UK Finance

Agentic AI and Regulatory Challenges

UK Finance response the Call for Views by the Digital Regulation Cooperation Forum

November | 2025

Introduction

UK Finance is the collective voice for the banking and finance industry. Representing 300 firms, we're a centre of trust, expertise and collaboration at the heart of financial services, championing a thriving sector and building a better society.

We welcome the opportunity to respond to the DRCF's call for views on agentic AI. This response synthesises feedback from our member organisations across the financial services sector, reflecting diverse perspectives from retail banking, payments, and financial technology providers. We support the DRCF's efforts to foster dialogue and understand emerging risks and opportunities in relation to this technology.

Agentic AI presents significant opportunities to enhance productivity, automate complex workflows, and unlock new forms of digital commerce, with the potential to make financial services safer, smarter, and more personal. This supports the UK's wider growth agenda. However, the autonomous characteristics of agentic AI require appropriate monitoring, robust controls, adequate human oversight, and dynamic goal setting.

We have identified various areas in respect of which the application of existing principles to agentic AI is – at least for now – unclear, particularly around liability frameworks, definitional consistency, human oversight requirements, and transparency and explainability requirements.

Although we raise a number of areas of uncertainty, we caution against issuing premature agentic AI guidance. Rather than seeking additional regulatory guidance at this stage, we would welcome opportunities for industry dialogue and discussion with regulators to

explore how existing frameworks apply to agentic AI use cases. We do, however, see merit in reviewing some payments regulations to add flexibility and update safeguard requirements to better enable agentic commerce. We also see merit in considering whether financial advice rules are able to adapt to agentic AI.

We also note that supervisory processes in financial services might in some cases slow adoption.

If you have any queries about our response, please contact Walter McCahon, Principal: Privacy and Data Ethics, at **walter.mccahon@ukfinance.org.uk**.

Key messages

1. **Regulators should focus on working with industry to explore how existing regulatory frameworks can apply to agentic AI, rather than introducing new rules or guidance, and should ensure a balance between regulatory certainty, safety and fostering innovation.**

- We support the UK's principles-based approach to regulating AI, and the Government's five cross-cutting principles with respect to AI. We are also mindful of, and supportive of, the FCA's stated focus on growth as a cornerstone of its strategy, and the role that digital innovation will play in that growth[1]. The technology-agnostic existing legislation and regulation provides a strong foundation for safe innovation in agentic AI, and the PRA, Bank of England[2] and FCA[3] have mapped existing rules to the UK's five principles in their AI updates, demonstrating that the existing regulatory framework can accommodate AI innovation. The ICO has also already provided extensive guidance on the application of data protection principles to AI. In principle, agentic AI does not need to be treated differently.
- That said, the increased autonomy and opaqueness of agentic AI do present particular risks, and existing laws and regulations were not drafted with agentic AI in mind. This results in a lack of clarity in various areas, particularly around liability frameworks, definitional consistency, human oversight requirements, and transparency and explainability obligations. For the financial services sector to confidently leverage the benefits of agentic AI, it is important that regulators work with industry to explore these areas of uncertainty further. Clarity is required, but additional regulation and guidance should not be introduced prematurely, and any regulatory intervention should be made with input from industry. This will ensure regulators have a sufficiently deep understanding of the technology to enable them to take pragmatic and tech-aware approaches to regulation that protect against consumer harm without stifling innovation. It is important that supervisory approaches and expectations for agentic AI are proportionate, targeted to specific points of genuine uncertainty and risk-based to avoid slowing development and deployment of beneficial use cases.

[1] FCA letter on a new approach to ensure regulators and regulations support growth
[2] The Bank and the PRA's response to DSIT/HMT: update on our approach to AI | Bank of England
[3] AI Update

2. Clear and interoperable definitions are needed for consistent risk assessment.

- Currently there is no consistent definition of 'agentic AI', which makes it difficult to classify risk and to communicate clearly internally, across the sector and with regulators. We would encourage establishing a shared understanding of what 'agentic AI' means, both within the UK and internationally, ensuring alignment and interoperability with international standards, to enable firms to assess and classify associated risks consistently across borders. We note that this is not necessarily a straightforward task, with government already issuing research notes on the topic.[4]

3. Liability is a key area of uncertainty.

- AI in general presents liability and accountability challenges, and these are only increased when it comes to agentic AI. The involvement of multiple agents, acting autonomously (and potentially on behalf of consumers in ecosystems such as agentic commerce), makes it difficult to assign liability when something goes wrong within that chain. A flaw in one agent can cascade errors across multiple interconnected agents, amplifying negative impacts and making it hard to identify the source and responsibility for the error. We encourage collaboration between regulators, government, industry and civil society to clarify how existing rules and frameworks on liability apply, particularly when agents act on behalf of consumers. Consideration also needs to be given to the responsibility of the technology companies who are developing these technologies, e.g. when it comes to the use of LLMs and agents to perpetuate fraud.

4. Expectations as to human-in-the-loop requirements in agentic AI are unclear.

- A number of different AI-related laws and regulations (including the EU AI Act and the EU and UK GDPR) contain various provisions relating to human-in-the-loop requirements. Ethical principles also suggest that some level of human oversight is a fundamental requirement of responsible AI. However, this can be seen as being fundamentally at odds with agentic AI, which is predicated on agents taking decisions autonomously without human input. It may be extremely difficult to monitor every step of the decision-making process, so exploration of risk-based human supervision with appropriate oversight of outputs and outcomes would be welcomed. Firms do not currently have clarity on regulatory expectations as to what level of human oversight should be maintained, and how and when that

[4] See in particular: https://www.gov.uk/government/publications/ai-insights/ai-insights-agentic-ai-html

should be implemented, which makes it difficult to be confident of full compliance and could impact efforts to build consumer trust.

5. **It is not clear how transparency and explainability principles should be applied in agentic AI.**

- There is already helpful guidance on transparency and explainability in AI generally, including in the FCA's AI Update (referencing Consumer Duty and UK GDPR requirements) and the ICO's AI explainability guidance, developed in conjunction with the Turing Institute. However, there remains some uncertainty around the application of these principles in the specific context of agentic AI, with peculiarities for different use cases. In financial services, agents can make or support lending, pricing and eligibility decisions but best practice for explanations of lending decisions is evolving, even where AI is not involved. The extent to which firms need to explain lending decisions and reasoning when agentic AI *is* involved is more complicated still, particularly given the technical challenge of explaining complex, multi-agent decision-making processes in ways that can be easily understood by impacted individuals. Discussions between industry and regulators on this topic, and how these obligations can be managed, would be helpful.

Response to questions

Question 1: What are the key regulatory challenges you face when adopting or developing agentic AI?

Agentic AI intersects with multiple regulatory regimes, particularly in a heavily regulated sector like financial services. Existing regulations provide a strong foundation, but were not designed with autonomous agents in mind, and there is a degree of uncertainty around compliance, particularly in areas like liability, data protection, and consumer rights. This uncertainty can discourage investment and slow down innovation, and clarity is needed generally on how existing requirements should be applied in the context of agentic AI.

Specific regulatory challenges identified by members are set out below:

- **Definitional ambiguity and taxonomic inconsistency.** The absence of a consistent definition of 'agentic AI' represents a barrier to regulatory compliance and risk classification. There is no consensus on what constitutes 'agentic AI'; terms such as 'agentic AI' and 'AI agent' are often used interchangeably, with no clear agreement as to the distinctions between the different terms. This creates practical challenges. If firms do not know what makes a system 'agentic' from a regulatory perspective, it makes it difficult to classify risk consistently and communicate clearly with policy teams, auditors and regulators. We encourage

the DRCF and individual regulators to work with industry and government to create a shared vocabulary or framework for what constitutes 'agentic AI' (or potentially 'degrees' of agentic AI), ensuring it is aligned and interoperable with international standards, so that firms are able to confidently classify risk and accurately communicate in relation to the technology. We note that there is helpful work moving in this direction already.[5]

- **Transparency and explainability requirements.** Explainability is a critical challenge in the context of agentic AI, particularly for financial services where agentic AI can make or support decisions such as lending, pricing and eligibility. Regulators expect clear documentation on how decisions are made and detailed audit trails (to comply with Consumer Duty and UK GDPR requirements, as per the FCA's mapping work in its AI Update), but the reasoning used by agentic AI is often opaque, and the 'black box' nature of agentic workflows makes it difficult to meet expectations around transparency and accountability. There is also uncertainty around when and how to disclose the use of AI to impacted individuals, particularly for systems that sit behind the scenes and do not directly influence individual outcomes. As is broadly understood, there may also be difficult trade-offs between performance and explainability in some cases.
- **Human-in-the-loop requirements vs autonomous capabilities.** Current regulatory expectations around human oversight (stemming from laws such as the EU AI Act and GDPR) create fundamental tensions with the autonomous nature of agentic AI. It is often difficult to see where and how human oversight can be incorporated into workflows that are, by their nature, autonomous. There is potential for machine oversight, and for risk-based approaches that focus on monitoring outputs and outcomes rather than every step of the decision-making process. As with the other areas of challenge, we encourage public-private sector collaboration, including through the mechanisms detailed in our response to question 4 below, to explore how and when human oversight should be maintained in the context of agentic AI, especially in high-risk or consumer-facing contexts.
- **Liability and accountability frameworks.** Liability allocation is one of the most critical areas requiring improved clarity. This is already a challenge in AI in general, and is increased when it comes to agentic AI. Agentic AI involves distributed agency, with chains of multiple agents all acting autonomously to make decisions and take actions on behalf of organisations or individuals. There

[5] See for example: https://www.gov.uk/government/publications/ai-insights/ai-insights-agentic-ai-html and https://www.ibm.com/think/topics/agentic-ai

are currently no clear frameworks for determining liability within these chains. Particular examples cited by members include:

- Agentic commerce flows, where it is not clear who would bear responsibility if, for example, an agent acted on behalf of a customer to make a decision that resulted in financial loss. Members have raised the need for clarity regarding authentication protocols where an agent is acting on behalf of a customer.
- APP scam mandatory reimbursement requirements, where it is not clear whether the use of an agent impacts the threshold for 'gross negligence'.

- We encourage collaboration between government, regulators, industry and civil society to clarify how existing rules and regulations around liability apply to agentic AI use cases.

- **Consumer Duty, consumer protection and vulnerable customer considerations.** Agentic AI presents both risks and opportunities for consumer protection and management of vulnerable customers. Agentic AI can offer positive outcomes for customers and vulnerable groups, e.g. through improved accessibility. Conversely, unreliable results may have a disproportionate impact on vulnerable customers. It is important to be clear that technology companies making agentic tools available to consumers also have responsibilities to those consumers. Existing consumer protection frameworks may not fully account for these scenarios. For example, agents managing subscriptions should not auto-renew without clear consent and should not misrepresent terms.
- **Data protection and governance.** Agentic AI raises a number of questions about how the existing data protection framework in the UK should be applied, particularly in the context of ecosystems like agentic commerce. For example, concerns exist around the dynamic and autonomous nature of agentic decision-making, purpose limitation, transparency, data minimisation and consent. The UK GDPR is rightly technology-agnostic, and agentic AI should not be unnecessarily treated differently from other technology, but firms would welcome regulatory efforts to work with industry through dialogue and collaboration to clarify these areas.

Question 2: Are there specific areas (e.g. data protection, liability, consumer protection, competition, copyright) where clarity is most needed?

The key areas members have identified as requiring regulatory clarity (following on from the above) are set out below. However, we reiterate our position that these should not, at this stage, be clarified through the introduction of new guidance, but rather through engagement between regulators and industry (including as set out in our response to question 4 below):

- **Liability and accountability:** It may be unclear who is accountable when agentic AI systems cause harm or make errors, in particular when agents carry out action on behalf of a customer / individual without them being present (such as in agentic commerce). There may also be ambiguity in accountability for cascading failures and autonomous decisions; in agentic AI, a flaw in one agent can cascade errors across multiple interconnected agents, amplifying negative impacts and making it difficult to establish where the original failure occurred and where responsibility sits.
- **Consumer protection:** Existing consumer protection frameworks may not fully account for agentic AI. Potential misalignment with human intent is an overall challenge faced across sectors, as agentic systems may act in ways that conflict with the user's actual intent even when technically fulfilling assigned goals. This makes it crucial to get verification and documentation of consumer intent right (e.g. in agentic commerce where AI agents may be acting on behalf of consumers) and to adjust responsibilities to the role each actor plays in the agentic AI value chain. The responsibility of technology companies making agentic tools available to consumers to protect against fraud or consumer harm must also be considered.
- **Data protection and privacy:** Autonomous agents often process personal data in real time, and it is unclear how consent, purpose limitation, and data minimisation principles apply when decisions are made dynamically and without direct user interaction. Other key areas of uncertainty in the data protection context are transparency and explainability requirements, risks of oversharing across systems (or with humans), associated cross-border data flow concerns, and potential unwanted or unauthorised disclosure of information.
- **Vendor and third-party accountability:** There is a lack of clarity around obligations and oversight mechanisms for third-party AI components, and concerns about third / fourth party concentration risks and consequent resilience issues. Information sharing from AI providers is needed to support due diligence, requiring a balance between supporting consumers of AI and the IP of firms who are developing them. Various regulatory requirements, including the PRA's

Supervisory Statement on Model Risk Management[6], require firms to verify and validate vendor solutions to the same standard as internal expectations, but vendors may be unwilling to provide sufficient documentation to evidence compliance. Members have raised the issue of a lack of clarity as to third party accountability, including under the Critical Third Parties regime, when an agentic system drifts outside of its documented design. Firms may struggle to contractually influence big tech firms providing agentic technology; regulatory support in terms of setting out expectations on vendors as to due diligence and information provision may be helpful, to ensure firms have comfort from third parties as to their roles and responsibilities in ensuring compliant and safe systems.

- **Competition:** We understand that the CMA intends to release further guidance on AI, following its previous guidance on algorithms published in 2021 and strategic update on AI in 2024. This further guidance is expected to reflect the CMA's approach to AI generally, which is to ensure that commercial use of AI continues to support the objective of maintaining and promoting competitive dynamism in markets and achieving pro-competitive effects. There is a risk that the agentic AI market could see a small number of players exercise market power, and competition considerations will be important to ensure a diverse and competitive market.
- **Know your agent (KYA) protocols and verification standards:** With the introduction of agentic commerce, a key challenge is ensuring that agents acting on behalf of users are legitimate, authorised, and aligned with user intent. The risk of malicious or misconfigured agents undermines confidence in agentic AI ecosystems. Members support the development of know your agent (KYA) protocols that help verify agent identity, intent, and authorisation when an agent is acting on behalf of a consumer - supporting secure and transparent agentic transactions. Input on concepts such as Know Your Agent (KYA), dynamic consent, and liability frameworks would help shape proportionate safeguards aligned with consumer protection and data governance principles. Agentic commerce relies on secure, verifiable identity frameworks, and accelerating digital ID initiatives, in line with the Data (Use and Access) Act 2025, is also essential to ensure transactional integrity and accountability.

[6] https://www.bankofengland.co.uk/-/media/boe/files/prudential-regulation/supervisory-statement/2023/ss123.pdf

Question 3: How do current regulations support or hinder innovation in agentic AI?

Supporting elements:

- **Principles-based regulation:** The UK's regulator-led, principles-based approach provides flexibility for firms to interpret and apply existing rules to novel technologies like agentic AI, while avoiding premature or overly prescriptive regulation. This helps to support innovation, encourage agility and enhance UK competitiveness, and supports the government's growth agenda.
- **Existing regulatory frameworks:** The UK already has robust regulatory frameworks in place, particularly in data protection and consumer protection, which provide a strong foundation for responsible AI innovation, including with agentic AI.
- **Regulatory sandboxes:** Initiatives such as the FCA's Digital Sandbox and Supercharged Sandbox offer a safe environment for firms to test agentic AI applications under supervisory oversight, providing access to synthetic data, technical resources, and regulatory guidance, helping firms validate agentic models before live deployment.
- **FCA approach:** We support the FCA's approach of not creating new rules for AI, which acknowledges the technology-neutral nature of existing laws and regulations, and avoids regulatory overreach.

Hindering elements

- **Regulatory uncertainty:** While we support the UK's flexible principles-based approach, there is inevitably a level of regulatory uncertainty when applying existing rules (which were not drafted with agentic AI in mind) to new technological developments. Currently regulations do not explicitly stop innovation – and, indeed, the UK's approach is explicitly intended to foster innovation – but the lack of clarity can be a hindrance. For example, high-level principles like 'AI must be explainable and fair' are reasonable, but without practical examples or templates, they can be difficult to operationalise. The extra documentation and governance steps can slow down experimentation, even for low-risk use cases.
- **Specific regulatory barriers:** The financial services sector is juggling a significant volume of regulation. Certain elements of that regulation can make it difficult for firms to adopt agentic AI at the same scale and speed as other industries, and may even cause challenges for other sectors and organisations looking to leverage the benefits of agentic AI, particularly in the context of agentic commerce. For example, existing strong customer authentication rules limit the

extent to which agents can make multiple purchases on behalf of customers. We support requiring robust evidence of customer intent behind transactions, but the current rules around authentication (while providing some flexibility through the application of exemptions) are insufficiently outcomes focused. The development of agentic commerce provides an opportunity to update these rules in a way that manages the risk of unwanted or fraudulent transactions, while at the same time making it possible for agents to act with a greater degree of autonomy. This is one area where we do see merit now in a potential revision to existing rules, which can form a part of a review of the Payment Services Regulations 2017 and the Electronic Money Regulations 2011, including such key concepts as 'knowledge', 'possession' and 'inherence'.[7]

- **Potential for delays and high regulatory burden:**
    - The PRA's Model Risk Management supervisory statement applies to all models used for 'informing business decisions'. The rapid expansion of generative AI has greatly increased the number of models in scope, significantly increasing the workload for model risk management (MRM) teams. This trend will likely continue with the growth of agentic AI.
    - Related to this, it is important that supervisory approaches to agentic AI are reasonable and proportionate, to enable adoption of beneficial agentic AI use cases to be accelerated, rather than low-risk cases being slowed down unnecessarily. Though we note that experiences vary from member to member, some firms advise that some use case approvals have taken a longer time than the risk profile would seem to require, including an 18-month period in one case. We are uncertain of the reason for this, but speculate that staffing constraints, access to skills and regulatory caution might all be relevant. (We do note that the PRA is not a member of the DRCF, however).

[7] In reviewing these regimes more generally under the Future of Payments Review we have advocated for:
1. Outcome-based and market-driven approaches, shifting away from prescriptive regulation.
2. Consider outcomes for all end users, including consumers, merchants and corporates.
3. Same activity, same risk, same regulation, applied proportionately.
4. Technology-neutral regulation.
5. Creating a level playing field, recognising the need for commercial arrangements.

Question 4: What kinds of informal advice, engagement, or support would be most helpful?

We advocate for regulatory support that allows for exploration in safe environments, alongside regular regulatory dialogue and early industry engagement to ensure regulatory approaches are shaped around technological capabilities. This is a better approach than rushing to produce AI guidance, then needing to retrofit a complex new development into guidance provided prematurely. Particular areas that members have raised as being helpful include:

- **Early and continuous industry engagement.** For a technological development as complex as agentic AI, it is essential that regulators understand how the technology works, and we advocate for a 'bottom up' approach to regulation and regulatory intervention, rather than a 'top down' approach. Engagement with industry players through roundtables, sector-specific workshops, working groups, advisory groups and bilateral discussions would help regulators understand the complexities of agentic systems and their interaction with existing rules before formal interventions are considered. These processes also help inform industry of emerging regulatory thinking. Members also raise the possibility of short 'office hour' style sessions with regulators to help surface common challenges and build consensus on governance approaches. Generally, opportunities where different participants can work together and share knowledge and learnings about agentic AI benefits, risks and guardrails (within an appropriate competition law framework) would be welcomed.
- **Regulatory sandboxes and testing environments.** Regulatory sandbox environments (akin to the FCA's Digital Sandbox and AI Live Testing Service) will allow firms to test agentic AI applications under supervisory oversight, using synthetic data and receiving informal guidance during development. Some members called for a tailored sandbox specifically for agentic AI. This could help flush out and clarify issues around liability, data protection and consumer rights and how they apply in practice. Furthermore, it would be helpful if regulators could share learnings from these sandboxes to support firms' development of governance around agentic AI.
- **Standards development support.** We fully support the development of universal standards and industry engagement to help develop those standards, including through input on concepts such as Know Your Agent (KYA), dynamic consent, and liability frameworks to help shape proportionate safeguards aligned with consumer protection and data governance principles.
- **Cross-regulatory and international coordination.** Coordination between UK regulators (e.g. FCA, ICO, CMA) would support consistent treatment of agentic systems and avoid fragmentation across regimes. It would be helpful to

understand how the DRCF members propose to work together to understand and regulate agentic AI. Informal cooperation with international counterparts can further help align privacy standards and support interoperable data exchanges essential for the success of agentic systems. Establishing clarity on definitions on an international level can also help ensure consistency and interoperability across jurisdictions, and avoid fragmented or siloed approaches.

- **Specific policy areas.** As well as the general comments above, members identified the following specific policy areas where support would be valuable:
    - **Financial advice boundary:** Consideration and dialogue on best practice in the adoption of customer-facing agentic AI applications would be useful, with one live initiative being the ongoing FCA Advice Guidance Boundary Review. We feel this review of the financial advice regulations is useful, including a consideration of how those regulations apply in the context of agentic AI, and whether those requirements are flexible enough to facilitate further innovation in this space. We observe that LLMs that sit outside the financial services regulatory perimeter can be prompted by users to provide customers with bespoke and detailed financial advice. It is important that regulated providers are able to compete on a level playing field.
    - **Data sharing:** The Data (Use and Access) Act 2025 allows for smart data schemes to be set up, including cross-sectoral schemes. While the initial focus has rightly been on Open Banking, Open Finance and the Long-Term Regulatory Framework, increasing the momentum behind cross-sector schemes would support agentic AI use cases like agentic commerce.
    - **Digital identity and payments:** Informal support for industry-led initiatives on digital ID, tokenisation, and programmable payments (e.g. stablecoins) can accelerate adoption of agentic commerce or other agentic AI ecosystems involving agents acting on behalf of consumers, with regulatory clarity on their use – particularly around spending limits, oversight, and integration with existing payment security standards – being key to enabling safe adoption.
- As we have raised several times in this response, there is a strong feeling among members that there is not a need or desire for additional, agentic AI-specific guidance at this time, and focus should instead be on collaboration and dialogue as described above. However, if and when a need for further guidance does become apparent, it is important that such guidance is both practical, and developed in close collaboration with industry. Such guidance could help firms navigate uncertainty and avoid disproportionate compliance burdens. In the

future, sector-specific guidance on acceptable risk thresholds and mitigation strategies may be valuable, along with practical examples and case studies enabling firms to learn from others as to how they have balanced innovation with compliance. Templates, designed in collaboration with businesses, may also be helpful in due course to operationalise high-level principles.

Question 5: Are there sector-specific concerns (e.g. legal services, finance, telecoms) that should be considered?

As the collective voice for the financial services industry in the UK, we have focussed our response here on considerations and concerns specific to the financial services sector.

- **Payments and agentic commerce.** Agentic AI introduces a shift in how decisions are made and how actions are executed, often without direct human involvement. Initiatives like Google's Agent Payments Protocol and the Agentic Commerce Protocol codeveloped by OpenAI and Stripe, where AI agents initiate and complete transactions on behalf of users, create particular regulatory challenges from a Payments perspective. In particular:
    - **Strong Customer Authentication:** Existing SCA rules may need to be updated (as explained above).
    - **Real-time decision-making:** Agentic systems may approve or reject transactions instantly, requiring robust safeguards to prevent errors, abuse, or unintended outcomes. Informal engagement can help shape proportionate oversight mechanisms that reflect the speed and autonomy of agentic systems.
- **Fraud prevention and risk management.** Autonomous systems must detect and respond to threats without human input, while remaining auditable and accountable, with informal feedback on liability frameworks and agent authorisation protocols, such as KYA, supporting this balance. Agentic systems must be resilient to fraud, manipulation, and infrastructure failures, especially in high-stakes financial environments. Agentic systems could also be used to *perpetrate* fraud, and consumers may be put at risk of fraud and bad outcomes as a result of agentic tools being made available to them by tech companies or others in the absence of appropriate safeguards.
- **Consumer Duty.** In a future where customers rely on AI agents to manage their financial affairs, they may not have as much (or any) direct engagement with their financial services providers, who may themselves develop agents to interact with their customers' agents. In this type of environment, how financial services providers discharge their obligations under the Consumer Duty (for example, to

support customers to pursue their financial objectives, and avoid foreseeable harm) is likely to change considerably.

- **Operational resilience.** Agentic AI failures could breach PRA / FCA operational resilience requirements, for example if interconnected agentic systems experience correlated failures or if autonomous agents cause unexpected disruptions to critical business services. However, in keeping with members' limited appetite for additional guidance or legislative/regulatory changes at this stage, members feel that the most helpful solution would be dialogue with regulators on how firms can meet these requirements while deploying agentic systems, rather than amendments to operational resilience requirements to accommodate agentic AI.

Question 6: What future risks or opportunities do you foresee in the use of agentic AI?

Agentic AI brings significant opportunities for the financial services sector, and it is important that the supervisory approach across DRCF regulators enables firms to leverage those opportunities in line with the FCA's growth agenda.

Opportunities

- **Productivity and efficiency:** Significant productivity and efficiency gains could transform customer experiences. Agentic AI can simplify and speed up processes, and automate complex workflows, freeing up staff for higher-value activities and reducing cost.
- **Consumer empowerment:** Agentic systems can act on behalf of users to manage payments, subscriptions, and financial decisions. This offers greater convenience, control, and personalisation. Customer expectations may increasingly evolve to require the empowerment and efficiency brought about by agentic AI, but at the same time, customer trust will be vital to the success of agentic systems. A common understanding of necessary guardrails between industry and regulators can help to foster that trust.
- **Operational innovation:** Autonomous agents can streamline complex workflows, reduce friction in digital commerce, and enable new business models, particularly when paired with programmable payments and smart contracts.
- **Security enhancement:** With the right safeguards, agentic systems can strengthen transactional integrity through cryptographic authentication, auditable activity logs, and tokenisation frameworks.
- **Operational resilience:** Although operational resilience is raised above as a sector-specific challenge, agentic AI also has the potential to *improve* operational

resilience. For example, agentic systems can enable self-healing of operational activities and smarter decision support across front and back office, along with transformation of technology functions including acceleration of the software lifecycle.

Risks

- This response has already covered many of the risks of agentic AI in detail. For brevity, we do not propose to repeat all of the risks identified above, but members have identified the following additional risks that it is worth flagging here:
- **Security vulnerabilities:** Agentic AI increases the threat vector by expanding the potential attack surface; numerous autonomous agents mean an increased number of targets. The multi-agent nature of agentic systems also means that one compromised agent can rapidly propagate threats across different agents. This means the cybersecurity risk in connection with agentic AI systems is heightened, with the potential to cause failures across the market in the context of interconnected agentic AI ecosystems. This will necessitate robust security-by-design principles and continuous monitoring.
- **Workforce impact:** Workforce replacement is a risk of AI generally, but particularly agentic AI. Various organisations, including in the financial services sector, have recently announced workforce reductions as a result of increased use of AI. Agentic systems are inherently designed to increase automation and reduce human input, making further workforce replacement likely and raising potential employment concerns. There is also a risk of automation outpacing skills development, leading to a potential skills gap.
- **Consistency and non-deterministic outputs:** Reliance on unstructured data and multilayered decision systems complicates evidencing data lineage and ensuring consistency. Agentic systems often produce subjective, non-quantifiable outputs, complicating risk quantification and the definition of ‘good’ behaviour and outputs.
- **Management of agent sprawling:** As agentic AI systems scale, there is a risk of a proliferation of autonomous agents across business processes. This creates governance challenges, including:
    - Maintaining visibility and control over the number and scope of agents deployed.
    - Preventing duplication of functionality and ensuring agents operate within defined boundaries.
    - Implementing lifecycle management to retire or consolidate redundant agents.

- **Fragmentation across frameworks:** We note that there is a risk that regulatory expectations in relation to agentic AI could diverge, as authorities' positions evolve. This would reduce regulatory clarity, slow down innovation and lead to hesitancy in agentic AI adoption and experimentation. This would risk discouraging experimentation. Similarly, fragmented global standards could emerge, resulting in heightened complexity and costs, creating barriers for multinational firms.