Asked for lessons learnt, this attendee acknowledged that PoCs are resource and time intensive, and most fail, but investment is worthwhile because those that succeed can have material impact on the business.

Another attendee offered further use case. Again, his organisation was looking to handle systems and software written in the 1980s and 1990s when it was common to use in-house code. Many of the systems are nearing retirement, making them complex and expensive to manage ("We call the old coders mercenaries for the amount they charge."). Working on the basis that if "LLMs can translate languages, they can translate code", his firm has built a series of agents to do just that. He calls the output "building blocks of logic" and these can be applied to modernise old applications, replacing legacy and bespoke code with newer, better known languages.

Risk, regulation and other barriers to adoption

Asked why Agentic AI deployment has not been quicker, the conversation naturally turned to the risk profile of banks, insurers and other finance firms around the table. In the heavily-regulated world of financial services, firms are institutionally risk averse. Some acknowledge that a friction exists between risk and reward, between compliance and a fear of missing out, but most opt for caution.

Other barriers persist, too. These include legacy systems (a feature of most established financial services firms, as demonstrated above), poor quality data, and explainability obligations which rule out “black box” solutions.

Agentic AI and the future of the workforce

Opinions varied when attendees were asked how Agentic AI is likely to impact the workforce. Most agreed that the “slash and burn” approach to efficiency saving – use the technology to aggressively cut jobs – is counterproductive. It risks losing in-house skills and the opportunity to reimagine roles that can augment AI capabilities.

One attendee suggested that those providing business process outsourcing services would be hit hardest. These mundane, repeatable, and functional activities can be delivered more cheaply by the AI. Others weren’t so sure. One asked rhetorically: “Do you know how much you are paying for AI?” The inference: the human option might remain equally cost effective.

Elsewhere the tone bordered on existential. “Any knowledge-based act, the AI will do by 2030,” said one voice. “The social contract for human IT doesn’t survive this.” The only break on progress, he said, will be the lack of energy needed to build the data centres to power further AI resources. Others remained bullish, pointing out how humans have reinvented their roles and purpose in the wake of other technology revolutions. They would do so again. And, said another voice, business will always need IT leaders to enable technology, corral adoption across the business, and manage a new team of agentic colleagues.

Agentic AI in Financial Services: Your next transformation engine or your biggest compliance headache? – a CIONET executive roundtable in association with Thoughtworks – took place on Thursday 27 November 2025 at the South Place Hotel, London.



About CIONET

CIONET is the leading community of more than 10,000 digital leaders in 20+ countries across Europe, Asia, and the Americas. Through this global presence CIONET orchestrates peer-to-peer interactions focused on the most important business and technology issues of the day. CIONET members join over a thousand international and regional live and virtual events annually, ranging from roundtables, programs for peer-to-peer exchange of expertise, community networking events, to large international gatherings. Its members testify that CIONET is an impartial and value adding platform that helps them use the wisdom of the (IT) crowd, to acquire expertise, advance their professional development, analyse and solve IT issues, and accelerate beneficial outcomes within their organisation.

cionet.com