

RE • WORK

15 SEPTEMBER 2026

# AI in Finance Summit London

The must-attend summit for solving the engineering and governance hurdles of actual AI deployment.

AI adoption in financial services is no longer the challenge; scaling it safely, reliably, and with measurable impact is.

As **AI in Finance** returns to London for its 11th edition, it brings together the leaders and practitioners navigating this exact shift: **moving beyond pilots to production-ready systems that are tested, governed, and resilient by design.**

It is a practitioner-led environment focused on how AI is actually built, validated, and scaled within regulated institutions. Expect real-world use cases, hands-on sessions, and candid discussions—and leave with practical frameworks you can apply immediately.

## WHAT'S NEW IN 2026?

- **Execution Over Strategy:** A shift from high-level theory to a "working session" format—focused on the build, deployment, and control of AI in finance.
- **Technical Deep-Dives:** New interactive workshops on AI evaluation, stress-testing, and agent design with immediately applicable frameworks.
- **The "Failure" Files:** Candid deployment stories. Hear what broke between pilot and production and how top-tier firms bypassed those hurdles.
- **End-to-End AI Security:** Elevated focus on the threat landscape, addressing everything from prompt injection to deepfake defense.
- **Infrastructure and Resilience:** Dedicated sessions on cloud concentration, third-party dependencies, and maintaining compliance under DORA.
- **Practitioner-Led Networking:** A curated, senior-only audience of technical and innovation leaders designed for high-level peer exchange.

## Key themes:

- **From Pilot to Production:** Turning AI experiments into scalable, production-ready systems with measurable impact.
- **Optimal Architecture:** Navigating RAG, fine-tuning, and infrastructure to balance cost, speed, and control.
- **Testing, Evaluation and Observability:** Validating AI systems, reducing risk, and ensuring reliability in live environments.
- **Agentic AI:** Designing autonomous and semi-autonomous workflows with clear guardrails and human oversight.
- **Model Risk and Governance:** Building explainable, traceable systems that meet evolving regulatory standards.
- **Security and Financial Crime:** Defending against AI-driven threats and advancing fraud detection.
- **Resilience and Third-Party Risk:** Managing vendor dependencies, cloud risk, and resilience under DORA and evolving regulations.
- **Markets and Trading:** Reshaping alpha generation and risk modelling through advanced AI innovation.

Questions about the programme? Get in touch with Conference Producer **Emilia Malaj** at [emilia.malaj@coriniumgroup.com](mailto:emilia.malaj@coriniumgroup.com)

**Confirmed Speakers:**

Roxanne Howdle-Rowe, *Managing Director, Group Head of Data & Analytics Office* — **BRITISH BUSINESS BANK**  
 Annapoorani Lakshmi Narayanan, *Product Manager Lead, AI Research* — **J.P. MORGAN CHASE**  
 Parag Kumar, *Executive Director, Data Analytics Lead, Compliance, Conduct and Operational Risk* — **J.P. MORGAN CHASE**  
 Jai Ferguson, *AI Regional Lead for Europe* — **HSBC**  
 Kaushik Chaubal, *Senior Director, Investor Solutions* — **BLACKROCK**  
 Hemant Patkar, *Cybersecurity Lead Designer/Architect* — **VIRGIN MONEY**  
 Parinita Kothari, *Engineering Lead* — **LLOYDS BANKING GROUP**  
 Charles Phiri, *Executive Director, SME AI/ML Innovation* — **J.P. MORGAN CHASE**  
 Eshaan Salwan, *Cybersecurity Audit Analyst* — **LLOYDS BANKING GROUP**  
 Soung Low, *Model Risk Data Scientist* — **NATWEST GROUP**  
 Amit Kumar, *Associate Director, eFx Quant* — **HSBC**  
 Rajasree R, *Product Lead - Agentic Observability* — **LLOYDS BANKING GROUP**

<b>AI in Finance London 2026   Program</b> 15 <sup>th</sup> of September	
8:00-8:45	<i>Registration and Coffee in the Exhibition Area</i>
8:45-9:00	<b>Chairperson's Opening Remarks and Icebreaker</b>
9:00-9:40	<b>OPENING PANEL: AI in Production — Real Value vs. The Reality Gap</b> <ul style="list-style-type: none"> <li>Which use cases are delivering true scale (fraud, ops, servicing) and why are others hitting the "deployment wall"?</li> <li>Honest lessons from the "PoC graveyard"—where are firms rolling back investments, and what should practitioners stop doing immediately?</li> <li>How digital-native banks are outstripping incumbents, and what legacy institutions can realistically adopt to stay competitive.</li> </ul>
9:40-10:00	<b>KEYNOTE: Cracking The Last Mile Problem with Frameworks for Industrial-Grade AI</b> <ul style="list-style-type: none"> <li>Why great pilots fail in production and how to architect for success?</li> <li>Insights on overcoming the governance and testing hurdles that cause stalling.</li> <li>Methods for ensuring your AI remains observable, safe, and resilient in live financial environments.</li> </ul>
10:00-10:30	<b>CASE STUDY: The Production Blueprint from Inside a Real Deployment Journey</b> <ul style="list-style-type: none"> <li>A step-by-step look at aligning architecture and governance to a move-high stakes project out of the lab.</li> <li>Identifying where deployments fail between PoC and scale, and the proactive fixes used to bridge those gaps.</li> <li>Practical strategies for gaining 'yes' from risk, compliance, and cross-functional stakeholders without slowing down.</li> <li>How to define and communicate value in a way that secures multi-year funding and firm-wide adoption.</li> </ul>
10:30 – 11:00	<i>Mid-Morning Coffee Break &amp; Networking in Exhibition Area</i>
11:00-11:30	<b>TECHNICAL KEYNOTE: From Models to Systems — The New AI Engineering Stack</b> <ul style="list-style-type: none"> <li>Moving from standalone LLMs to integrated production systems—solving for latency, consistency, and state management.</li> <li>What does “gold standard” look like for AI engineering in financial services today?</li> <li>Exploring the evolution of tooling and pipelines—how automated evaluation and real-time monitoring are enabling hands-off deployment.</li> </ul>

11:30-12:10	<b>PANEL: From Paper Policies to Automated Engineering Guardrails — Governance and Model Risk Musts</b> <ul style="list-style-type: none"> <li>• How to move from abstract "responsible AI" principles to concrete, defensible engineering practices?</li> <li>• Strategies for identifying and securing "off-book" LLM usage across the enterprise.</li> <li>• Where does data privacy fails in complex, agentic financial ecosystems?</li> <li>• Defining clear accountability across Risk, Tech, and Business to avoid "governance by committee."</li> </ul>	
	<b>USE CASE SHOWCASE</b>	
12:10-12:20	<b>Showcase 1: Next-Gen Fraud and AML Detection for Precision</b>	
12:20-12:30	<b>Showcase 2: Elevating Customer Operations and CX Through Augmented Agents</b>	
12:30-12:40	<b>Showcase 3: The Intelligent Enterprise — RAG, Knowledge, and Internal Workflows</b>	
12:40-13:40	<i>Lunch and Networking in the Exhibition Area</i>	
	<b>MAIN CONFERENCE</b> <b>Workshops</b>	<b>TRACK C</b> <b>Roundtables</b>
13:40-14:20	<b>WORKSHOP A: The Pre-Flight Checklist for Stress-Testing and Validating Financial AI</b> <ul style="list-style-type: none"> <li>• Moving beyond "accuracy" to define rigorous KPIs for hallucination, bias, and financial reasoning.</li> <li>• Practical frameworks for uncovering edge cases and security vulnerabilities before they reach production.</li> <li>• Identifying intervention points where human oversight must be implemented into the automated workflow.</li> <li>• How to monitor and improve performance to prevent model drift post-deployment?</li> </ul>	<p><i>Two simultaneous roundtables running in parallel. Attendees remain seated at their chosen table throughout. Speakers rotate between tables at 40-minute mark, ensuring both groups hear from every facilitator.</i></p> <p><i>Each table will have a printed prompt card with 3 discussion questions to guide conversation and encourage audience participation.</i></p> <p style="text-align: center;"><b><u>Human + AI Operating Model and Adoption at Scale</u></b></p> <p><b>ROUNDTABLE 1: Embedding AI into Real Workflows — From Tooling to Daily Usage</b>  Focus: Operational adoption and workflow integration</p> <ul style="list-style-type: none"> <li>• Where are AI tools actually being used daily vs. abandoned after rollout—and why?</li> <li>• How do you redesign workflows (not just tools) to ensure AI becomes part of decision-making?</li> <li>• What incentives, training, or friction points determine whether teams adopt or reject AI?</li> </ul> <p><b>ROUNDTABLE 2: Owning the Outcome — Accountability, Trust, and Human Oversight in AI Decisions</b>  Focus: Decision ownership and human-AI collaboration</p> <ul style="list-style-type: none"> <li>• Who is accountable when AI influences or makes decisions—and how is that defined in practice?</li> <li>• Where should the line sit between human judgment and automated decision-making?</li> <li>• How do you build trust with business users who don't fully understand the models?</li> </ul>
14:20-15:00	<b>WORKSHOP B: Engineering Accountability — Designing Controlled Workflows for Agentic AI</b> <ul style="list-style-type: none"> <li>• Identifying the "Goldilocks" use cases for agents: tasks with high complexity but clear, programmable constraints.</li> <li>• How do you design workflows with hard-coded "Circuit Breakers" and human-intervention triggers?</li> <li>• Frameworks for ensuring consistency across agent systems to meet internal audit and reliability standards.</li> <li>• How to mitigate the risks of data leakage and "permission creep" as agents navigate internal banking silos?</li> </ul> <p>Senior Representative,  <b>IBM</b></p>	
15:00-15:30	<i>Afternoon Coffee Break and Networking in Exhibition Area</i>	

15:30-16:00	<p><b>FIRESIDE CHAT: Algorithmic Alpha vs. Systemic Stability — AI’s Impact on Capital Markets</b></p> <ul style="list-style-type: none"> <li>• How is AI fundamentally altering alpha generation, execution logic, and portfolio construction in 2026?</li> <li>• Exploring the risks of "Model Herding"—could synchronized AI decision-making trigger new forms of market volatility and flash-instability?</li> <li>• Where should firms draw the line between machine execution and human judgment?</li> <li>• How are regulators approaching AI in capital markets, and what should firms prepare for next?</li> </ul>
16:00-16:30	<p><b>KEYNOTE: The AI Security Beyond Phishing and Prompt Injection</b></p> <ul style="list-style-type: none"> <li>• Defending against "excessive agency" and unauthorised actions as autonomous systems become primary targets for cybercrime.</li> <li>• Moving past system prompts to secure the non-human identities and "digital insiders" managing high-value transactions.</li> <li>• Strategies for securing the AI supply chain—from data poisoning and model tampering to deepfake-driven social engineering.</li> <li>• Building "Defensible AI" that maintains operational integrity under the scrutiny of the EU AI Act and DORA.</li> </ul>
16:30-17:00	<p><b>PANEL: The New AI Partnership — Navigating Resilience, Sovereignty, and Vendor Lock-In</b></p> <ul style="list-style-type: none"> <li>• Navigating the systemic risk of high-reliance on a few dominant third-party AI and Cloud providers.</li> <li>• How can firms design for resilience under DORA, including failure scenarios and service outages?</li> <li>• Meeting the 2026 FCA/ESMA mandate for algorithmic accountability—who is liable when an external model fails?</li> <li>• How to maintain operational control when your core intelligence sits in a third-party environment.</li> </ul>
17:00-17:10	<p><b>Chairperson’s Closing Remarks</b></p>
17:10-18:00	<p><i>Networking Drinks Reception and Prize Draw</i></p>
<p><b>END OF SUMMIT</b></p>	