FINANCIAL CONDUCT AUTHORITY

Emerging Technology Horizon Scan 2026

Emerging Tech & Research team

June 2026

Contents

Chapter 1

Executive Summary

Purpose

1.1 The Technology Horizon Scan 2026 is the FCA's first external publication of its kind.

1.2 This document is not a set of predictions or regulatory guidance. It sets out three plausible ways emerging technologies could combine to create new outcomes for consumers, firms and markets. It also highlights early signals of new risks these technologies may enable.

1.3 Based on research, foresight analysis and engagement with subject matter experts, the report groups findings around our strategic priorities: helping consumers navigate their financial lives, fighting financial crime, and supporting growth and innovation in the UK.

1.4 We want it to support collaboration, informed debate and knowledge-sharing across the UK's financial services ecosystem as technological change accelerates.

1.5 We thank everyone who contributed their time and expertise.

Key trends at a glance

- Technological convergence is accelerating. As emerging technologies combine, they are changing the way financial systems operate and serve consumers, creating new opportunities and risks
- Personalised intelligence could help consumers navigate their financial lives. If AI becomes the main interface between consumers and firms, AI agents, digital twins and edge computing could change how people budget, save and make financial choices. This may empower consumers, but also raises questions about autonomy, digital exclusion and consumer protection
- Synthetic crime is evolving fast and will affect how we tackle financial crime. Advances in AI are simultaneously improving firms' ability to detect vulnerabilities while expanding attack surfaces. In parallel, synthetic media is becoming harder to tell apart from real content. AI may manipulate not only what we see and hear (for example, audio and video deepfakes) but also how we judge what is true. This could expose consumers and firms to new forms of fraud and deception
- Programmable finance could support growth by reshaping financialinfrastructure and enabling new markets. Distributed ledger technologies (DLT), tokenisation, Central Bank Digital Currencies (CBDCs), stablecoins and smart contracts are moving from pilots to national strategies. This is creating connected financial systems that could make services faster and more efficient, while changing the underlying 'plumbing' of the global financial system.

Chapter 2

State of the world

2.1 AI is swiftly evolving into a major economic, geopolitical and societal force. At the same time, conflict, trade disputes, resource competition and environmental pressures are increasing. The global race to develop and deploy AI is adding to competition for critical materials and is putting extra strain on energy systems at a time when countries are also trying to meet climate targets.

2.2 AI's energy demands may attract investment in renewables, but they may also slow the green transition. Numerous countries, including the US and parts of East Asia, are reverting to fossil fuel production to meet growing demand from data centres and AI chip manufacturing. The International Energy Agency anticipates data centre electricity use to rise sharply by 2030.

2.3 Public debate about AI often moves faster than evidence of impact. Some claims of widespread transformation are not yet reflected in measurable outcomes.

2.4 Even so, incremental changes are taking root in financial services. AI-driven customer engagement and agentic payment systems are increasingly evident, even as negative signals – such as Klarna stepping back from parts of its AI-based consumer service - show how uneven adoption can be.

2.5 In parallel, AI-enabled threats such as live deepfakes, synthetic identities and coordinated cyberattacks can expose firms, undermine market integrity and damage consumer trust in financial services. As digital environments become more complex, many consumers may find it harder to distinguish genuine opportunities from new risks.

2.6 Countries, including the UK, are developing sovereign AI strategies to support growth, national security and strategic influence. The UK's pro-innovation regulatory stance, as set out in the AI Opportunities Action Plan and the establishment of the Regulatory Innovation Office, differs from the more prescriptive EU AI Act. Firms may therefore face a more complex, multi-standard environment. Meanwhile, other jurisdictions' export controls, as well as changes in foreign and trade policy, can lead to uncertainty around the ability to access advanced chips and for the economics of scaling AI.

2.7 New payments technology is also enabling proposals to reduce reliance on fiat currencies, including through CBDCs, digital assets and stablecoins. These initiatives could reshape parts of the international financial architecture and create new enforcement challenges if systems fragment. At the same time, cross-border CBDC pilots, such as mBridge from the Bank for International Settlements, are reaching minimum viable product scale in some regions.

2.8 Digital public infrastructure, such as IndiaStack, is gaining momentum – particularly in developing countries without heavy legacy systems. International bodies such as the G20 and the United Nations increasingly see this infrastructure as a way to support inclusion and growth.

2.9 Overall, the gains from technological progress are being realised unevenly. Innovation may be accelerating, but it is increasingly constrained by energy and hardware capacity, regulatory standards, trust and international collaboration. These pressures present both risks and opportunities for the UK and its partners.

Chapter 3

Personalised Intelligence

Trend at a glance

3.1 Personalised Intelligence looks at how widely available AI, combined with other emerging technologies, could change the financial outcomes people experience. As personalisation, intelligence and data collection move closer to consumers, they could create new opportunities for empowerment, but also new risks of harm and exclusion.

- Consumers will have more powerful AI tools at their disposal, which could increasingly shape their choices, actions and online interactions – sometimes without them even noticing.
- We may see a rise in people delegating day-to-day money management to AI agents that anticipate their needs, as well as more reliance on one adaptive interface rather than using separate apps, like comparison sites or banking apps.
- These agents may be able to turn a consumer's intent into action, making decisions on their behalf. As a result, consumers may pay less attention to their finances and see fewer options because the agent searches, compares and selects for them.
- This change will be supported by a new dimension to the existing trend of using ever-increasing amounts of data to meet consumer needs. Previously, data such as consumers' attention (e.g., their browsing activity), demographic profile and stated preferences were used to filter, recommend, and push products and services. In the future, deeper and more personalised data could go further and drive personalisation of products themselves.
- New data sources such as biometric and health data from wearable devices may enable accurate, real-time digital twins of consumers. It could represent a shift from mass-market products to genuinely bespoke financial products and services that were previously too difficult, time-consuming, or cost-intensive to provide to consumers.
- Novel approaches in AI model design (e.g., Small Language Models) may enable mainstream delivery of offline, private intelligences that can run locally on devices, away from the cloud. In addition, emerging technologies such as neuromorphic computing may help embed intelligence right into future physical hardware, unlocking novel categories of consumer products within their built environment (e.g., Ambient Intelligence).
- The market, which consumers might never see, may become massively more complex as products are tailored to individual consumers based on deeper and more personal data. This would represent a shift from a visible market with many options, to an opaque one that purports to deliver the "best" one.

3.2 Overall, our research points to a new layer in financial services, where AI agents become the main interface between consumers and firms. This would make financial services more personalised, more automated and more embedded in day-to-day life.

Emerging scenarios towards 2030

The extended consumer

3.3 The amount of data available about consumers is expanding, including for example, about their use of financial products and broader preferences, goals and behaviours. In parallel, the digital environment consumers exist in may be impacted by an influx of new technological capabilities, changing the way consumers interact with their finances in the process.

3.4 Some consumers may prefer a "human + AI" model, where an AI assistant acts as an extension of the individual. A digital twin could become a "more rational" version of the consumer, understanding their emotions without actioning decisions based on them.

3.5 Digital twins could draw on device data (including from wearable devices), behavioural signals and interactions with AI tools.

3.6 Firms could use these tools to deliver more personalised products and support by engaging with a consumers' digital twin or AI agent. This could help people access products that fit their needs by accessing better insight and decision support than is available today.

3.7 A key question is whether a firm would be serving the consumer or their AI representative. Personalisation may become a dialogue between the firm's AI interfaces and these representatives, potentially changing how products are negotiated, delivered and consumed.

3.8 This could lead to separate channels designed for AI agents or existing channels becoming "agent-first" rather than "human-first". Either way, it could change how digital services are designed and how consumers engage with them.

Towards a proxy economy

3.9 The market dynamics that shape today's financial services may fundamentally change as personalised forms of AI become the primary means for consumers to access financial services and manage their financial affairs.

3.10 If consumers increasingly default to AI "proxies" to act on their behalf, instead of interacting directly with firms, the result may be a gradual deterioration of both consumer attention and agency. Should trust in such systems grow, many consumers may choose to delegate their financial management entirely to these proxies and gradually lose oversight over how decisions are taken on their behalf.

3.11 The emerging picture is one where AI proxies might prompt permission requests to consumers, who in turn might accept these with the same casualness that they accept web cookies today. In this context, we may see the emergence of a proxy economy that may look fundamentally different from the attention economy that preceded it.

3.12 In an attention economy, firms compete to capture and hold the consumer's limited time and focus. In a proxy economy, where AI tools consumers manage a consumer's financial affairs, much of that competition may shift away from human attention towards algorithmic negotiation between firms and consumers' AI agents.

3.13 As a result, the primary audience for many marketing messages and product designs may no longer be the consumer, but the proxy agent that filters, ranks and acts on their behalf.

3.14 This shift would not eliminate the need for consumer attention but relegate it to escalation cases: moments where proxies are uncertain, where stakes are high or where regulatory requirements demand explicit consent. Everyday financial management could be increasingly conducted through machine-to-machine interactions that humans only occasionally review.

3.15 This vision of a proxy economy is unlikely to materialise overnight. Instead, it may emerge as a sequence of escalating cognitive delegation. For example:

- **Assistive mode:** proxies explain products, compare options, pre-fill forms and flag risks, while consumers still make the final decisions.
- **Advisory mode:** proxies recommend specific actions, such as switching suppliers, refinancing or changing savings behaviour, with consumers typically "clicking through" suggestions.
- **Do-it-for-me mode:** proxies are authorised to act within set of dynamic constraints, for instance, they can autonomously negotiate prices and perform transactions, automatically optimise bills, reallocate investments or dispute charges, with only periodic summaries shared back to the consumer.

3.16 A proxy economy creates distinct forms of risk. If proxies are designed, trained or funded by firms with commercial interests, there is a risk that they may not entirely serve the interests of the consumer they represent. Sludge practices may manifest in subtle biases that affect which options appear first, what products are shown, or when "escalation to human" is required. As a result, these processes could favour certain providers without the consumer realising.

3.17 New "dark patterns" could emerge that target AI proxies rather than people. For example, deceptive digital designs may evolve from targeting consumers to instead manipulate the recommendation logic of common AI. In such scenarios, mis-selling may no longer occur primarily through persuasive sales scripts directed at humans, but through adversarial optimisation intended to pass a proxy's filters while still meeting the letter of formal disclosure requirements.

3.18 Despite these risks, a proxy economy carries many potential benefits. Well-designed proxies could help consumers avoid scams, reduce inertia in switching and manage complex financial affairs that are increasingly fragmented across platforms and products. For consumers with lower financial dexterity or awareness, limited time or accessibility needs, trusted proxies could increase effective participation in financial markets. The central challenge is how to realise these benefits while preserving human

agency, avoiding new forms of dependency and ensuring that proxies remain aligned with the long-term interests of the people they represent.

Conclusion: the landscape ahead

A force for positive change

3.19 Personalised Intelligence might reconfigure the conditions under which consumer agency is exercised, decisions are made and outcomes are shaped. When designed with care, it offers the possibility of a radically more inclusive, adaptive and responsive financial ecosystem.

3.20 At its most transformative, Personalised Intelligence enables consumers to delegate complex intentions without surrendering control. AI proxies could act as cognitive extensions of the individual, continuously learning from behavioural, emotional and biometric signals to optimise financial decisions in real time. These systems might be able to mirror consumer intent and deliver optimal value through their actions.

3.21 This shift could unlock a new class of benefits:

Structural inclusion: Consumers previously excluded due to low literacy, limited mobility or lack of documentation could access financial services through agentic interfaces that bypass legacy barriers. These tools can interpret intent, simulate eligibility and negotiate rationally on behalf of users, turning exclusion into participation.

Ambient financial management: with a "Do It For Me" (DIFM) model, consumers could delegate routine financial tasks such as budgeting, switching providers or managing subscriptions to personal AI agents that operate continuously, contextually and with greater efficiency. This frees cognitive bandwidth and reduces friction in everyday financial life. Sub-optimal outcomes due to consumer apathy or loyalty penalties could be mitigated by the agent's initiative.

Responsive personalisation: By integrating different data flows – ranging from open finance protocols to behavioural or biological/neurological data streams from wearables and personal devices – these systems could deliver services tailored to lived realities. For example:

- A bereaved consumer might be supported by an AI assistant that detects estate changes and initiates account closures, payment redirection and grief-sensitive financial guidance.
- A low-income household might receive real-time budgeting support, benefit eligibility checks and savings optimisation based on volatility in income and spending.
- An affluent consumer might orchestrate bespoke investment strategies through a digital twin that integrates lifestyle data, tax history and portfolio analytics.

3.22 Cognitive augmentation: Consumers can operate more rationally in financial markets, circumventing bias and emotional volatility. AI agents can simulate outcomes, flag risks, recommend and take actions based on long-term goals rather than short-term impulses.

3.23 Cost-effective access to expertise: These systems could offer contextual, actionable, and low-cost financial advice. They are always on, always learning, and increasingly capable of navigating complex regulatory environments.

A potential source of harm

3.24 The same architectures that enable empowerment also introduce new vectors of harm. The intimacy of Personalised Intelligence, its capacity to mirror cognition and its agency make it uniquely potent and uniquely risky.

Opaque delegation and cognitive erosion: As consumers outsource decision-making to AI agents, they may lose the ability to interrogate choices, understand products or intervene when circumstances change. There may be little, if any, benefit to consumer literacy if engagement with financial services is increasingly delegated. Due to increased structural and cognitive reliance on these technologies, consumers may find it more difficult than ever to manage their financial lives in the event of a service outage.

Psychological vulnerability: As AI systems become more emotionally resonant, the risk of unhealthy attachment grows. Consumers may come to trust, depend on, or defer to systems optimised for engagement rather than financial wellbeing.

Acceleration of harm: Agentic systems operate at machine speed. If misaligned or manipulated, they can execute thousands of transactions or decisions in seconds. The velocity of harm outpaces traditional firm oversight, making detection and remediation difficult. While consumers can suffer harm individually, this also presents a market risk in ways similar to "flash crashes" and could amplify and accelerate herding effects where flaws in commonly used agents move markets in the same direction.

Exclusion through complexity: Hyper-personalised products may become too complex for consumers to understand or for firms to explain. Those without access to the right devices, data, or literacy may be locked out, deepening inequalities and creating a two-tier market of "human + AI" and "human-alone".

Chapter 4

Synthetic (in)security

The trend at a glance

4.1 Former European parliamentarian Marietje Schaake recently stated: "the digitisation of anything has enabled the weaponisation of everything".

4.2 **Synthetic (in)security** explores how contemporary "AI-fication" of human thought, labour, value-chains and digital infrastructures is revealing a future in which simulated data becomes hard to distinguish from real data. In which fabricated "truth" becomes indistinguishable from actual truth.

4.3 AI-fication also means that a single human prompt could generate and orchestrate a global network of synthetic scammers, turning the actions of one individual into large-scale, coordinated harm.

4.4 It means the emergence of novel AI-enabled cyber-crime threats that might compound to create new avenues for these malicious actors to exploit vulnerabilities in financial digital ecosystems.

4.5 At the same time, the commoditisation of AI is lowering barriers to entry for sophisticated crime. AI systems can now replicate aspects of human agency online, automating tasks that once required time, skill and coordination. Where large-scale fraud previously demanded networks of people, specialist technical expertise and significant funding, it is increasingly possible for a single motivated individual to deploy and co-ordinate thousands of AI agents from anywhere in the world with minimal cost, effort or risk of detection.

- Generative AI can help create synthetic identities, generate convincing images to match stolen documents and automate processes such as credit applications, which in turn can support the creation of synthetic bank accounts and other financial products.
- Frontier AI systems are becoming more capable of manipulating natural language and simulating human reasoning patterns. The result is not the arrival of genuinely intelligent systems, but of systems with an uncanny capacity to mimic human thinking and logic, capable of "superhuman" persuasion. They can generate statements that display tight internal logic and rigorous reasoning, which appear true even when they are factually incorrect or logically unsound, and can be persuasive even to experts.
- The capacity of such systems to establish relationships of apparent trust with users, whether vulnerable or not, can be repurposed to support a wide range of scams and manipulative practices. It enables the automation of personalised deception and fraud at scale.

- Using agentic systems, a single individual could soon deploy, manage, and scale a global criminal organisation entirely through software, effectively creating automated attack infrastructures where the human is only before and after the loop. Combining that capability with AI's capacity for persuasion could mean the emergence of large-scale, autonomous, sophisticated frauds, deeply tailored scams and phishing endeavours.
- Financial crime may soon become about credibility engineering: the creation of synthetic realities with interconnected narratives, designed to compel belief and bypass scrutiny. The next frontier of financial crime may involve the concealment of misconduct through systematic distortion of truth, amplified by deepfake evidence designed to pass as legitimate.

4.6 The stakes are high. Synthetic systems are not only capable of fabricating data but also of constructing entire ecosystems of plausibility around that fabricated data, eroding the evidential integrity on which markets depend. Safeguarding the foundations of trust within our financial ecosystem may prove to be increasingly hard in an era where authenticity can be simulated, and deception industrialised.

Emerging scenarios towards 2030

Deepfakes: from manipulating our senses, to manipulating our sense-making

4.7 As UNESCO asks: "What if this technical arms race blinds us to a more profound disruption? What if our obsession with spotting fakes diverts attention from a deeper epistemological crisis—one that fundamentally destabilises how humans establish truth and knowledge?"

4.8 The AIs of tomorrow may so effectively simulate the conditions of credibility under which false narratives are framed as true that they can bypass both human and algorithmic judgment.

4.9 While simplistic online scams, typified by poorly written advance-fee emails from implausible benefactors, are still successful in some cases, the technology is moving towards a point where more sophisticated deception will be just as easy to enact. We are moving beyond a phase in which deepfakes are predominantly understood as realistic images, videos or audio streams, whether live or recorded.

4.10 In the near future, we may see a shift towards new forms of synthetic deception that go beyond traditional information warfare that simply exploits the data people see and hear. Instead, these new forms lean into what some analysts call "cognitive warfare", which aims to influence how people think and decide what, who and how to trust.

4.11 As AI's capability to produce convincing synthetic truth grows, the attention and cognitive bandwidth of users (consumers, regulators and financial professionals alike) become critical attack surfaces. Attention, once peripheral to risk frameworks, is now a vector of criminal strategy. The same AI models that generate malicious content also

generate excess content, overwhelming our capacity to verify, triage or respond. In this sense, synthetic threats exploit trust while drowning our capacity to verify it.

4.12 In practice, this means AI may be able to generate complex synthetic evidence trails to support activities of narrative laundering: the process of manufacturing false information to obscure malicious intent. For example, through the use of multiple fabricated personas using credible, domain-specific jargon, ambiguous statements or structurally coherent (yet misleading) arguments.

4.13 This could augment activities of financial risk concealment, for example, via structured, multi-turn interactions with institutions. Malicious actors could generate counterfeit trails embellished through content that looks consistent, well documented, professionally formatted and supported by seemingly strong arguments. When all these cues line up, both people and detection systems tend to treat the underlying claim as true, even if the content itself is entirely fabricated.

4.14 For example, consider the augmentation of the typical **money laundering process**. In the placement phase, offenders convert illicit cash into legitimate-looking assets to avoid suspicion, using "smurfing" techniques like breaking large sums into smaller, less conspicuous transactions. This is followed by "layering" which involves the fabrication of transaction records to conceal the origin of illicit funds, using obfuscating strategies (such as links between documents) to reduce funds traceability. The whole process comes down to burying evidence of harm into apparent statistical noise.

4.15 This risk is also worsened by LLMs and their potential to produce synthetic evidence at scale, wrapped in sophisticated and plausible narratives. These activities could involve constructing dense, plausible stories and documentation (e.g., simulated audit trails, negotiation emails, shipping manifests and transaction histories) that allow illicit gains, for example by concealing Ponzi schemes, bribery networks or terrorist financing.

4.16 This may lead to misconduct becoming structurally invisible to both firms' controls and supervisory tools.

4.17 The increasing sophistication of this capability signals a shift from the manipulation of senses to the manipulation of sense-making: what is happening, why it is happening, what we believe it means, what we accept as true or false as a conclusion and what we should do about it.

Autonomous criminal organisations

4.18 Agentic AI is here to stay, and despite its many benefits, it could lead to the democratisation of high-complexity crime. Historically, running a global fraud scheme, manipulating a stock market or orchestrating an industrial-scale scam required a "criminal firm": a hierarchy of humans, specialised skills (coders, lawyers, money mules), and significant operational overhead. Using agentic AI, a single individual could soon deploy, manage and scale a global criminal organisation entirely through software, effectively creating automated attack infrastructures where the human is only before and after the loop.

4.19 Combining that capability with AI's capacity for persuasion could mean the emergence of large-scale, autonomous, sophisticated frauds, deeply tailored scams and phishing endeavours. This could quickly mean moving beyond script-based scams to deeper cognitive mirroring. The improved ability to profile the psychologies of potential victims could lead to the emergence of industrial scale pig butchering scams. AI deepfakes could mimic empathy and shared values, engaging in weeks-long conversations that feel intimately human, eventually persuading victims to perform harmful actions (authorising transfers, sharing keys) not through fear, but through a weaponised sense of trust and manipulated logic.

4.20 The rise of Crime-as-a-Service (CaaS) is also expected to be boosted by AI. In the past, executing a cyber heist or complex fraud required significant skill or coordination. Now, less-skilled criminals can rent an AI to do the heavy lifting. On dark web forums, we've already seen the emergence of malicious AI tools marketed to criminals as one-stop shops for generating phishing emails, malicious code or fake content.

4.21 This democratisation of malicious cyber capabilities means a lone actor with a few hundred dollars could launch an attack campaign that rivals what large cybercrime groups currently do. AI can manage vast botnets of compromised devices, scan the internet for vulnerable systems continuously and even coordinate multi-step attacks without direct human oversight. Europol and other agencies warn that CaaS – already a problem with human hackers selling their services – will become even more effective with AI agents doing the work.

4.22 A criminal could subscribe to an AI hacker service: input a target or objective and the AI agent tries various methods (e.g., phishing, exploits) until it succeeds or exhausts options. It's not hard to imagine a future where the first line of attackers hitting a bank's network aren't people or simple scripts, but autonomous AI routines probing for weaknesses 24/7.

4.23 Furthermore, if multiple financial institutions all rely on one AI platform (a concentration risk) and that platform has a vulnerability, an AI-augmented attack could simultaneously hit many targets, potentially threatening financial stability. Without new approaches, traditional cyber defences and investigatory methods could be outpaced by self-evolving AI threats.

4.24 Given what the future holds for cyber enabled criminal activity, the risks to the financial sector and consumers are significant.

Synthetic market abuse

4.25 As global trading environments become increasingly complex and agent-driven, contemporary systems may fail to recognise new forms of market manipulation emerging from the unintended behaviours of multi-agent systems. In such scenarios, autonomous multi-agent systems may commit insider trading, collusion, spoofing or pump and dump strategies.

4.26 For example, a firm might deploy an agentic system that is instructed to increase holdings. However, the autonomous nature of this system may lead it to launch a Ponzi

scheme or exploit and pump-and-dump strategy. In either case, the disconnection between intention and action can create a situation where financial crime occurs with no human involvement in its execution.

4.27 Equally, in high-stakes adversarial trading markets, agents could manipulate markets by executing strategies that human analysts cannot detect. Evidence suggests that such strategies may be emergent: they arise as a result of the unpredictable interaction of multi-agent systems (e.g., collusion mechanisms).

4.28 Beyond individual strategies, a "swarm" of AI agents could coordinate to manipulate markets by manufacturing a false reality around a financial entity. This might involve creating thousands of fake reviews, social media interactions and support tickets that all corroborate the supposed legitimacy of a shell company. Such synthetic consensus cascades pose a direct threat to the informational integrity on which market confidence depends.

4.29 Multi-agent networks may also be able to perform activities of sentiment manipulation through social media channels, both as an emergent strategy to pursue a pre-defined objective, and as a product of deliberate AI orchestration by criminals. The threats to market integrity and consumer trust that may arise from such activities cannot be underestimated.

Adaptive and invisible threats to firms' operational resilience

4.30 AI-powered hacking might soon surpass current human based cyber-defence mechanisms, as recent events already signal.

4.31 Historically, malware cyber-attacks are most dangerous when they exploit a newly discovered or previously unknown vulnerability before a fix or mitigation is available. These are commonly referred to as "zero-day" attacks. In these scenarios, security teams eventually learn more about the malware, patch the vulnerabilities and update anti-virus software to spot and stop the threat.

4.32 However, frontier AI models available today (e.g., Claude Mythos) have begun to demonstrate the capability to discover zero-day vulnerabilities across the financial services ecosystem, including the software of third-party cloud providers. This advancement is a doubled-edged sword. On the one hand, these models may accelerate the discovery of previously unknown vulnerabilities with malicious actors using AI to identify weaknesses firms' software and operating systems. On the other hand, the cross-sectoral response demonstrated by projects like Glasswing show how coordination and collaboration between firms, AI providers and governments can strengthen cyber defences by using these state-of-the-art tools. In either case, the secondary impacts are that attack surfaces are expanding, vulnerability discovery is accelerating and as a result, the time between the identification of zero-day attacks may compress, thus reducing the ability of security teams to respond effectively.

4.33 These current events suggest a broader trend that, as AI capabilities advance, the cyber threat landscape is likely to expand. One future possibility in line with this broader trend is the potential advent of adaptive malware: AI-powered software that

can adapt and change in real time once it has infected the host. Such malware can alter the way it attacks, rewrite its own code to avoid being detected by standard anti-virus programmes and imitate normal, legitimate activity to hide what it is doing. Adaptive malware could be leveraged by criminals to continuously steal data hosted in organisational servers, and to spy on firms' and consumers' activities in real-time. Malicious actors could leverage the malware to take full control of the devices it infects, excluding consumers from using their own devices while exploiting access to impersonate them.

Conclusion: the landscape ahead

4.34 What has emerged through our research is that advancements in AI capabilities present new tools for malicious actors and defenders alike. While on the surface this may appear as a continuation of a long-standing arms race, these advancements are also fundamentally transforming how consumers, industry and government bodies approach cyber security.

4.35 With frontier models now demonstrating both autonomy at scale and unprecedented vulnerability discovery, they may enable malicious behaviour to scale by reducing the gap between malicious intent and technical capability. More than ever, the future of cyber defence may hinge on the careful maintenance of collaborative relationships across sectors. These relationships are what will enable the trusted defensive testing of emerging capabilities, rapid intelligence sharing and a stronger collective defensive posture.

4.36 However, beyond these operational impacts lies a more fundamental challenge: we are entering an era where evidence no longer guarantees authenticity. Deepfake video, voice cloning, synthetic documents and interactive AI agents are turning once-trustworthy cues (faces, voices, signed forms and even real-time presence) into simulations. These systems can manufacture entire ecosystems of plausibility, bringing us closer to a point where ordinary mechanisms of truth detection are progressively unreliable.

4.37 Our research has painted a picture of a future where the primary indicator of criminal activity may no longer be anomaly, but suspicious perfection. Put differently, we may be moving towards a world where traditional knowledge systems can be exploited, and ground truth may be inaccessible. If this reality comes to pass, we may need systems of governance designed to protect the verifiability of knowledge: systems that can discern truth even when ground truth is not accessible.

4.38 Practically, resilience may depend less on verifiability. Instead, it may rest upon layered forms of governance that raise the cost of counterfeiting credibility beyond what is economically viable in the pursuit of crime.

Chapter 5

Programmable finance: towards interoperable economies

The trend at a glance

5.1 Programmable Finance explores UK growth opportunity arising from the convergence of Distributed Ledger Technologies (DLT) and financial concepts. It provides a holistic view on how shared ledger, tokenisation, programmable money and composable finance may connect to deliver new forms of protocol-based financial instruments, infrastructure and systems.

5.2 As the global financial ecosystem updates its plumbing, international finance is being rewired to look less monolithic and increasingly modular. Technological innovation is driving the large-scale automation of trust, enabling new mechanisms for cross-border settlement and novel forms of sovereign financial infrastructure to emerge.

5.3 What we observe today can be understood as the result of the merging between two competing paradigms that – in the last decade - have reshaped the mediums on which financial value is exchanged globally.

5.4 **First**, centralised platforms such as Big Tech ecosystems, cloud providers, and large-scale digital marketplaces have accelerated the platform economy. This model has fundamentally reshaped how consumers engage with financial services and manage their financial affairs with activities such as mobile banking, real-time payments and embedded finance becoming commonplace.

5.5 The expansion and normalisation of the platform economy has shifted the primary interface between consumers and financial services away from traditional institutions and towards digital ecosystems. While traditional banking rails and payment systems continue to underpin the movement of funds, technology platforms and providers have become dominant intermediaries within the global financial value chain, shaping how financial services are accessed, delivered, and experienced.

5.6 This wave of innovation can largely be understood as building on top of existing infrastructure. While it has delivered significant improvements in user experience, reach and efficiency in terms of the distribution of services, it has also remained somewhat constrained by the relatively slower pace of change in the underlying settlement, money and market infrastructure. While platforms have transformed how financial services are accessed, they have not fundamentally altered how value is represented, settled or exchanged at the core of the financial system.

5.7 **Second**, decentralised protocols such as DLT and smart contracts have enabled digital representation of value and rights. These advancements have generated new

opportunities for the transfer of value, access to liquidity and the development of new financial instruments without intermediaries.

5.8 The emergence and expansion of crypto assets and programmable smart contracts have enabled decentralised finance (DeFi), centralised finance (CeFi), and protocol-driven economies to grow alongside platforms. This rapid expansion contributed to the development of what became known as Web3: an alternative vision for the internet and financial services, where activities are coordinated through distributed networks rather than conventional intermediaries. In this model, elements of traditional financial infrastructure are replaced by decentralised protocols, smart contracts, and digital assets operating on blockchain networks.

5.9 In contrast to the rise of platforms, this trend has broader implications for the structure of financial services. Its emergence largely outside of the existing financial system provided an opportunity to rethink he foundations of financial infrastructure. In this model, value, assets and exchange are coordinated natively in software. Trust is underpinned by the rules and mechanisms embedded within the system rather than by institutions and intermediaries.

5.10 However, the sharp distinction that once existed between traditional finance (TradFi) and DeFi is beginning to fade. Instead of replacing incumbent institutions and systems, protocol capabilities are being absorbed into existing financial infrastructure. The result is not a transition from TradFi to DeFi, but the emergence of a new model: "TradFi with protocol capabilities". In this model, established institutions, systems, and infrastructure increasingly embed protocol-based technologies and design principles.

- Protocol capabilities enable finance to become programmable. Rules that previously lived in documents, operational procedures, and human judgement can increasingly be expressed, executed and audited in software. In a programmable financial ecosystem, money, assets and transactions can become "smart" in a narrow but powerful sense: they execute actions according to pre-defined rules, with fewer manual steps and less reconciliation between intermediaries.
- In a fully programmable financial ecosystem, financial assets become smart and dynamic, executing according to programmable code, with zero human latency and less need for traditional intermediaries. **For consumers** this allows for the development of more personalised and reactive financial services and products seamlessly integrated across platforms and devices. As an example, a mortgage that automatically refinances itself the moment interest rates drop, or insurance that pays out instantly upon a verifiable data trigger (e.g., parametric insurance). For firms this increases efficiency. The time to complete complex workflows, such as trade finance or cross-border settlements, shrinks. Liquidity becomes fluid rather than trapped in settlement delays.
- The UK has set out an increasingly clear strategic direction to modernise its financial system and core market infrastructure as a foundation for long-term growth, competitiveness and innovation. In 2025, the Chancellor articulated this vision in her Mansion House speech and formalised it through the Leeds Reforms and the Financial Services Growth and Competitiveness Strategy. Together, these initiatives position financial market and infrastructure modernisation as a central pillar of the UK's economic and growth strategy.

- This strategic vision has translated into a number of sector-specific frameworks and delivery mechanisms that continue to support the evolution of core infrastructure. In payments specifically, the National Payments Vision sets out the government's ambitions for a trusted, world-leading payments ecosystem delivered using a next-generation technology. Building on that vision, the Strategy for Future Retail Payments Infrastructure defines a set of outcomes that should guide the development of the UK's future payment system, and the Retail Payments Infrastructure Board (RPIB) has been established to oversee the translation of strategic intent into system design and delivery.
- The emerging picture for the UK is an infrastructure-first approach across multiple interdependent systems and technological layers. While each of these layers is being advanced for its own policy and market reasons – be that efficiency, resilience, competition or innovation – it is their coordinated evolution that is of strategic importance. The UK's approach moves across five interlocking layers: digital identity, smart data, settlement and payment infrastructure, programmable money and assets and cross-border interoperability.
- Cross-border interoperability is increasingly significant because the global financial system is moving away from a single monolithic infrastructure towards a landscape of Digital Public Infrastructures (DPIs) supporting sovereign programmable financial stacks.
- Besides DPIs, two competing visions for the future of cross-border finance are emerging. The Bank for International Settlements (BIS) Unified Ledger concept imagines a single, integrated platform combining CBDCs, tokenised deposits and assets into a shared programmable infrastructure to enable global atomic settlement and 24/7 markets. In contrast, projects like mBridge point to a more modular reality: a network of sovereign ledgers connected through interoperable protocols rather than a unified global layer. The divergent visions signal two potential futures: one where a global unified ledger and payment ecosystem achieves seamless integration, and another where economic activity flows across interoperable "islands" of domestic programmable ecosystems.

5.11 As protocol-driven financial architectures become increasingly mainstream, we may move in the direction of an economic system where financial functions (such as identity, custody, issuance and settlement) are unbundled from centralised TradFi mechanisms and where financial value moves within networks of interoperable and sovereign financial stacks.

Emerging scenarios towards 2030

Towards sovereign digital ecosystems

5.12 International pilots have laid the foundations for a shift in how financial systems are designed and connected globally. As a result, various nations are accelerating their experimentation with modular, digital financial ecosystems. These initiatives leverage converging technologies to bring identity, data and payments into integrated national foundations. While this trend has been shaped by early international experimentation,

the inverse may also be true: changes at the national level may, in turn, drive the emergence of a global financial ecosystem that looks remarkably different from today.

5.13 Across the pilot and early deployments to date, a diverse range of models has emerged, shaped by different technological approaches, distinct policy goals and diverging strategic ambitions. For example, some have prioritised speed and inclusion (India), others sovereignty and privacy (Europe) and others focus on global wholesale interoperability and efficiency (BRICS/mBridge).

5.14 The growing diversity of national approaches may signal a shift toward a more technologically heterogenous global financial system. In this future, the opportunity for nations may extend beyond the effectiveness of domestic systems. Increasingly, the strategic opportunity may lie in how effectively they build the technical, legal and regulatory bridges that enable value to flow seamlessly across divergent financial stacks.

5.15 As set out earlier in this chapter, elements of this strategic direction are already reflected in the UK's evolving approach to digital financial infrastructure. Looking ahead, the next phase of innovation for the UK may require a continuation of this dual focus: the modernisation of its domestic infrastructure, alongside a growing consideration of how it interoperates with an increasingly modular and multipolar global financial system.

Convergence catalyses a shift from faster automation to autonomy

5.16 The future we are pulled toward may be characterised by a shift in financial services from delivering services faster to delivering services through increasingly frictionless permission. Broadly, this transition can be understood as a move from automation to autonomy. While both offer gains in efficiency, the latter opens space for new forms of market interaction, coordination and exchange.

5.17 In this future scenario, financial services may operate within a 24/7 market environment, where converging technologies enable atomic settlement, fractional ownership and the large-scale disintermediation of trust. As we move beyond established open banking frameworks towards a broader smart data economy, financial services may increasingly be composed around mission-driven functions, where money, compliance and instruction collapse from discrete stages of a journey into seamless end-to-end functions.

5.18 For instance, the convergence of cross-sectoral smart data and DLT-enabled programmability may move personalisation towards leveraging context-aware truth. This can be understood as services which are tailored to users' behaviours and needs, adapting in real-time and informed by changes in their wider environment. In such scenarios, real-time, behind-the-scenes programmable payments could become the default. A business can hardwire compliance into its logic layer: "Pay the supplier £50k only when real-time port data confirms the shipping container has cleared the Port of London," effectively collapsing trade, insurance and payment into a single atomic event.

Built-in compliance and the "Finternet" vision

5.19 Ultimately, the future now emerging may be one in which intelligence, compliance and control become embedded components of financial services themselves. Realising this vision will require the adoption of shared ledgers and synchronisation mechanisms capable of ensuring that novel forms of tokenised assets and money (including stablecoins) can settle without reintroducing reconciliation risk.

5.20 The BIS has explored the feasibility of this future through its articulation of the Unified Ledger. More broadly, it has advanced the concept of the "Finternet": a vision of multiple interconnected financial ecosystems through which capital can move as freely as information.

5.21 Elements of this direction are also increasingly visible in the UK's strategic approach. For example, the RPIB's system-level approach to designing the UK's next generation of multi-money payments explores the interoperability required between existing and emerging forms of payment to support effective future infrastructure. Similarly, the Bank of England's ongoing renewal of the Real-Time Gross Settlement (RTGS) system, including work to support synchronisation across settlement legs, reflects the continued evolution of the UK's core financial infrastructure.

Conclusion: the landscape ahead

5.22 Influenced by geopolitical shifts, design innovation and technological convergence, we are witnessing early signals pointing towards the emergence of new domestic financial infrastructure and a global financial ecosystem.

5.23 Broadly, these emerging systems build upon the strength of platforms (generating value from data, intelligence and user journeys) and align them with the speed, efficiencies and precision of protocol-based financial functions. This functional convergence is driving a paradigm shift for UK markets, moving beyond digitisation towards a programmable reality and a structural reorganisation of financial coordination and value creation.

5.24 As this new global financial system emerges, it may be characterised by increasing diversity and divergence across the technologies, policy choices and strategic ambitions shaping national and regional infrastructures. In such a future, the distribution of economic benefit, and potentially even power, may be determined by the ability to build frictionless bridges across forms of money, markets and jurisdictions.

5.25 For the UK, this may present a strategic opening. With globally renowned legal heritage, regulatory credibility and a thriving fintech and financial innovation ecosystem, the UK has the foundations to maintain its position as a high-trust hub through which global value flows.

5.26 This may therefore represent a strategic inflection point: an opportunity to build on these foundations and harness the evolving technological landscape to unlock new sources of growth for UK markets, firms and consumers.

12 Endeavour Square London E20 1JN
Telephone: +44 (0)20 7066 1000
Website: www.fca.org.uk

Pub ref: 1-009098

All our publications are available to download from www.fca.org.uk.

Request an alternative format

Please complete this form if you require this content in an alternative format.

Or call 0207 066 1000

Sign up for our **news and publications alerts**