

Trend Topic: The next frontier of global finance

By



Colin Payne

Head of Innovation at the Financial Conduct Authority (FCA) and Chair of the Global Financial Innovation Network (GFIN)

Colin Payne is a distinguished leader in financial innovation, currently serving as the Head of Innovation at the Financial Conduct Authority (FCA) and Chair of the Global Financial Innovation Network (GFIN) convening more than 90 national regulators to build a powerful, collaborative network. In these roles, he spearheads strategic initiatives to foster growth through responsible innovation, such as the newly launched AI Lab, Sandboxes and TechSprint Programme.

Before joining the FCA, Colin was the Senior Director at PwC, where he delivered programs across the Gulf Cooperation Council, including national payment strategies and digital asset initiatives with Central Banks in the region. He has also held Partner and Executive Vice President positions at EY and Cargemini with an early tenure at Apple working with Steve Jobs.

Explore more from
The Kyndryl Institute
kyndryl.com/institute

Agents of change: Agentic AI and the future of financial services

From automation to autonomy

Financial services globally are on the cusp of a profound transformation.

Until now, artificial intelligence in finance has largely meant powerful algorithms that assist human decision-makers – refining credit scores, detecting fraud, or generating reports on command. But a new breed of AI is emerging that does more than assist; it can act.

This is agentic AI.

This brand of AI is capable of independent decision-making, collaboration, and continuous learning without constant human prompts. In practical terms, agentic AI systems can perceive, reason, and act autonomously, executing tasks and adapting strategies in real time. Unlike today's generative AI tools (which wait for our instructions), these AI “agents” operate with a degree of agency that promises to revolutionize financial services.

The implications are enormous. Imagine AI-driven portfolio managers that adjust investments 24/7 to market shifts, or personal financial assistants that automatically optimize a user's finances across banks and services. We are talking about an era of autonomy in finance, moving beyond

automation. This shift could unlock unprecedented efficiency and personalization — bringing finance closer to an autonomous, self-driving paradigm.

But it also raises fundamental questions about trust, control, and risk.

As we stand at the threshold of this agentic AI era, financial leaders and regulators face a high-stakes balancing act: how to harness AI's transformative potential while safeguarding markets and consumers.

Defining agentic AI: A new autonomous frontier

Agentic systems include autonomous agents in the true sense, ones that are able to set goals, interact with their environment, learn from experience, and collaborate with humans or other AI agents. It is this quality of agency (the doing, not just calculating or predicting) that marks a new frontier in financial innovation.

This concept is not science fiction; it builds on real technological breakthroughs in contextual understanding, memory and multi-tasking capabilities. In the financial world, agentic AI is already moving from pilot projects toward deployment. Potential future use cases range from autonomous trading algorithms that self-adjust strategies, to AI-driven compliance bots that monitor transactions and flag anomalies without being asked, to personalized robo-advisors that proactively manage a customer's day-to-day finances within preset guardrails. These agents would operate with minimal human intervention, aiming to augment or even replace certain routine decision processes. Crucially, they are adaptive, learning from each interaction to improve over time.

Before diving into the challenges, let's consider how these autonomous AI agents could disrupt business models and the very structure of financial markets.

Disrupting business models and market structure

Agentic AI has the potential to reshape financial business models in ways not seen since the advent of the internet.

One immediate impact is on the “Do It For Me” economy in finance — where customers delegate tasks to automated agents. With agentic AI, consumers might each employ their own AI financial proxy: a personal bot that shops for the best insurance, manages bill payments, optimizes savings across accounts, or even negotiates mortgage rates. In such a world, the competitive landscape could shift dramatically. Banks and insurers may find that they are no longer marketing directly to humans, but to AI agents acting on behalf of humans. Competition could intensify, as switching providers becomes frictionless when an AI agent can instantaneously scout the market for better deals. Indeed, industry analysts suggest competition will “tick up” as startups and tech-savvy new entrants deploy agentic AI to challenge incumbent banks. Established firms might be forced to adapt their offerings to attract algorithms as much as human customers.

Market structure could also be upended.

If autonomous trading agents proliferate, markets might become more efficient in processing information — or conversely more volatile, if many agents

This concept is not science fiction; it builds on real technological breakthroughs in contextual understanding, memory and multi-tasking capabilities.

respond to the same signals in tandem. Lower barriers to automated market interaction mean even small firms (or individuals with AI advisors) can execute complex strategies that once required teams of traders. This democratization of sophisticated finance could blur the lines between professional and retail market participants. But it may also introduce new forms of systemic risk. For example, if many AI agents are tuned to similar data or strategies, their synchronized actions could create herding effects and sudden swings in markets. Market stability could be tested by fast-moving “flash” events driven by algorithms trading at machine speed, demanding new kinds of guardrails in market infrastructure.

Business models within firms will transform as well.

Many roles and processes currently handled by outsourced service providers or junior staff might be taken over by AI agents. Consulting, accounting and auditing functions, heavy on data processing and rule-based analysis, are likely to be among the first domains disrupted by autonomous AI, according to industry observers. In banking, everything from customer onboarding to risk management could be streamlined by AI that learns and improves over

time. This promises leaner operations and reduced inefficiencies. But it also means firms must rethink their workforce and skill needs. If routine tasks are handled by AI, human roles will shift to higher-level oversight, strategy, and exception handling. The nature of work in financial services could change profoundly, and institutions will need strategies for reskilling employees and integrating human and AI workflows seamlessly.

From a competitive standpoint, agentic AI could lower entry barriers in some areas of finance.

FinTech innovators can leverage AI agents to provide services at scale with relatively modest human headcount. On the other hand, we might see consolidation in certain functions if AI at scale favors those with access to the best data and computing power. The market structure may thus bifurcate: nimble AI-enabled startups on one end, and big incumbents or tech firms (with vast data resources) on the other, potentially squeezing mid-sized players. Regulators will need to watch how these dynamics play out to ensure healthy competition and prevent digital monopolies or undue concentration of AI capability in a few hands.



Rethinking risk paradigms: Trust, transparency and new threats

Financial systems run on trust and confidence, and that extends to trusting the tools and models firms use.

As the UK's Financial Conduct Authority's (FCA) own experience convening industry experts has shown, "trust isn't just a buzzword — it's the whole deal when it comes to embracing AI", and it must be earned through safe and responsible use. If clients and markets don't trust AI-driven processes, adoption will stall, and the benefits of agentic AI will never fully materialize.

Building trust starts with transparency and explainability.

Yet agentic AI systems, often based on complex machine learning, can be "black boxes", making choices that even their creators struggle to fully explain. This opacity is more than a technical annoyance; it strikes at accountability and fairness. Stakeholders need clear insight into how an AI agent is making high-stakes decisions (granting a loan, flagging a fraud, executing a trade).

Without explainability, when something goes wrong, who or what is accountable?

Unexplained algorithmic decisions could also hide or even amplify biases.

For instance, an AI agent learning from biased historical data might systematically disadvantage certain groups of customers in credit or insurance decisions. In a world of agentic AI, firms must update their risk frameworks to include rigorous AI model validation, bias testing, and ongoing monitoring of AI behavior. Concepts like "model risk management" take on heightened importance as models are not static tools but self-updating agents, so the risk is a moving target.

New operational and security risks also loom.

An autonomous AI that can initiate actions could potentially go awry in unpredictable ways. For example, a trading algorithm that "learns" a harmful strategy, or a compliance bot that misflags and halts legitimate customer transactions, at scale. The cybersecurity dimension is critical. These AI systems often require vast data (including personal data). That raises privacy issues as their autonomy could be exploited by malicious actors. Regulators in global financial hubs have warned that the explosive growth of AI agents amplifies and rapidly distributes security risks, bringing novel threats to every organization's doorstep. Firms will need to bolster their defenses and monitoring to prevent AI systems from becoming conduits for fraud or cyber-attacks. For example, imagine an AI agent that's manipulated by a sophisticated deepfake or data poisoning attack. The risk paradigm in finance must expand to cover these AI-specific vulnerabilities. Crucially, there is a growing recognition that many of these risks have a systemic dimension.²

If numerous institutions rely on similar AI models (perhaps from a handful of big tech providers), a single flaw or bad decision logic could propagate widely, creating systemic shocks. Market volatility could be exacerbated by herding behavior among AI agents. Liquidity crunches or flash crashes might be more frequent and harder to unravel when algorithms are interacting with each other at high speed. All this suggests that both firms and regulators will need to develop new stress tests and safeguards for an AI-driven market — for instance, scenario analyses of AI-agent behavior under extreme conditions,

Firms will need to bolster their defenses and monitoring to prevent AI systems from becoming conduits for fraud or cyber-attacks.



or circuit-breaker mechanisms triggered by detected AI malfunctions.

Finally, ethical and consumer protection risks cannot be overlooked.

Autonomous AI engaging directly with customers (say, an AI investment advisor managing a retirement portfolio) must act in the customer's best interest. Who ensures the AI doesn't take excessive risks or mis-sell products? How do we prevent unintended discrimination by AI agents in lending or underwriting? These questions highlight why human oversight "above the loop" remains essential, as one international expert put it. Agentic AI should complement human judgement, not replace it entirely, especially in matters of ethics and values. The risk paradigm for financial institutions, therefore, must evolve into a socio-technical one: not only addressing technical reliability and security, but also embedding ethical guardrails and human accountability into AI deployments.

Regulatory frameworks under pressure: A new frontier for regulators

The advent of agentic AI doesn't just disrupt industry players, it also tests the limits of current regulatory frameworks across the world.³

In the European Union, policymakers are opting for new rules specific to AI. The upcoming EU AI Act will impose requirements on high-risk AI systems, including transparency, risk assessment, and human oversight, with provisions to ensure clear responsibility and liability for AI decisions. This reflects a more prescriptive stance — setting baseline standards for AI systems before they can be deployed at scale.

Across the Atlantic, U.S. regulators have so far taken a sectoral and principles-based approach, leaning on existing laws (like anti-discrimination statutes or securities laws) to address AI outcomes, coupled with frameworks like the National Institute of Standards and Technology's AI risk management guidelines.



Meanwhile, international bodies such as the Financial Stability Board and the International Organization of Securities Commissions are studying AI's systemic implications and could develop coordinated guidance. There is also discussion of novel ideas — for instance, the International Monetary Fund has floated the concept of “automation taxes” to help society adapt to AI-driven disruption, underscoring how even fiscal policy might play a role in managing this transformation. What's clear is that global coordination and knowledge-sharing will be vital. No single regulator has all the answers, and AI agents will operate across borders and sectors, so a fragmented regulatory approach could leave dangerous gaps.

The UK's philosophy to date has been distinctive: no rushed new law for AI, but an emphasis on clarifying how existing rules apply. As I noted at a global fintech summit recently, the FCA does not see a need for a standalone AI-specific regulation at this time. Instead, the focus is on providing “clarity and regulatory confidence” within the robust principles and outcomes-based regime already in place. The FCA believes that its existing frameworks — such as the

Senior Managers and Certification Regime (SMCR) and the Consumer Duty — already cover AI innovations. In other words, accountability and consumer protection obligations apply equally to AI-driven activities: if a robo-advisor makes poor decisions, senior managers can and will be held accountable under SMCR, just as they would for a human advisor's failings. Rather than writing a raft of new rules that might quickly become outdated, the FCA is focusing on clarifying expectations (for example, how firms should govern AI models, validate outcomes, and protect customer data).

At the same time, the regulatory toolkit is expanding in innovative ways.

Regulators are recognizing that understanding fast-evolving AI technology requires closer collaboration with industry and academia.⁴

Sandboxes, labs, and experimental forums are becoming as important as rulebooks.

This is where the FCA has been taking a pioneering stance.

The FCA's strategic approach: AI Lab and "Supercharged" Sandbox

At the Financial Conduct Authority, we have adopted a strategy of "open innovation" in regulating AI, engaging industry experts, testing new ideas in controlled settings, and learning by doing.

A cornerstone of this strategy is the FCA's newly launched AI Lab, unveiled in late 2024 as part of the Innovation Hub. The AI Lab is not a physical laboratory but a program of initiatives that adds an AI-specific focus to the FCA's innovation services. Its mission is twofold: enable the safe and responsible use of AI in UK financial markets while driving growth and competitiveness, and simultaneously help the regulator itself deepen its understanding of AI's risks and opportunities in a practical, collaborative way. In other words, the Lab acts as a bridge between the regulator and the regulated: convening data scientists, fintech firms, banks, and academics to experiment with AI solutions under the FCA's gaze, and to share insights that can inform future regulatory policy.

Under the AI Lab umbrella, the FCA has rolled out several initiatives.

One is the AI Spotlight program, which invites firms to showcase new AI use cases in a controlled environment, giving the FCA and other stakeholders a "peek under the hood" of cutting-edge AI applications.

Another was a high-profile AI Sprint (essentially a multi-day workshop/hackathon) held in early 2025 that gathered more than 300 experts to consider AI's impact on financial services. The outcome of that Sprint reinforced the importance of focusing on trust and transparency for AI. The FCA has published a summary of these discussions and identified follow-up actions, which notably included expanding the AI Lab and launching a "Supercharged Sandbox" and live testing consultation paper⁵ – both of which have been delivered within H1 2025.

The Supercharged Sandbox is an evolution of the regulator's digital sandbox concept, tailored for AI era needs and was launched in June this year at London Tech Week.

It provides a safe testing environment where firms can work alongside regulators to trial AI models with greater computing power, access to richer datasets, and AI-specific evaluation tools. By enhancing the Sandbox's infrastructure, the FCA aims to let AI developers experiment at scale (for example, running an AI trading algorithm on simulated market data, or testing a machine learning credit scoring tool on synthetic loan portfolios) without risking real-world consequences. This initiative is essentially an "AI sandbox" that will run dedicated tech sprints and pilots focused on AI innovation. The FCA has expressed that it "looks forward to inviting firms to collaborate and experiment in new ways" in this environment⁶. Importantly, the Sandbox isn't just for fintech startups – it's open to incumbents and even

Sandboxes, labs, and experimental forums are becoming as important as rulebooks.



Big Tech entrants, provided they meet criteria like having robust post-deployment monitoring plans. The goal is twofold: accelerate beneficial innovation (by helping firms get AI products to market faster, with confidence they've been vetted), and develop shared learning between firms and supervisors on what responsible AI deployment looks like in practice.

Through efforts like the AI Lab and Sandbox, the FCA is effectively “learning by doing” — supervising in real time and adjusting its approach based on evidence. This agile approach is critical, given the pace of AI advancements. It also exemplifies how regulation can become more dynamic: not just writing rules, but also facilitating experimentation under watchful eyes.

The FCA is not alone in this; globally, regulators from Singapore to Canada are launching innovation hubs and sandboxes with similar intent.

But the UK's approach stands out for trying to bake AI considerations into the existing regulatory ethos of outcomes-based regulation.

If successful, it could serve as a model for balancing innovation and risk in the age of AI.


High stakes: The promise and peril of getting it right (or wrong)

The emergence of agentic AI in finance brings extremely high stakes.

The opportunities, as described, are transformative.

Used wisely, agentic AI could widen access to financial services, bringing more people into the financial system through personalized, low-cost advice and automated solutions. It could reduce inefficiencies, saving billions in operational costs and eliminating friction in everything from payments to compliance. Customers could enjoy hyper-personalized products (think tailored investment strategies or credit offers uniquely optimized for an individual's situation in real time).

If we steer this transformation well, “success” would mean a financial sector that is more inclusive, efficient, and dynamic than ever. Imagine a world where small businesses get instant, algorithmically optimized loans because an AI agent can reliably assess their real-time cashflows; or where fraud is detected (and stopped) the moment it begins by AI systems monitoring transactions across the network. In a successful scenario, trust in AI-driven finance would be high because firms and regulators have implemented strong governance, clients know that AI advisors are acting in their best interests, and regulators have the tools to intervene when needed. Market participants would have confidence that AI is a positive sum game: improving outcomes for customers and stability for the system. In short, if done right, agentic AI could usher in a new era of financial prosperity and innovation, akin to a “self-driving finance” revolution where routine financial management is largely automated, safe, and accessible.



This agile approach is critical, given the pace of AI advancements.

However, the perils of getting it wrong are equally profound.

Overreliance on AI without proper oversight could undermine trust and destabilize the system.

In a worst-case scenario, unrestrained AI agents might make decisions that lead to discriminatory outcomes (e.g. denying credit to protected groups due to biased algorithms), massively eroding public trust in financial fairness. Lack of control over autonomous trading bots could contribute to a major market incident, for example, cascading AI-driven selloffs causing a flash crash that ripples through



the global financial system. If firms deploy AI irresponsibly and customers are harmed, we could see scandals that set back the industry's reputation by decades. Systemic risks could materialize if many institutions are blindly relying on the same AI technology. Imagine a software glitch or malicious attack impacting a widely-used AI platform, suddenly knocking out critical services across multiple banks. The consequences of mismanaging this transition include not only financial losses but a chilling effect on innovation: a few high-profile failures could lead to a loss of confidence that derails the positive potential of agentic AI for years. In essence, the industry would face a crisis of legitimacy ("Can we trust AI with our money?") and regulators might resort to punitive measures rather than the proactive, enabling approach we have the opportunity to pursue now.

The difference between these futures comes down to actions taken today.

The high stakes sharpen the imperative for both industry and regulators to actively shape the outcome. It is not an exaggeration to say that the future stability and integrity of financial markets will hinge on how we manage the rise of AI agents in the present.

We must anticipate the pitfalls and build the guardrails before a disaster forces our hand.

As one expert succinctly put it, agentic AI's promise is enormous, but so are its challenges and indeed, "one thing is certain: the agentic AI era of financial services is here, and the time to act is now."

Navigating the agentic AI era – toward success or crisis

Finance stands at a crossroads in the age of agentic AI.

Down one path lies a future where AI-powered agents drive a golden age of efficiency, inclusion, and innovation in financial services. Down the other path lies a landscape of missteps and crises, where unchecked AI undermines the very trust on which finance is built.

The role of leadership, both in industry C-suites and regulatory bodies, is to steer us toward the former and away from the latter. This will require vision, rigor, and humility: vision to embrace the game-changing potential of autonomy; rigor to put in place robust safeguards and governance; and humility to continually learn and adjust course as we uncover what we don't yet know about AI's impact.

The financial world has navigated big technology shifts before — agentic AI is no different in its

disruption, but it is unique in its capacity to act for us and with us in real time. That dual potential, as collaborator and as risk, makes it one of the most consequential innovations of our age for financial services. By adopting a globally informed, yet locally tested approach, drawing on the UK's experience and international expertise, we can chart a course where autonomous AI systems become a force for good within finance.

The stakes could not be higher, and the responsibility could not be clearer: agentic AI is here — now we must make it work for everyone. —

References

- 1 Colin Payne, "AI through a different lens: what 115 experts taught us about AI innovation," FCA Blog, Apr. 23, 2025.
- 2 Bryan Zhang & Kieran Garvey, "Agentic AI will be the real banking disruptor," The Banker, Feb. 25, 2025.
- 3 Kay Firth-Butterfield et al., "How Agentic AI will transform financial services," World Economic Forum, Dec. 2024.
- 4 Citi Global Perspectives & Solutions, "Agentic AI" (Industry Report), Oct. 2024.
- 5 Compliance Corylated, "FCA announces live AI testing service: 'It's not about new regulation,'" Apr. 2025.
- 6 FCA, "AI Lab - Innovation Pathways," Oct. 2024 (updated Apr. 2025).

