Solutions by Jewel – Written Evidence (RSB0011)

House of Lords Financial Services Regulation Committee
Inquiry: Regulatory sandboxes in financial services

Submitted by Jewel Darlington, Chief AI Governance Officer & Founder
Organisation Solutions by Jewel (SBJ)
Contact **[REDACTED]**
Website **[REDACTED]**

Questions addressed 8, 8(a), 9 and 10
Deadline 5:00pm, 23 October 2026

Summary

1. Solutions by Jewel (SBJ) welcomes the Committee's inquiry. This submission answers Questions 8, 8(a), 9 and 10. It concentrates on agentic AI: systems that plan and take actions, including payments, within limits set by a firm or a consumer.
2. Agentic AI is now a stated focus of FCA testing programmes. The second AI Live Testing cohort includes agentic payments use cases, and the second Supercharged Sandbox cohort is testing solutions for safer agent-led payments and commerce [1][2].
3. Sandbox tests have long asked whether a product works, whether consumers are treated fairly and whether risks are contained. Agentic AI adds a fourth question: what the agent is authorised to do, on whose behalf, and whether that authority still holds at the moment the agent acts.
4. SBJ recommends five measures for agentic tests:
   (a) an authority record for each agent, completed before testing begins;
   (b) authority checks at the point of consequential execution, with each check logged;
   (c) reassessment triggers that continue to apply after the firm leaves the sandbox;
   (d) a common evidence format across FCA testing programmes, so findings can be compared and published; and

(e) a published, anonymised report of agentic AI findings, timed to inform the review of payments regulation.

About SBJ and the basis of this evidence

5. SBJ is an AI governance advisory firm serving mid-market and enterprise organisations. Its work covers authority, accountability and evidence for AI systems and AI agents. SBJ publishes governance resources at solutionsbyjewel.com, including an AI Agent Governance Checklist and an AI Decision Record template [8][9].

6. SBJ has not participated in an FCA sandbox. This evidence draws on public FCA and parliamentary documents and on SBJ's governance work. In September 2026 SBJ responded to Question 15 of HM Treasury's consultation on Modernising Payment Services Regulation, which concerned agentic payments [10]. Parts of this submission build on that response.

Question 10: To what extent are regulatory sandboxes being used to test agentic AI, and does this introduce any new risks for the use of regulatory sandboxes?

Extent of use

7. Public FCA announcements show agentic AI in three AI Lab programmes during 2026:

   (a) **AI Live Testing, second cohort** (announced 21 April 2026). Eight firms are testing use cases that include agentic payments, on models ranging from agentic AI to small language models. Testing concludes by the end of 2026, and the FCA will publish an evaluation report in Q1 2027 [1].

   (b) **Supercharged Sandbox, second cohort** (announced 22 July 2026). Twenty-one organisations, selected from 199 applications, are exploring use cases that include safer agent-led payments and commerce [2].

   (c) **Agentic Academy.** A 10-week specialist programme for selected firms, delivered by the FCA with the Centre for Finance, Technology and Entrepreneurship [2].

8. The Committee's terms of reference also cover the Regulatory Sandbox, the Digital Sandboxes and the Digital Securities Sandbox. SBJ has found no published figures on agentic AI testing within those three. The Committee may wish to ask the FCA and the Bank of England for them.

9. The Mills Review, published by the FCA in July 2026, reports that more than half of industry survey respondents are now piloting or deploying agentic AI, mostly in limited scope and under human oversight [3, p. 14]. It describes a spectrum of autonomy that runs from the human as operator to the human as observer, and notes that accountability becomes harder to trace as systems move along it [3, pp. 6, 12]. The risks set out below grow with the level of autonomy under test.

New risks

10. **Authority risk.** Authentication establishes who is acting. Authority establishes what that party may do. A correctly functioning agent can act outside the authority its principal granted. A test that records outcomes alone may leave that question unanswered. The Mills Review notes that, once the human becomes an approver, reliance shifts onto how well permissions and safeguards are set in advance [3, p. 12].
11. **Change during and after the test.** Models update, and agents acquire new tools and permissions over time. The Mills Review observes that a model tested at deployment may perform differently as its environment changes [3, p. 17]. Sandbox evidence describes the agent as tested. Its continuing value to supervisors depends on knowing when the agent has materially changed.
12. **Boundary risk.** Agents call external systems, data sources and other agents. A test defines which firms, accounts and counterparties are in scope. An agent with live tool access can reach systems outside that scope unless its permissions are bounded and logged. Third-party and multi-agent chains also make it harder to identify which party is responsible for an action; the Mills Review records consumer uncertainty about whether responsibility sits with the firm, the model provider, the platform or the agent [3, p. 48].
13. **Evidence at scale.** Agents can take many actions in a short period. Sampled review can miss a pattern of actions that sit just inside or just outside a limit. Agentic tests need complete, machine-readable records of material authority events.
14. **Signalling risk.** Agentic products may be marketed on the strength of sandbox participation. A clear public statement of the scope of each test, and of the conditions it covered, helps investors, partners and consumers read participation accurately.

Question 9: How effectively do regulatory sandboxes manage risks to consumers and markets, and what safeguards and accountability arrangements are in place?

15. The FCA applies its existing frameworks to AI, including the Consumer Duty and the Senior Managers and Certification Regime, and has stated that it will avoid additional regulations for AI by relying on those frameworks [2][4]. Agentic sandbox tests give the FCA a practical setting in which to show how those frameworks apply to delegated authority.

16. SBJ suggests the following safeguards for each agentic test, scaled to its risk:

    (a) **Authority record.** Before testing, the firm records each agent's identity, the principal it acts for, its authorised actions, its out-of-scope actions, any spending or exposure limits, the source and date of its authorisation, and the accountable senior manager. SBJ's AI Agent Governance Checklist covers these elements [8].

    (b) **Approval points matched to autonomy.** The firm states where the agent sits on the autonomy spectrum and which actions require human approval [3, p. 12].

    (c) **Authority check at execution.** Authority is checked when a consequential action executes, and the check is logged. Consent given at set-up covers the scope and period stated at set-up.

    (d) **Traceability of authority events.** Grants, changes, uses, refusals and revocations of authority are each recorded with a time stamp and a source.

    (e) **Reassessment triggers.** Named events require re-review: a model or version change, new tools or permissions, new counterparties, an incident, or a material change in volume. SBJ's AI Decision Record template includes fields for these triggers and for the owner's sign-off [9].

    (f) **Exit record.** Each test closes with a record of the authority position at exit and the reassessment triggers the firm will apply in live operation.

17. On accountability, each agent under test should map to a named senior manager whose statement of responsibilities covers it. The Mills Review records that firms do not seek to change the Senior Managers Regime accountability model, and that firms and regulators will need to work through how it operates as autonomy increases [3,

p. 9]. Agentic sandbox tests are a suitable place to do that work, and to publish the results.

Question 8: What role should regulatory sandboxes play in informing the development of regulation, particularly for emerging technologies? (including 8(a))

18. Sandboxes inform regulation when their findings are comparable, timely and published. The FCA's 2017 lessons-learned report on the first two Regulatory Sandbox cohorts, which covered 146 applications and 50 supported firms, is a useful precedent [5]. On launching this inquiry, the Committee's Chair observed that sandboxes have received little evaluation of how they work in practice [6].

19. Timing matters for agentic AI. The FCA's 2026 Payments Regulatory Priorities report states that the FCA will work with HM Treasury to modernise payments regulation, including considering whether change or development of regulation is needed to support agentic AI payments [7, p. 7]. The AI Live Testing evaluation report is due in Q1 2027 [1]. Findings from agentic payments tests are most useful to policy if they reach both processes in time.

20. SBJ suggests four ways to improve how regulators learn from sandboxes and share that learning:

    (a) **A common evidence format** across the Regulatory Sandbox, the Digital Sandboxes, the Supercharged Sandbox and AI Live Testing, so that results on authority, consent, failure modes and redress can be compared across tests.

    (b) **A published agentic AI findings report,** anonymised, covering authority failures, near misses, controls that worked and open questions. It could sit within, or alongside, the Q1 2027 evaluation report.

    (c) **Sharing with other regulators.** Agentic AI raises questions for the Information Commissioner's Office, the Competition and Markets Authority, Ofcom and the Prudential Regulation Authority. The Mills Review recommends strengthening system-wide coordination and oversight [3, p. 10].

    (d) **A defined route from finding to policy.** Each agentic test closes with a short note on any rule, guidance or perimeter question it raised, and the FCA publishes how it handled those notes.

21. Published learning matters most to firms outside sandboxes. Most firms that deploy agents will never take part in a sandbox. Published findings let them adopt controls that have already been tested.

Conclusion

22. Agentic AI is already present in FCA testing programmes. Treating authority as a tested property, with records that continue after the test, would let sandboxes inform the payments review and wider AI policy with evidence the whole market can use.

23. SBJ would be glad to provide further information or to give oral evidence if that would assist the Committee.

References

[1] Financial Conduct Authority, "FCA announces second cohort for AI Live Testing", press release, 21 April 2026. https://www.fca.org.uk/news/press-releases/fca-announces-second-cohort-ai-live-testing

[2] Financial Conduct Authority, "Anthropic to support FCA's Supercharged Sandbox", press release, 22 July 2026 (including notes to editors). https://www.fca.org.uk/news/press-releases/anthropic-supercharged-sandbox

[3] Financial Conduct Authority, AI and the future of retail financial services (The Mills Review), July 2026. https://www.fca.org.uk/publication/corporate/the-mills-review.pdf

[4] Financial Conduct Authority, "AI and the FCA: our approach". https://www.fca.org.uk/firms/innovation/ai-approach

[5] Financial Conduct Authority, Regulatory sandbox lessons learned report, October 2017. https://www.fca.org.uk/publications/research/regulatory-sandbox-lessons-learned-report

[6] House of Lords Financial Services Regulation Committee, "Financial Services Regulation Committee launches inquiry into regulatory sandboxes", 17 September 2026. https://committees.parliament.uk/work/10006/regulatory-sandboxes-in-financial-services/news/217930/financial-services-regulation-committee-launches-inquiry-into-regulatory-sandboxes/

[7] Financial Conduct Authority, Payments Regulatory Priorities, March 2026. https://www.fca.org.uk/publication/regulatory-priorities/payments-report.pdf

[8] Solutions by Jewel, AI Agent Governance Checklist. https://solutionsbyjewel.com/ai-agent-governance-checklist

[9] Solutions by Jewel, AI Decision Record template. https://solutionsbyjewel.com/ai-decision-record-template

[10] Solutions by Jewel, response to HM Treasury consultation "Modernising Payment Services Regulation", Question 15, September 2026.

17 October 2026