

GAPCO

Agentic AI

The new frontier in financial
services innovation

Introduction

Since the launch of OpenAI's ChatGPT in late 2022, generative AI (GenAI) has quickly established itself as a transformative force and continues to dominate industry conversations. Agentic AI now promises to kickstart a new phase of financial services innovation.

With organisations now looking to move beyond pilot projects, advance to production-level implementations, and unlock measurable, impactful benefits, expectations are rising in terms of how much further this technology can go.

Today, using chatbots and copilots, we can get responses to flesh out ideas, answer customer queries, and execute more creative tasks such as generating images and presentations – but the appetite to go one step further and leverage AI to complete end to end tasks is rapidly growing.

Enter Agentic AI.

The accessibility and power of today's large language models (LLMs) allow for Agentic AI – AI agents that can understand users' requests and act on their own with minimal guidance. These AI agents can

make decisions, take actions, and work with other tools and systems to help users achieve their goals – all without the need for constant human input or oversight.

Agentic AI is already here, with Microsoft incorporating AI agents into their products and OpenAI launching their Operator agent within its Pro subscription. Gartner has predicted that 33% of enterprise software applications will include Agentic AI by 2028, up from less than 1% in 2024, enabling 15% of day-to-day work decisions to be made autonomously.¹

We are already seeing our clients turning their attention to Agentic AI and asking a range of questions. What is it? How does it differ from GenAI? What are the benefits? How can they use it in financial services? How can they manage any risks? In this article, we will address these points.

¹'Intelligent Agents in AI Really Can Work Alone. Here's How.'

What is Agentic AI?

Agentic AI – or AI agents – are systems that act autonomously to achieve human-defined goals. Current AI systems act as assistants or copilots to help humans in their work, whereas an AI agent can remove the need for humans for given tasks.

Agents can interact with multiple downstream agents and execute multiple tasks via a single interface and prompt. An investment advisor agent could be asked by the user to figure out investment risk and call market analysis agents to analyse and synthesise market data to then provide a response, while automated trading agents could employ analyst bots to provide trade recommendations to execute. Always-on support agents personalised to the customer can consider the customer's profile and previous interactions to give hyper-personalised advice.

The AI agent ecosystem has developed significantly in the last year, with advancements in memory, tool usage, secure execution, and deployment.

With greater compute power and accessible storage, AI agents can retain and make better use of more information, allowing them to handle more complex, long-running, and context-dependent tasks and provide more personalised assistance, where previously user interactions with models were limited to smaller, narrower tasks.

AI agents can also be configured to seek human input if they encounter novel problems to mitigate hallucinations and learn from these experiences.

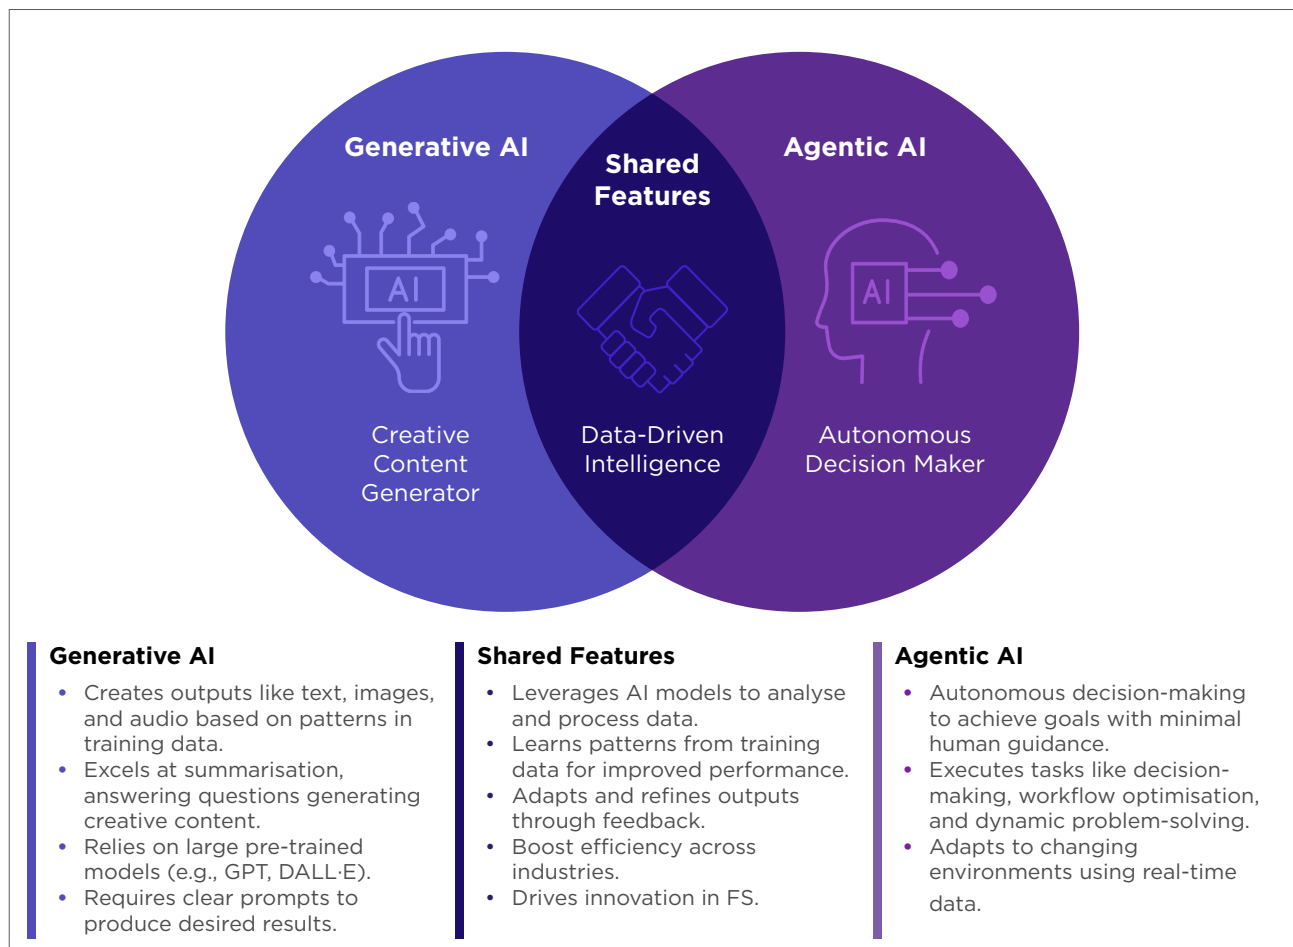


Figure 1: Comparison of GenAI and Agentic AI's capabilities

As we outline below, there are a range of different agents that can be deployed, depending on the task at hand.






Type	Description	Illustrative example: fraud detection system
 Simple reflex agents	<p>These agents act purely based on the current situation or inputs. They do not consider past experiences or learn over time.</p>	<p>Flags a transaction as fraudulent if it exceeds a predefined amount (e.g. \$10,000), without considering the customer's past transaction behaviour.</p>
 Model-based reflex agents	<p>These agents take current inputs and combine them with a simple 'internal model' of how the system works (e.g. basic memory of past interactions).</p>	<p>Flags a transaction as fraudulent not only based on its value but also by checking whether similar transactions have occurred in the customer's account history.</p>
 Goal-based agents	<p>These agents aim to achieve a specific goal by planning actions and evaluating future outcomes.</p>	<p>The goal is to reduce fraudulent transactions to below 1% of all transactions. The agent identifies patterns (e.g. geographical fraud hotspots) and implements targeted rules to achieve the goal.</p>
 Utility-based agents	<p>These agents choose actions based on a utility function to maximise overall satisfaction or achieve the highest net benefit.</p>	<p>The agent evaluates the trade-off between blocking potentially fraudulent transactions and minimising inconvenience to legitimate customers. It assigns a utility score to different actions (e.g. blocking, flagging, or approving) and chooses the one with the highest utility.</p>
 Learning agents	<p>These agents improve over time by learning from past experiences and adapting to new patterns.</p>	<p>A machine learning model learns from flagged fraudulent transactions and adapts its detection algorithm to identify new types of fraud, such as evolving phishing schemes.</p>

Figure 2: Overview of types of Agentic AI

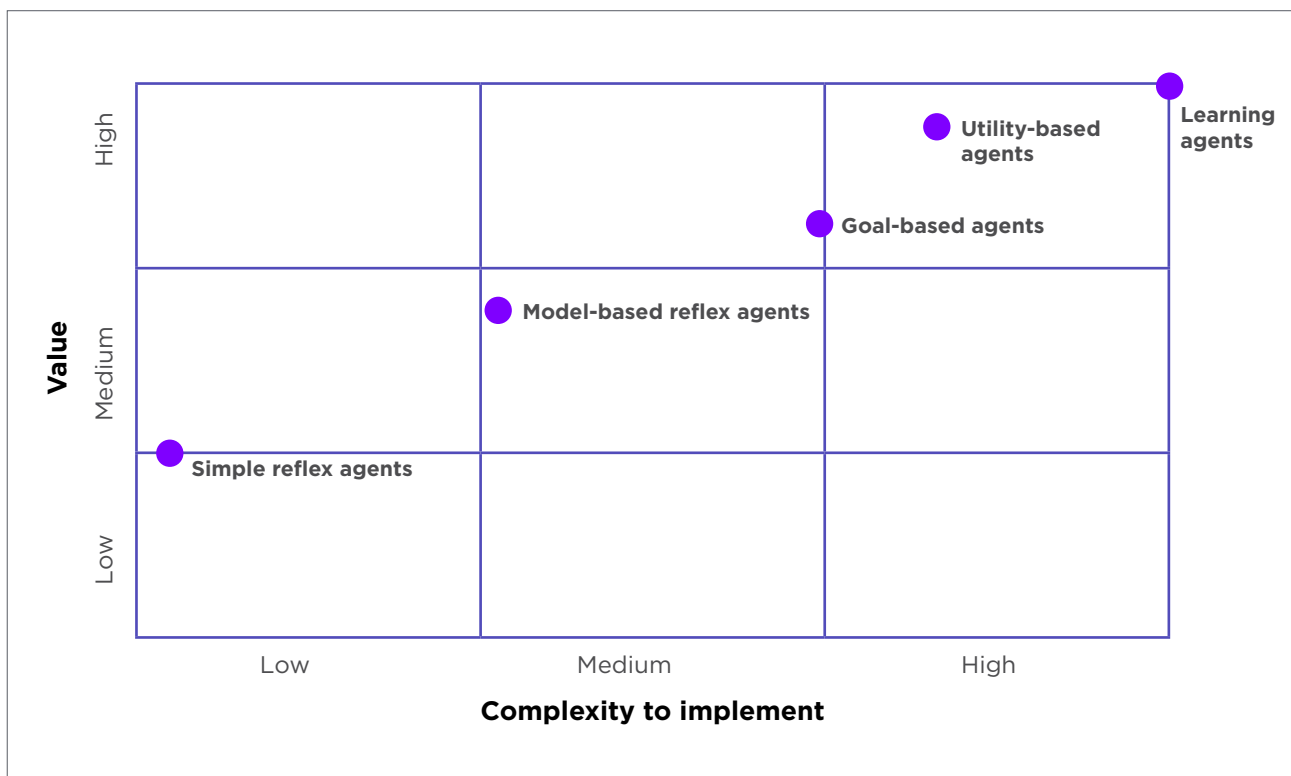


Figure 3: Comparison of Value vs Complexity of types of Agentic AI

Looking at some of the tasks defined as agentic, one might question whether this is not just traditional automation rebranded. However, there is an intuitive progression from traditional automated processes to systems with agentic capabilities.

- **Automated systems** – these rely on deterministic algorithms to execute tasks within a defined process and follow well-defined inputs and produce predictable outputs based on explicit logic.
- **AI-augmented automated systems** – while inputs and final outputs remain well-defined, the outputs of these AI components are generated by models trained on data, rather than solely relying on predefined rules allowing for more dynamic behaviour.

- **Agentic systems** – powered with memory and human feedback, agents can personalise and change their behaviour and apply judgement to decide and take action.

Investment in GenAI serves as a natural precursor to Agentic AI systems, as the two technologies are closely interlinked, with the latter building upon the former's foundational capabilities.

This distinction is crucial in financial services, where the ability to act independently within defined parameters can create significant competitive advantages but carries significant risks that must be addressed upfront.

Agents in action within financial services

The implementation of Agentic AI in financial services is already delivering value by unlocking operational efficiencies and entirely new opportunities. These include:

Autonomous trading system

Modern trading platforms can employ agentic AI to analyse market conditions, identify opportunities, and execute trades autonomously.

These systems can adapt to changing market conditions in real-time, adjusting strategies based on multiple variables while maintaining compliance with regulatory requirements.

Personalised financial advisory

Agentic AI can enhance wealth management through systems that not only provide recommendations but actively manage portfolios based on individual client goals, risk tolerance, and market conditions.

These platforms can proactively rebalance portfolios, identify tax-loss harvesting opportunities, and adjust strategies based on changing client circumstances.

Enhanced risk management

In the realm of risk management, Agentic AI systems continuously monitor transactions, market conditions, and customer behaviour to identify potential risks before they materialise. These systems can autonomously adjust risk parameters and take preventive actions when necessary.

While this creates several exciting new opportunities agents come with increased risk and their use and outputs need rigorous testing as part of any organisations AI governance process to ensure alignment with regulation.



Benefits of Agentic AI in Financial Services

Task automation for operational efficiency

Complex tasks currently requiring human resource can be executed by agents in a scalable way, reaching goals at lower cost and latency. From back-office processes to customer service, these systems can reduce human error and free up human resource for more impactful work. Due care should be exercised to ensure the agents act reliably within an explicit scope of operation to avoid unexpected impacts in the system.

Hyper-personalisation at the customer level

Agents can be integrated into support and sales channels to act as personal advisors or assistants at the moment of asking. Using user data, activity and preferences, coupled with product and policy information in their knowledge base, they can provide tailored recommendations for financial products, give advice to resolve issues, and route to human support agents in complex cases,

resulting in improved customer satisfaction and servicing rates. As with all personalisation use cases, responsible use of customer data (privacy, fairness, inclusiveness, retention) should be followed.

Enhanced decision-making

By allowing an AI agent access to data and sub-agents, they can break down business questions into sub-tasks, leveraging the tools and delegating to sub-agents to conduct the appropriate analysis, and respond with results and decision recommendations. Similar leveraging of knowledge capture of documents and policies can help process cases involving natural language. Human approval of the recommendations is non-negotiable to ensure subsequent actions align with business objectives and has the intended impact.

Challenges and Risks

With automation comes increased risk. While the benefits are clear, the stakes are higher with AI agents, and we anticipate these risks initially will slow down adoption as organisations get comfortable with wider use of AI and GenAI systems maintain a human in the loop. Proper planning and oversight during development, integration, and governance is needed to use them while being compliant with regulations and making sure they stay on task.

Ethics and regulatory compliance

The deployment of Agentic AI must navigate complex regulatory requirements, particularly considering the EU AI Act and similar emerging regulations. As with other AI solutions, financial institutions must triage use cases appropriately and ensure their AI systems maintain Responsible

AI principles (fairness, reliability and safety, privacy and security, inclusiveness, transparency, and accountability) and employ effective privacy and security measures. Any AI agent system inherits the same risks of its components (e.g. transparency of inference in LLMs, hallucinations) with accelerated accumulation of errors if not handled properly.

Integration to business processes

Additional infrastructure, such as the framework to build agents on, and integration with existing tools and applications in the business, require meticulous planning, execution and QA. Deploying agent systems must be approached with care, ensuring testing and rollout is robust, and subsequent monitoring is clear.

Conclusion

As organisations in financial services explore the transformative potential of AI, the concept of Agentic AI raises critical questions: is leveraging AI for process efficiency truly Agentic AI, or does it represent a different paradigm altogether?

From our perspective, the deployment of fully autonomous, decision-making AI remains high-risk, which financial services are not yet ready to adopt at scale. Instead, the current focus lies on automation with human-in-the-loop governance frameworks, aligning with regulatory preparedness and operational trust.

The journey towards Agentic AI needs deliberate steps, clear guardrails, and a maturity in organisational readiness and regulatory clarity. For now, the balance between innovation and control is where progress will be most effectively realised.

Agentic AI represents the next frontier in financial services innovation. By embracing this technology while carefully addressing associated challenges, financial institutions can unlock new levels of efficiency, personalisation, and service quality.

The key to success lies in viewing Agentic AI not as a replacement for human expertise, but as a powerful tool that augments human capabilities in the financial sector.

Authors

Charlotte Byrne

Leo Ferrari

Jessica Forbes

Contact

Charlotte Byrne, Capco UK Gen AI Lead, charlotte.byrne@capco.com

About Capco

Capco, a Wipro company, is a global management and technology consultancy specializing in driving transformation in the energy and financial services industries. Capco operates at the intersection of business and technology by combining innovative thinking with unrivalled industry knowledge to fast-track digital initiatives for banking and payments, capital markets, wealth and asset management, insurance, and the energy sector. Capco's cutting-edge ingenuity is brought to life through its award-winning Be Yourself At Work culture and diverse talent.

To learn more, visit www.capco.com or follow us on LinkedIn, Instagram, Facebook, and YouTube.

Worldwide Offices

APAC

Bengaluru – Electronic City
Bengaluru – Sarjapur Road
Bangkok
Chennai
Gurgaon
Hong Kong
Hyderabad
Kuala Lumpur
Mumbai
Pune
Singapore

MIDDLE EAST

Dubai

EUROPE

Berlin
Bratislava
Brussels
Dusseldorf
Edinburgh
Frankfurt
Geneva
Glasgow
London
Milan
Paris
Vienna
Warsaw
Zurich

NORTH AMERICA

Charlotte
Chicago
Dallas
Houston
New York
Orlando
Toronto

SOUTH AMERICA

São Paulo

