BCLP. Client Intelligent

Insights

THE MILLS REVIEW REVISITED: WHERE AUTONOMOUS AI IS TESTING THE UK REGULATORY FRAMEWORK

Sep 15, 2026

Summary

The FCA's position has been that the UK's existing regulatory framework can accommodate AI and that AI-specific regulation is not currently required. The Mills Review (the "**Review**"), published by the FCA on 6 July 2026, broadly supports that view. It should not, however, be read as signalling regulatory restraint. Rather than proposing new rules, it anticipates heightened expectations around governance, oversight and evidencing compliance as firms adopt increasingly autonomous AI systems.

At the heart of the Review is a five-stage "autonomy spectrum", under which humans move from making decisions to supervising AI-generated outcomes. As delegation increases, accountability becomes harder to trace and meaningful human oversight more difficult to demonstrate. Against that backdrop, the Review raises a broader question: whether firms can continue to demonstrate effective oversight, accountability and regulatory compliance as decision-making becomes increasingly delegated to AI. The FCA's September 2026 review of frontier AI and cyber resilience brings that evidential challenge into sharper focus: its practical observations emphasise that effective AI deployment depends on the governance, controls, human judgement and operational environment around the model, rather than model capability alone.

The autonomy spectrum

The Review explains AI deployment through a five-stage "autonomy spectrum": operator (AI assists with tasks), collaborator (human and AI work together), consultant (AI recommends and the human decides), approver (AI proposes actions for human approval) and observer (AI acts within agreed parameters while the human oversees outcomes). At each stage, the human role shifts further away from direct decision-making and towards supervision.

The significance of the autonomy spectrum is not technological but regulatory. The Review suggests that existing frameworks, including the Consumer Duty, Senior Managers & Certification

Regime ("**SM&CR**") and operational resilience requirements, are generally capable of accommodating AI at the lower end of the autonomy spectrum. The more significant point is that traditional assumptions about decision-making become less reliable as autonomy increases. As humans move from making decisions to supervising outcomes, concepts such as oversight, challenge and accountability may become increasingly difficult to demonstrate in practice.

SM&CR, accountability and autonomous AI

The Review emphasises the importance of personal accountability under the SM&CR in relation to the implementation and use of AI systems viewing it as a potential bulwark against a 'growth at all costs' approach. The Review concludes that, for the time being, the SM&CR remains broadly appropriate for ensuring accountability in AI-enabled environments and does not recommend any fundamental changes to the regime.

However, the report acknowledges that demonstrating accountability will become more challenging where AI systems operate at greater levels of autonomy (i.e. the approver and observer stages), where humans oversee outcomes rather than making decisions. The report characterises the “observer” stage as a "pressure point", where meaningful human control is likely to be difficult to evidence and accountability becomes harder to exercise in practice. This raises the question of how senior managers demonstrate “reasonable steps” in higher-autonomy environments. The strain on accountability may be compounded in an environment where there are interconnected AI systems and with the rise of third-party model providers.

The Review notes two developments that may serve to mitigate the risk of the SM&CR becoming less effective as AI systems become more autonomous. First, future AI models might become less probabilistic and less opaque than current models. Whilst ultimately a matter for technologists, this observation appears counterintuitive given the corresponding increase in complexity that more autonomous systems are likely to entail. Second, assurance tools (including pre-deployment and ongoing checks) could better enable senior managers to take (and demonstrate) "reasonable steps" to prevent regulatory breaches or misconduct.

The paper acknowledges the industry’s call for clearer guidance on what the regulator considers to constitute “reasonable steps” in the oversight and management of AI systems, and that providing such guidance will increase firms’ confidence in adopting delegated and autonomous AI. In turn, this will promote long-term growth and competitiveness.

Senior managers responsible for AI systems will need to familiarise themselves with the regulator’s evolving expectations of how such systems are to be safely implemented, monitored and controlled. The paper makes clear that these expectations extend to the monitoring of system behaviour, testing beyond deployment, the implementation of controls that detect systems moving outside expected bounds and responding to evolving model risks, including drift, bias and opacity. The practical challenge for senior managers will be ensuring that they are receiving the appropriate management information, are sufficiently equipped (and educated) to interpret it and are reacting appropriately to issues as they arise.

Operational resilience

The Review concludes that the UK's operational resilience framework remains broadly fit for purpose. However, it warns that AI may introduce a new category of systemic risk that sits beyond the reach of traditional firm-level resilience assessments.

That concern has become more immediate. In September 2026, the FCA published observations from its multi-firm review of frontier AI and cyber resilience. While expressly introducing no new rules, guidance or regulatory expectations, the FCA found that frontier AI is already testing firms' governance, risk ownership, remediation capacity and dependency mapping. Its central practical message is that the value of frontier AI depends less on the model itself than on the governance, tooling, controls, human oversight and operational environment around it.

The concern is not simply that individual firms may experience outages or technology failures. As firms increasingly rely on common AI models, datasets, cloud providers and infrastructure, a single point of failure could affect multiple firms simultaneously. Equally, the use of similar models for underwriting, pricing and investment recommendations may lead to correlated decision-making, herding and the narrowing of consumer options, amplifying market-wide risks.

The Review therefore shifts the focus from firm-specific resilience to ecosystem resilience. Shared dependencies may be difficult for regulators to identify until disruption occurs, creating risks that no individual firm can see in isolation. In response, the Review calls for greater cross-firm monitoring and regulatory coordination to identify common vulnerabilities and concentrations of risk.

Many of the solutions identified in the Review are directed at the regulator rather than the industry. The report recognises that the FCA's own Supervision team may need to look across the system as a whole: monitoring correlations between firms and mapping shared dependencies, whilst stress-testing how systems behave together. However, given the likely dependency on a concentrated number of models, and the present absence of a dedicated AI regulator, coordination between UK sectoral regulators and international partners will be necessary. This will involve collective monitoring of concentration and ecosystem vulnerabilities across critical providers. Looking ahead, the Review identifies a need for more structured international supervisory coordination and, where appropriate, sector-specific international standards to complement national frameworks. However, that remains some way off.

The regulatory perimeter

The Review concludes that the UK's activity-based regulatory perimeter remains broadly fit for purpose. However, it identifies growing pressure points as consumers increasingly rely on general-purpose AI tools to understand and compare products and services and make financial decisions. Such tools may exert significant influence over consumer outcomes without being subject to the same regulatory safeguards and accountability mechanisms, because they are not authorised persons and therefore sit outside direct FCA regulation. The underlying concern is a potential disconnect between influence and accountability. Consumer decisions may increasingly be shaped by AI tools sitting outside the regulatory perimeter, while responsibility for resulting outcomes remains with regulated firms.

The Review also highlights the risk of competitive asymmetry. While regulated firms remain subject to conduct requirements and financial promotion restrictions, AI platforms and model providers may shape consumer decisions without being subject to equivalent obligations. As AI-mediated customer journeys become more prevalent, traditional distinctions between regulated and unregulated activity may become increasingly difficult to maintain. In recognition of this, the Review recommends a perimeter review to examine how existing regulatory concepts operate in AI-enabled customer journeys, including the advice/guidance boundary, the "by way of business" test and the financial promotions regime.

The Agentic Supervisory Model

The Review's proposed Agentic Supervisory Model may ultimately prove to be one of its most significant recommendations. If firms increasingly rely on AI, regulators will need to do the same.

The proposal is not to replace human supervisors. Rather, the FCA would deploy AI tools across the regulatory lifecycle, including authorisation, supervision and enforcement, while retaining human responsibility for judgment and decision-making. AI would assist in analysing information, identifying risks and prioritising supervisory activity, enabling the FCA to intervene more quickly and effectively.

The broader significance of the recommendation is that it moves supervision beyond the individual firm. The Agentic Supervisory Model is intended to identify cross-firm patterns, common model dependencies and system-wide risks that may not be apparent through traditional supervisory approaches. For firms, that is likely to result in greater expectations around monitoring, governance and the ability to provide structured evidence of AI-related decision-making and outcomes.

What should firms be doing now?

Most of the Review's recommendations are directed at the FCA and the Government rather than regulated firms. Whether through a perimeter review, further guidance on accountability or the development of the Agentic Supervisory Model, much of the proposed change will take time to materialise. That does not mean firms should wait. The FCA's more recent frontier AI work reinforces the practical importance of clear ownership, proportionate guardrails, effective human oversight and the capacity to act on AI-generated findings.

The Review's central message is that the existing frameworks already apply in AI-enabled environments. Firms that defer consideration of governance, oversight, resilience and perimeter issues until formal regulatory change arrives may find themselves behind supervisory expectations.

Firms should in particular:

1. Review governance and accountability frameworks to ensure that responsibility for AI deployment, oversight and decision-making is clearly allocated and capable of being evidenced.

2. Ensure that senior managers responsible for complex and increasingly autonomous AI systems understand the risks and controls across the AI lifecycle, receive appropriate management information, and are equipped to interrogate it.

3. Maintain a reliable inventory of AI systems and models, with consistent risk-based classification and an assessment of the level of autonomy at which each system operates.

4. Assess whether controls for higher-autonomy models remain sufficient to demonstrate compliance with obligations including the SM&CR, Consumer Duty and operational resilience requirements.

5. Review the firm's reliance on third-party AI providers and infrastructure, particularly where multiple business-critical processes depend on the same models, datasets or cloud providers.

6. Assess customer-facing AI journeys, with particular attention to advice, guidance and financial promotion risks, and the extent to which meaningful human oversight remains possible to deliver good customer outcomes.

Conclusion

While the headline of the Review is that the current regulatory framework does not require reform, its findings identify the points at which that framework may come under strain. The report recognises that agentic AI will put pressure on the suitability of the framework. The FCA's more recent work on frontier AI suggests that some of those pressures are already emerging in practice, particularly in relation to governance, remediation, third-party dependencies and operational resilience.

In the meantime, firms are left to apply the existing regulatory expectations to an ever-evolving landscape, in which (inevitably) there is commercial pressure to implement AI and maximise its potential. As firms' usage becomes more autonomous, and decision-making becomes increasingly delegated, distributed and difficult to interrogate, the challenge will be to maintain, and evidence, effective accountability and oversight. Firms should be able to explain where AI is used, who owns the risk, how performance and outcomes are monitored and what happens when systems or providers fail. The existing regime may remain sufficient, but existing approaches to compliance may not.

Related capabilities

- Financial Services
- Financial Regulation Compliance & Investigations
- Investigations

Meet the team

David Rundle

Partner, London

david.rundle@bclplaw.com
+44 (0) 20 3400 4027

Samantha Paul

Knowledge & Innovation Counsel, London

samantha.paul@bclplaw.com
+44 (0) 20 3400 3194

Caroline Cwiertnia

Associate, London

caroline.cwiertnia@bclplaw.com
+44 (0) 20 3400 2144