CSAI Foundation | Cloud Security Alliance

FCA on Agentic AI: Senior Managers Own Third-Party Agents

How the UK regulator's November 2026 position reinforces that vendor-supplied agents are a senior-manager accountability question

2026-11-26

AI-assisted Rapid Research

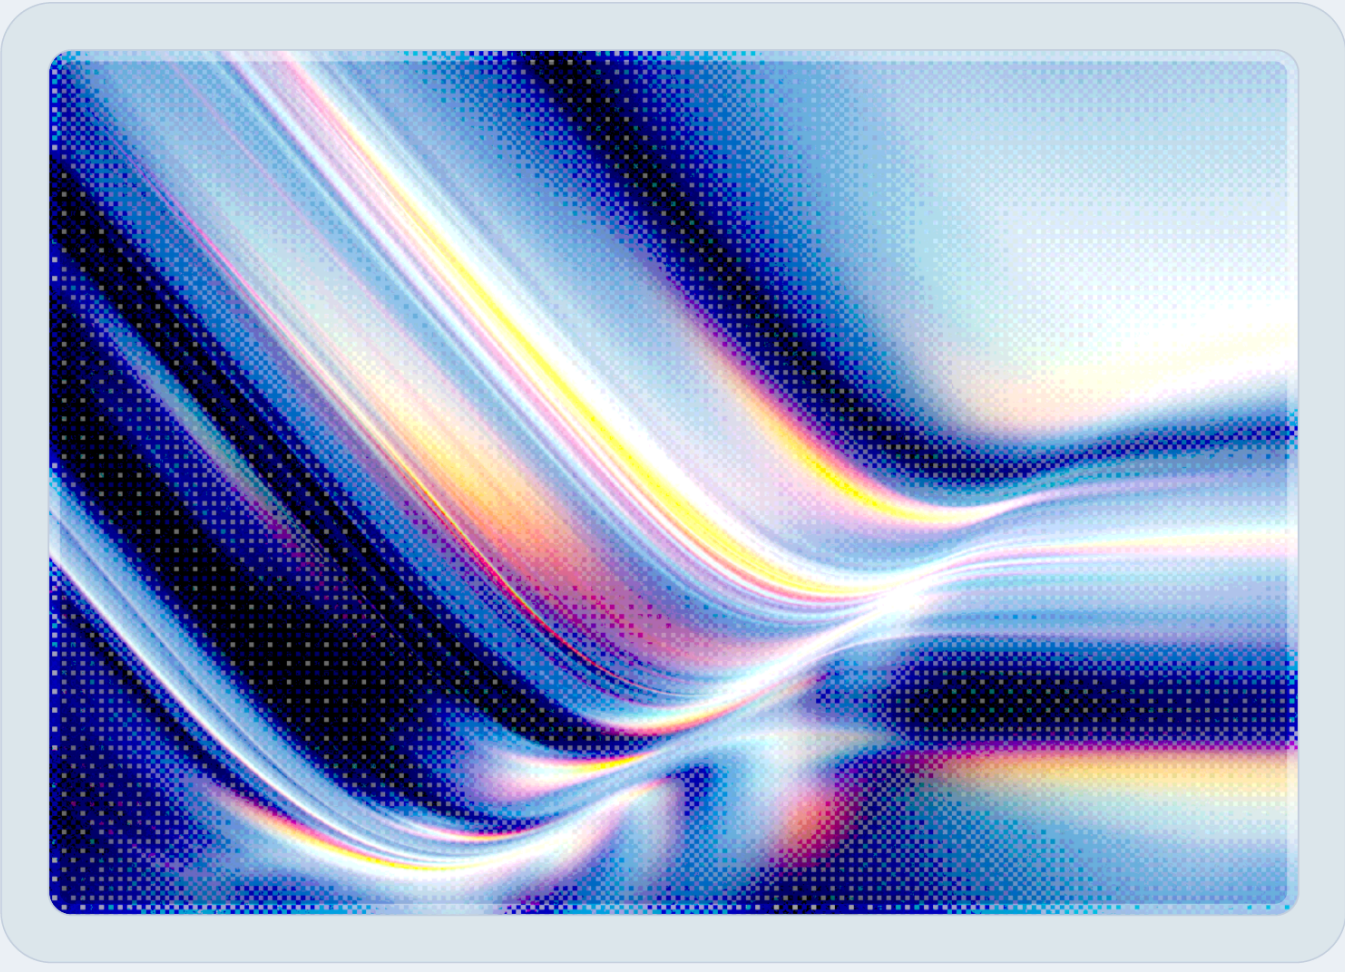

CSAI | CSA cloud security alliance®

© 2026 Cloud Security Alliance. Some rights reserved.

You may download, store, display, view, print, redistribute, and link to this document in its original, unmodified form, provided that attribution to the Cloud Security Alliance is maintained and all trademark and copyright notices remain intact.

This document may not be modified or altered. You may quote portions of the document as permitted by the Fair Use provisions of the United States Copyright Act, provided that attribution is given to the Cloud Security Alliance.

This document may be shared on professional and social media platforms in its original form with attribution.

This document was generated with AI assistance and has not undergone official CSA review and approval processes.

Key Takeaways

On 24 November 2026 the UK Financial Conduct Authority (FCA) published a webpage titled "Our approach to Agentic AI." It sets out expectations for regulated firms, technology providers, and social media platforms that develop or deploy AI agents [1]. The page creates no new rules and sets no compliance deadlines. Its central statement is that when a regulated firm "develops or deploys agentic AI or relies on a third-party provider," responsibility for meeting regulatory obligations "remains with you and senior managers within your firm" [1]. The FCA also says senior managers are responsible for any AI models deployed in their business area, "whether these are developed in house or by a third party," and must take reasonable steps to understand the risks, oversee them effectively, and challenge deployment [1]. Five points follow for security and risk leaders.

- **The emphasis is new; the rules are not.** The page does not name the Senior Managers and Certification Regime (SM&CR), but it uses the regime's "reasonable steps" language [2][8] and builds on the October good-practice document, which places AI ownership with senior manager function (SMF) holders [3]. Read together, they mean an accountable executive should be able to show they understood how much autonomy an agent had and which controls constrained it "in real time and in stress situations" [1].
- **Buying an agent does not transfer accountability.** This follows the FCA's long-standing outsourcing principle [9]. The difference with agents is that the outsourced component makes decisions, and its behaviour can change when the vendor updates the underlying model [3][4].
- **Agents that arrive from customers are now an explicit concern.** Firms are asked how they will recognise an AI agent, check what the customer consented to, and limit what the agent may do. They are also asked whether to accept payments instructed by an agent or made from an agent-controlled wallet [1]. The UK has no settled, interoperable framework for verifying agent identity [5].
- **Calling yourself a technology provider is no defence.** Technology providers whose agents carry out payment services, financial promotions, credit broking, or advice may need FCA authorisation. The FCA says it "will not hesitate" to use its intervention powers [1].
- **The framework is still being built.** The FCA has promised a "trusted framework for agentic finance" that will address agent identity, delegated authority, accountability, and execution [4]. It has also announced an Agentic Policy Sprint [1]. Neither has a published timetable, so firms must meet today's accountability expectations before those mechanisms exist.

Taken together, these points move agent governance from a technology procurement question to one that accountable executives will need to evidence personally. The sections that follow trace how the FCA arrived at this position, analyse its security implications, and set out the actions firms can take now.

Background

The November webpage is the latest of several UK publications on AI agents in financial services during 2026. The FCA has said that it does not plan "to introduce extra regulations for AI" [6]. Instead it relies on its existing principles-based framework, which includes the Principles for Businesses, the Consumer Duty, the SYSC systems-and-controls rules, operational resilience requirements, and the SM&CR [3][6]. During 2026, however, the FCA's attention moved from generative assistants to systems that act. The FCA-commissioned Mills Review, published on 6 July 2026, made seven recommendations. They included "Monitor the transition to autonomous models and adapt regulatory frameworks" and "Enable the foundations for agentic finance" [7]. Commentary on the review's earlier engagement stage noted that it questioned whether the SM&CR and the Critical Third Parties regime suit future AI risk, including "chains of accountability between firms and model developers" [10].

The FCA's answer came in two steps. On 19 October 2026 it published "AI in financial services: good practice and areas for improvement." The document applies to machine learning, generative AI, and agentic AI. It states that firms "are responsible for AI systems they develop, procure and deploy – including their outputs and outcomes" [3]. On the same day, Jessica Rusu, the FCA's chief data, information and intelligence officer, gave a speech saying that "greater autonomy also changes the risk." She said agentic systems "require continuous, conditional assurance" and that the obligation to report major incidents "must extend to significant AI failures and near-misses" [4]. She also previewed the joint Bank of England and FCA survey, published on 26 October 2026, which found that 86% of respondent firms now use AI [4][19]. In parallel, the Digital Regulation Cooperation Forum (DRCF), whose members are the FCA, CMA, ICO, and Ofcom, issued a joint statement on 30 October 2026. It said that an AI agent's autonomy "does not diminish the responsibility of organisations that deploy it" [11]. A follow-up DRCF letter to the Chancellor and ministers on 16 November urged firms to engage regulators before launching personal agent products [12].

The 24 November page turns these signals into sector-specific expectations. It is organised around three audiences: regulated firms, the customers who interact with those firms through agents, and technology businesses whose agent products may cross into regulated activity [1]. The page coincided with a consumer product milestone. According to Beri's reporting, it went live the same day Meta's Muse personal agent became available to UK shoppers [13]. Muse completes purchases through Stripe's Link

wallet for agents [20], and Rusu had already cited its launch as evidence of "the transition towards agentic finance" [4]. The table below summarises the sequence of publications that led to the November page.

Date (2026)	Publication	Relevance to agent accountability
6 July	Mills Review final report [7]	Recommends monitoring autonomous models and building foundations for agentic finance
19 October	FCA good and poor practice on AI [3]	Firms are responsible for AI they "develop, procure and deploy"; assign senior manager (SMF) ownership from design onwards
19 October	Rusu speech, "Building the future of finance in the age of AI" [4]	Continuous assurance, near-miss reporting, trusted framework for agentic finance
30 October	DRCF joint statement on the agentic economy [11]	Agent autonomy does not reduce deployer responsibility across regulators
16 November	DRCF letter to the Chancellor and ministers [12]	Firms should engage regulators before launching personal agents
24 November	FCA, "Our approach to Agentic AI" [1]	Senior managers accountable for third-party agents; expectations on inbound customer agents; perimeter warning

Security Analysis

Personal accountability now reaches the vendor's agent

Under the SM&CR, a senior manager can face personal enforcement action if the firm breaches a requirement in their area of responsibility and they did not take reasonable steps to prevent it [8]. The November page applies that test to AI models in the senior manager's business area, whether built in house or by a third party [1]. The page itself does not name the SM&CR, but it uses the regime's reasonable-steps language, and the October good-practice document places AI accountability explicitly

under the SM&CR [3]. This changes how the security function should think about vendor agents. A reasonable-steps defence depends on evidence: what the senior manager knew, what information they received, what they challenged, and what controls they approved. If a firm cannot produce an inventory of its deployed agents, a description of each agent's permissions and autonomy, and a record of how those were monitored, it may struggle to show that its accountable executive took reasonable steps. Security telemetry about agents therefore becomes evidence of governance as well as an operational tool.

The October good-practice document shows what supervisors consider adequate. Firms doing well assigned "clear SMF ownership for AI risks" and held senior managers accountable "from the design stage onwards for each AI use case" [3]. In one case study, ultimate accountability sat with the SMF for the business area where a tool operated, while a separate model owner handled the technical build [3]. Among firms developing agentic use cases, the stronger ones built governance layers with escalation routes before deployment. Retrofitting governance after go-live "left gaps in accountability" [3]. The FCA does not expect senior managers to be technical specialists. It does expect them to obtain "appropriate information and expertise where necessary" [3]. This suggests the security team's role is to turn agent behaviour into information a senior manager can act on: how autonomous an agent is, which tools it can call, which data it can reach, how far it has drifted from expected behaviour, and which incidents or near-misses have occurred.

Third-party agents differ from conventional outsourcing

Under the FCA's outsourcing rules in SYSC 8, a firm that outsources critical or important functions remains fully responsible for its regulatory obligations [9]. The November page applies the same principle to agents [1]. What differs is the nature of the outsourced component. A conventional outsourced service behaves according to a contract and a specification. An agent supplied by a third party makes decisions inside a delegated scope, and those decisions depend on a foundation model that the deploying firm typically does not control and that the vendor may update. Rusu's speech states the problem directly: "a system that is safe at deployment may behave differently at scale, under stress or as the models beneath it change" [4]. The good-practice document observes that some firms "relied on vendors' assurances without challenging or testing them," and it flags "Overreliance on hard-to-replace AI providers" and "Single-platform dependence without exit options" as areas for improvement [3]. It describes good practice as tracking vendor model updates and keeping viable alternatives [3][5].

In practice, this means a vendor attestation at procurement is unlikely to satisfy the expectation that senior managers understand controls "in real time" [1]. Firms will need assurance that keeps pace with the vendor's release cycle. That includes notice of model changes, the right to re-test after material updates, access to action-level logs, and the technical ability to suspend an agent without disrupting regulated

services. CMS raised the last point when commenting on the Mills Review, asking whether a firm "can suspend an agent without causing disruption" [10]. The Critical Third Parties regime gives regulators direct oversight of designated providers [14]. Designation does not relieve a firm or its senior managers of their own obligations, however, and the Mills Review's question about whether that regime suits AI supply chains remains open [10].

Inbound agents raise identity and authority questions

The most novel part of the November page concerns agents that customers bring to the firm. Firms are asked to consider "how you establish that you are dealing with an agent, what consents are in place, what an AI agent is allowed to do and how customers can get support if something goes wrong" [1]. They are also asked whether it is appropriate to accept payments instructed by an AI agent or made from an agent-controlled digital wallet, and what liabilities may follow from acting on agent instructions [1]. Each of these is a security question as well as a conduct question. Identifying an agent requires a way to distinguish it from the human customer. Verifying consent requires a durable record of what the customer delegated. Limiting an agent's actions requires authorisation scopes that the firm can enforce at runtime rather than infer afterwards.

These capabilities do not yet exist as a common standard. Freshfields notes that the UK has no settled, interoperable framework for verifying AI agent identities. The Mills Review favoured "trusted agent standards" built into Open Finance [5]. North American banks show the risk of improvised approaches. In late October 2026 several large U.S. and Canadian banks began blocking personal agents that log in with customers' own credentials [15][21]. Their main objection is that such agents cannot be told apart from the customer and inherit all of the customer's privileges [15]. UK firms face the same design choice: block agents, allow approved agents, or identify agents and limit their scope. They face it under a Consumer Duty that still requires them to avoid foreseeable harm and to support customers when something goes wrong [1][13]. A firm that blocks all agents may meet its security goals but create Consumer Duty friction, for example by preventing customers from using tools that serve their interests. A firm that accepts credential-sharing agents without separate attribution may find it cannot reconstruct who authorised a disputed payment.

Calling a product a technology service does not move the perimeter

The page's message to technology providers and social media platforms is that the perimeter "is set by Parliament," not by how a business describes itself [1]. Activities named as potentially requiring authorisation include payment services, financial promotions, authorised information services, credit broking or intermediation, and financial advice [1]. The page also notes that the personalised context some large language models rely on "can amplify this complexity" [1]. A general-purpose agent that drifts

from describing products to recommending them, or from showing balances to initiating transfers, may cross the boundary without a deliberate product decision. For security and product teams at AI providers, this suggests that tool permissions and capability boundaries now carry regulatory as well as safety consequences. An agent's tool manifest may effectively describe its regulatory status: a tool that can initiate payments, for instance, signals possible payment-initiation activity regardless of how the product is marketed.

What remains unresolved

Several gaps should be noted plainly. The trusted framework for agentic finance has been announced but not published, and no consultation or timetable is public [4][13]. The Agentic Policy Sprint has no dates [1]. The November page does not address how existing incident-reporting obligations apply to AI near-misses; that expectation comes from the October speech rather than from the guidance itself [4]. Beri also observes that neither the FCA's agentic AI page nor the Payment Systems Regulator has addressed how authorised push payment (APP) fraud reimbursement applies when the customer's agent, rather than the customer, is deceived [13]. Until these questions are settled, firms are held to an outcomes-based standard without detailed technical specifications. In CSA's assessment, this makes contemporaneous documentation of design decisions especially valuable.

The table below maps each expectation the FCA has stated, together with the direction it has signalled on incident reporting, to the security evidence a senior manager may need to demonstrate reasonable steps.

FCA expectation or signalled direction	Primary source	Security evidence a senior manager may need
Understand systems, autonomy, risks, and controls "in real time and in stress situations"	[1]	Agent inventory with autonomy levels, tool scopes, and data access; runtime monitoring dashboards; stress and failure test results
Senior managers accountable for in-house and third-party AI in their area	[1][3]	Named SMF owner per agent use case; documented approval and challenge records
Avoid relying on untested vendor assurances; keep exit options	[3]	Independent testing after vendor model updates; contract rights to logs and change notice; tested suspension procedure

FCA expectation or signalled direction	Primary source	Security evidence a senior manager may need
Respond to customer agents: identification, consent, permissions, support	[1]	Agent authentication method; consent records; enforceable scope limits; dispute reconstruction capability
Decide whether to accept agent-instructed payments and agent-controlled wallets	[1]	Documented accept, allowlist, or block policy; transaction limits; fraud analytics tuned for agent traffic
Signalled direction (speech, not guidance): extend incident reporting to significant AI failures and near-misses	[4]	Agent-specific incident taxonomy; near-miss capture; escalation thresholds

Recommendations

The FCA has set no deadlines, so the timing below reflects CSA's suggested sequencing rather than a regulatory requirement. The recommendations progress from establishing ownership and visibility, through strengthening vendor and customer-facing controls, to shaping the standards that will eventually define reasonable steps for agentic systems.

Immediate Actions

Firms should first establish who owns each agent and what each agent can do. Every agentic use case, whether built in house, licensed from a vendor, or embedded in a third-party platform, should be mapped to a named SMF holder, with that ownership recorded in the senior manager's statement of responsibilities where appropriate. Security teams should build or extend a central agent inventory that records each agent's autonomy level, the tools and APIs it can invoke, the data it can reach, the foundation model and version it depends on, and the vendor responsible. The October good-practice document cites a central inventory as a feature of firms developing agentic use cases well [3]. Firms should also decide on a provisional policy for inbound customer agents, covering blocking, allowlisting, or identify-and-scope access. That policy should be documented with its Consumer Duty reasoning, because the November page expects firms to have considered the question [1].

Short-Term Mitigations

Over the next one to two quarters, which is CSA's suggested horizon, firms should change vendor arrangements so that assurance continues after procurement. Contracts for agent-capable products should require advance notice of material model changes and provide rights to re-test after updates. They should also give access to action-level logs that let the firm reconstruct what an agent did and on whose authority. Firms should rehearse suspending each material agent, confirm that regulated services can continue without it, and record the results. On the customer side, firms should begin separating agent traffic from human sessions wherever the channel allows. They should prefer scoped, revocable, agent-specific credentials over shared customer passwords. Incident management should gain an agent-specific category that captures near-misses such as scope violations, unexpected tool calls, or blocked high-value instructions. This anticipates the direction the FCA signalled in October, when Rusu said incident reporting "must extend to significant AI failures and near-misses" [4]. Senior managers should receive regular summaries that link these signals to the risks they own.

Strategic Considerations

In the longer term, firms should take part in shaping the mechanisms the FCA has announced, through the Agentic Policy Sprint and the trusted framework for agentic finance [1][4]. Agent identity, delegated authority, and execution standards developed there will probably become the reference point for judging reasonable steps. Firms whose architectures already separate agent identity from human identity, narrow delegated authority at each step, and keep tamper-evident action records are likely to be better placed to adopt whatever standards emerge. AI providers and platforms that offer agents to UK consumers should run a perimeter analysis of each agent's actual capabilities rather than its marketing description, and should use the FCA's pre-application services where the analysis is uncertain [1]. Boards should expect the "AI did it" explanation to carry little weight with the UK's digital and financial regulators. The four DRCF members have made that point jointly across competition, data protection, online safety, and financial services law [11].

CSA Resource Alignment

CSA's 2026 survey of financial institutions, *The State of Cloud and AI for Financial Services 2026* [16], offers empirical context relevant to the FCA's concerns. The survey, commissioned by Anjuna, drew on 340 respondents across several regions, so its results are not UK-specific. Its results suggest that adoption of AI agents is outpacing governance maturity: 62% of respondents reported deploying AI agents, 85% anticipated autonomous AI-driven financial transactions, and third-party and supply-chain

risk ranked among the top cloud security challenges. The survey highlights non-human and agent identity governance as a gap. That is the capability the FCA now expects firms to show when it asks them to understand agent autonomy and to identify inbound agents. Firms can use the survey's findings to benchmark their own maturity against peers when preparing reasonable-steps evidence for senior managers.

CSA's research note *Banks Block Personal AI Agents: Liability, Credentials, and Fragmentation* [15] examines the North American response to the inbound-agent problem the FCA has now raised for UK firms. Its analysis argues that credential-sharing agents cannot be attributed as separate actors and finds evidence that liability is shifting to consumers through account agreements. It is useful for UK firms deciding whether to block, allowlist, or identify-and-scope agents, especially where the Consumer Duty limits how much risk can be pushed onto customers. The companion note *Agents at the Checkout: Loss Allocation in Frontier-Lab Payment Wallets* [17] addresses the agent-controlled wallet scenario that the FCA asks firms to evaluate. It analyses how losses are allocated when a personal agent completes purchases through a frontier lab's wallet product.

CSA's research note *AI Liability Inflection: Enterprise Accountability in the Agentic Era* [22] examines the broader legal trend behind the FCA's position. It argues that accountability for autonomous agents concentrates on the deploying enterprise rather than on the model developer, which is the principle the FCA has now applied to financial services and the DRCF regulators have endorsed collectively [11].

The AI Controls Matrix (AICM) v1.1 [18] offers one control structure for turning the FCA's expectations into auditable evidence. Its governance, identity and access management, supply chain, and logging domains correspond to the evidence a senior manager would need: who owns an agent, what it can access, how vendor dependencies are managed, and how its actions are recorded. AICM's shared-responsibility model across model providers, application providers, and AI customers sets out which controls a deploying firm keeps when it relies on a third-party agent. That is the responsibility the FCA says remains with the firm and its senior managers [1].

References

[1] Financial Conduct Authority. "Our approach to Agentic AI." FCA, November 24, 2026.

[2] Financial Conduct Authority. "Senior Managers and Certification Regime." FCA, accessed November 26, 2026.

[3] Financial Conduct Authority. "AI in financial services: good practice and areas for improvement." FCA, October 19, 2026.

[4] Jessica Rusu, Financial Conduct Authority. "Building the future of finance in the age of AI." FCA speech, October 19, 2026.

[5] Freshfields. "AI in financial services: FCA good and poor practice and the road to agentic finance." Freshfields, October 22, 2026.

[6] Financial Conduct Authority. "AI and the FCA: our approach." FCA, last updated November 18, 2026.

[7] Financial Conduct Authority. "FCA publishes landmark review into impact of AI on retail financial services." FCA press release, July 6, 2026.

[8] Stephenson Harwood. "FCA Sets Expectations for AI: What Do 'Reasonable Steps' Look Like for Senior Managers?" Stephenson Harwood, October 26, 2026.

[9] Financial Conduct Authority. "SYSC 8.1 General outsourcing requirements." FCA Handbook, accessed November 26, 2026.

[10] CMS. "FCA Mills Review: early indications for AI development in financial services." CMS, February 16, 2026.

[11] Burges Salmon. "AI agents and responsibility – a message from regulators." Burges Salmon, October 30, 2026. (Summarises the DRCF joint statement "Protecting people and businesses in an agentic economy," October 30, 2026.)

[12] Norton Rose Fulbright. "DRCF letter on personal AI agents." Regulation Tomorrow, November 19, 2026.

[13] Rajesh Beri. "FCA Asks UK Firms to Set AI Agent Permissions and Payment Rules." The D*AI*LY Brief, November 25, 2026.

[14] Financial Conduct Authority. "PS24/16: Operational resilience: Critical third parties to the UK financial sector." FCA, November 2024.

[15] Cloud Security Alliance. "Banks Block Personal AI Agents: Liability, Credentials, and Fragmentation." CSA AI Safety Initiative, November 2026.

[16] Cloud Security Alliance. "The State of Cloud and AI for Financial Services 2026." CSA, June 2026.

[17] Cloud Security Alliance. "Agents at the Checkout: Loss Allocation in Frontier-Lab Payment Wallets." CSA AI Safety Initiative, November 2026.

[18] Cloud Security Alliance. "AI Controls Matrix (AICM) v1.1." CSA, 2026.

[19] Bank of England and Financial Conduct Authority. "Artificial intelligence in UK financial services - 2026." Bank of England, October 26, 2026.

[20] Stripe. "Stripe helps Muse, Meta's new personal AI agent, shop across the internet with Link." Stripe Newsroom, 2026.

[21] Malwarebytes. "Banks are blocking AI agents from customer accounts." Malwarebytes Labs, October 2026.

[22] Cloud Security Alliance. "AI Liability Inflection: Enterprise Accountability in the Agentic Era." CSA AI Safety Initiative, 2026.