

CO-HOSTED BY

A wide-angle photograph of the Singapore skyline at dusk. The Marina Bay Sands hotel is the central focus, with its three towers and the SkyPark connecting them. To the left, the Singapore Flyer is visible. The foreground shows the water of Marina Bay and some lower-level buildings. The sky is filled with soft, grey clouds.

AI in Finance: From Use Cases to Agentic AI

Practical adoption today. Institutional
readiness for the next frontier.

CONTEXT AND SCOPE

About the roundtable

Artificial intelligence was moving rapidly from experimentation to practical deployment across financial services. Banks, asset managers and other financial institutions were already using AI to improve productivity, strengthen risk management, enhance customer experience, support compliance, accelerate research and redesign internal workflows.

At the same time, a new question was emerging. The next phase of AI in finance may not be defined only by better models or more powerful copilots, but by the rise of agentic AI: systems that could plan, take multi-step actions, interact with tools, coordinate workflows and escalate decisions to humans when needed.

For financial institutions, this created both opportunity and complexity. AI use cases were no longer only about automation or analytics. They increasingly raised deeper questions about workflow design, accountability, governance, human oversight, operational resilience and the future shape of financial work.

This roundtable brought together a small group of senior leaders to discuss where AI was creating real value at the time, and how agentic AI could reshape financial services in the coming years.

Because the roundtable took place in Singapore and was convened with the British Chamber of Commerce, Standard Chartered, and CFTE, the discussion was designed to bring together both Singapore-based insights and UK perspectives. The aim was to create a discussion that was practical, globally relevant and grounded in real financial services experience.

Core framing: use cases at the time showed where AI was creating value; agentic AI raised the next set of questions about workflow, control and institutional readiness.



HOSTED BY British Chamber of Commerce, Standard Chartered, and CFTE	FORMAT Closed-door leadership roundtable	VENUE Standard Chartered, Singapore
OPENING REMARKS Deputy High Commissioner, British High Commission Singapore Lucy Hughes	CO-CHAIRS Tram Anh Nguyen David Sharratt	PARTICIPANTS Senior leaders from banking and asset management

CONTEXT AND SCOPE

Purpose of the Roundtable

The purpose of the session was to create a practical senior-level discussion on AI adoption in finance, grounded in real use cases and forward-looking enough to explore the implications of agentic AI.

Identify where AI is creating real value in financial services today

Participants shared practical examples of AI use cases across banking, asset management and related areas of financial services, including what was working, what remained difficult and what separated serious deployment from experimentation.

Explore how agentic AI could change financial workflows

The discussion examined how AI systems that could take action, coordinate tasks and interact with tools might affect operations, compliance, risk, research, advisory, customer journeys, internal productivity and technology functions.

Bring together UK and Singapore perspectives

The roundtable drew on Singapore's role as a major financial and technology hub, while also comparing experiences from international institutions and markets. The goal was to understand which AI developments were local, which were global and what financial leaders could learn across jurisdictions.

Confidentiality and reporting scope

The discussion was conducted under the Chatham House Rule. Findings were anonymised and aggregated and should not be read as the position of any participant, employer, regulator or co-organiser. The proposed third discussion area was not held as a standalone segment because of time constraints; cross-border and geopolitical issues that surfaced organically are addressed in the conclusion.

Co-chairs and keynote



CO-CHAIR

Tram Anh Nguyen

Co-Chair, Finance & Fintech Committee, British Chamber of Commerce Singapore
Co-Founder, CFTE



CO-CHAIR

David Sharratt

Co-Chair, Finance & Fintech Committee, British Chamber of Commerce Singapore
Managing Director, Corporate & Institutional Banking, Standard Chartered



OPENING KEYNOTE

Lucy Hughes

Deputy High Commissioner, British High Commission Singapore

Set the UK–Singapore context for the session and emphasised the importance of scaling AI responsibly through cooperation between industry, government and regulators.

At the table



Lucy Haydon

Deputy Executive Director

British Chamber of
Commerce Singapore



Huy Nguyen Trieu

Co-Founder

CFTE



Brittney Buckley

Head of Business
Transformation, Corporate &
Investment Banking

Standard Chartered



Shayan Hazir

Managing Director, Chief
Digital Officer, Asia, HK &
MENAT

HSBC



Ian Macdonald

CEO APAC

Aberdeen PLC



Tina Berg-Nielsen

Vice President, Digital Capital
Markets

NatWest



Dennis Quah

Managing Director

BlackRock



Steve West

Head of Asia

Endava



Balaji Natarajan

Regional Head of Product &
Channels, APAC & Middle
East, GTB

Barclays

OBSERVERS



Megan Otway

FCA Financial Services Attaché

British High Commission
Singapore



Eugene Goh

Deputy Director, Financial
Infrastructure and AI
Monetary Authority of Singapore



Vivian Fang

Associate, Financial
Infrastructure and AI
Monetary Authority of
Singapore



Amy Ross

Head of Financial, Professional
and Business Services, APAC
British High Commission
Singapore



**Catriona Lindsey-
Renton**

Head of Financial Policy
British High Commission
Singapore



Monica Chen

Executive Director, Corporate &
Investment Banking
Standard Chartered



Paridhi Agarwal

Graduate Analyst
CFTE

A COMMON LANGUAGE

The FCA AI in Finance Observatory


To make the discussion as concrete and useful as possible, participants were invited to prepare in advance.

Before the session, participants received a AI use case taxonomy currently used by the FCA AI in Finance Observatory, developed with CFTE. This taxonomy was shared as a reference framework to help participants structure their examples and compare use cases across different parts of financial services.

Participants were invited to prepare examples, observations or questions related to AI in finance. Where possible, these could be mapped to the taxonomy. The aim was not to prepare a formal presentation. The aim is to ensure that each participant arrived with concrete examples and perspectives that could enrich the discussion.

How to use it

Step	Question	Output
1	Describe the use case without starting from the technology.	A business-led use-case description, such as AI-supported document review for onboarding.
2	Identify where it operates.	Sector, function category, sub-function, process step or workflow type
3	Identify the primary value driver.	Cost efficiency, revenue enablement, risk mitigation, productivity, customer experience, or capability and innovation.
4	Separate use case from implementation.	A general use-case pattern, with implementation details recorded only where available or appropriate.
5	Record evidence carefully.	Source type, maturity signal, public or confidential status, and any limitations.



KEY FINDINGS

Executive summary

The room broadly agreed that financial services had moved beyond the earliest experimentation phase, but not into uniform, controlled scale. Meaningful results were emerging while adoption remained uneven across institutions, business lines and markets. Several participants placed their own sectors closer to experimentation, particularly where customer-facing deployment, fragmented regional operations or limited investment capacity raised the threshold for trust and control.

01 Value was emerging in bounded workflows

The clearest examples augmented defined tasks: preparing client context, supporting sales, analysing onboarding requirements, processing claims and communications, and synthesising research. The strongest cases released capacity without removing human judgement.

02 Workflow redesign matters more than the model

AI cannot simply be added to processes designed around manual checks, legacy systems and changing rules. Institutions need to map the end-to-end workflow, determine where probabilistic outputs are useful and preserve deterministic decision boundaries.

03 Human oversight remains the operating default

Participants repeatedly distinguished colleague-facing assistance from autonomous customer-facing action. Human review, verification and escalation were seen as necessary for high-impact decisions, regulated processes and early production deployments.

04 Economics is becoming a strategic constraint

Token consumption, licence costs, compute, model evaluation, data access and integration costs may change which use cases remain viable at scale. Institutions need total-cost and outcome measures, not only demonstrations of technical feasibility.

05 Agentic AI expands the governance perimeter

When systems can plan and act, questions of ownership, permissions, auditability, incentive alignment, operational resilience and agent-to-agent interaction become as important as model performance.

06 Education is a primary adoption bottleneck

A lack of practical understanding among professionals – what AI can do, where it fails, how to verify it and how to redesign work around it – emerged as one of the strongest findings. Capability building must reach leaders, managers, domain experts and junior talent.

OVERALL SYNTHESIS

AI adoption in finance is becoming an institutional transformation challenge rather than a model-selection exercise. Sustainable value depends on use-case discipline, workflow and architecture design, clear controls, measurable economics, and a workforce able to exercise judgement alongside AI.

Implications for decision-makers

Financial institutions should manage AI as a portfolio of business changes rather than a collection of tools. Regulators and policymakers can support this transition through shared evidence, testing environments, interoperable expectations and cross-border learning. Education providers and employers need to build role-specific pathways that combine AI fluency with domain knowledge, systems thinking and responsible decision-making.

Part 1 FROM EXPERIMENTATION TO MEASURABLE VALUE

AI use cases in finance today

Where adoption currently stands

The opening pulse check showed a clear centre of gravity: AI was beginning to produce meaningful results, but adoption remained uneven. The divergence was not simply between early and advanced firms. It also appeared within institutions — between office productivity tools and embedded workflows, between internal and client-facing applications, and between use cases with clear controls and those that remained difficult to trust.

Scale and organisational context mattered. Larger institutions may have dedicated AI offices and stronger infrastructure, while asset managers, smaller firms and multi-market businesses can face proportionally higher costs and more fragmented regulatory or data environments. The discussion also distinguished access from adoption: making a copilot available does not mean it has been integrated into work or is generating measurable value.

Use-case snapshot from the discussion

USE-CASE PATTERN	PRIMARY FUNCTION	MAIN VALUE	IMPLEMENTATION SIGNAL / CONTROL
AI-supported onboarding and KYC document preparation	Front / middle / back office	Customer experience, productivity, risk mitigation	Early-stage results; evolving rules, cross-border data and human validation remain central.
AI-enabled sales context and content preparation	Front office	Productivity and revenue enablement	Colleague-facing workflow moving into production after end-to-end process mapping; human verification retained.
Insurance claims processing and B2B communications	Operations and servicing	Speed, capacity and customer experience	Reported efficiency gains; successful scaling depends on data quality and process integration.
Internal GPTs and enterprise knowledge support	Corporate functions / technology	Productivity and capability	Useful for synthesis and access to knowledge; privacy, permissions and cost require active management.
Investment research and systematic insight	Front / middle office	Decision support and investment outcomes	AI adds consistency and speed; human insight and accountability remain complementary.
Consulting, research and document-intensive work	Corporate / advisory	Productivity and cost efficiency	Material capacity gains were reported, but quality assurance and economic measurement are required.

Theme 1 AI as a colleague-facing productivity layer

The most mature examples placed AI around employees rather than directly in front of customers. A sales workflow brought together unstructured information from messages, market sources and relevant announcements to create a contextual view for relationship teams. The value was not simply faster content generation: it released capacity for higher-quality client engagement.

Internal GPTs, document synthesis, B2B communications and consulting-style research reflected the same pattern. AI was strongest when it aggregated context, drafted or summarised, and enabled a professional to apply domain judgement — narrow enough to control, but broad enough to remove meaningful friction from daily work.

Part 1 IMPLEMENTATION LESSONS**Theme 2 Customer journeys and operational processing**

Digital onboarding illustrated both the promise and the difficulty of AI-enabled operations. Existing KYC and client-lifecycle processes were often designed around human validation, repeated document requests and rigid workflow stages. AI could analyse a target customer profile, anticipate which documents were likely to be needed and sequence requests more clearly, reducing back-office effort and improving the customer experience without replacing the underlying control process. The same logic applied to claims processing and other operational communications. However, requirements change, data may be distributed across jurisdictions and the workflow cannot be treated as static: successful implementation depends on continuous rule maintenance, data permissions, exception handling and clear responsibility for final decisions.

Theme 3 Decision support and investment insight

Investment-related examples showed AI being used to extract and update insight more quickly than traditional research processes alone. Participants did not frame the future as human judgement or AI, but as a changing combination of both. Systematic methods can bring consistency, absorb new information and reduce friction, while experienced professionals retain responsibility for objectives, interpretation, suitability and judgement under uncertainty. This complementary model may alter where competitive advantage sits: as information gathering becomes faster and more widely available, differentiation shifts towards the quality of questions, proprietary data, portfolio construction, client understanding and the ability to interpret AI-supported evidence responsibly.

Theme 4 Enterprise enablement, measurement and economics

Broad access to copilots and internal AI tools was creating valuable experimentation, but participants cautioned against equating usage with value. Institutions still need to measure outcomes, redesign work, evaluate multiple models, integrate legacy systems and decide who owns the operating environment. Behavioural adoption can take as long as technical deployment. Cost emerged as a material concern: token consumption, model and licence fees, data access, testing, integration and monitoring can change the business case as usage expands. The risk is that pilots are judged on impressive outputs while the full cost of production, control and recurring use remains unclear. Participants therefore called for outcome-based measurement — time released, error reduction, cycle-time improvement, revenue protection, customer outcomes and total cost of ownership.

What separated promising pilots from scalable value

Use-case fit	Workflow readiness	Control and trust	People and behaviour	Economics and evidence
Probabilistic systems suit context, language and judgement support — not unguarded substitution for deterministic rules or approvals.	Value required process mapping, data access, permissions, integration and redesign rather than adding AI to an unchanged workflow.	Human review, testing, monitoring, explainability and escalation were necessary to move from demonstration to operational use.	Employees needed practical education, protected experimentation and leadership support to adopt new ways of working.	Clear baselines and measures of value, including recurring model, token, licence and assurance costs.

PART 1 SUMMARY

AI is creating real value where it is embedded in a bounded workflow, close to a clear user need and supported by human judgement. The central challenge is no longer finding possible use cases; it is selecting the right ones, redesigning the process around them, and proving that value and controls survive at scale.

Part 2 FROM ASSISTANCE TO ACTION

Agentic AI in financial services

Agentic AI was framed as a shift from tools that respond to prompts towards systems that can plan, take multi-step actions, interact with tools, coordinate workflows and escalate decisions. The discussion was deliberately practical: where could this capability create value, and what new risks arise when AI begins to act rather than only advise?

Applications and use cases discussed

Bounded colleague-facing agents

Agents can assemble client context, draft responses, retrieve information and coordinate parts of a sales or service workflow — a credible near-term pathway because employees remain responsible for verification and action.

Client lifecycle and KYC orchestration

Agents could gather documents, query data sources, prepare case information and route exceptions across onboarding. The opportunity is substantial, but deterministic validation and human escalation are essential at decision points.

Research, advisory and investment engagement

Agents may monitor information, prepare analysis, support portfolio or product comparisons and personalise engagement, raising questions about suitability, transparency and the future role of relationship managers.

Enterprise workflow coordination

As organisations deploy multiple models and tools, agents may orchestrate tasks across systems and eventually interact with other agents — a new operating layer requiring identity, permissions, ownership and monitoring.

Opportunities identified

- Shorter end-to-end cycle times by coordinating multiple steps rather than accelerating only one task.
- Greater capacity for teams to focus on judgement, client relationships, exceptions and higher-value work.
- More proactive service through continuous monitoring, contextual retrieval and personalised preparation.
- Reusable operating capabilities that connect unstructured information with existing systems, rules and workflow tools.
- New service and market structures as software agents search, compare, negotiate or interact on behalf of institutions and clients.

Strategic implications for the model

Participants raised the possibility that agentic systems could change intermediation. A client-side agent may compare products, prices and advice more efficiently; an institutional agent may respond to requests or manage workflows at machine speed. This could compress information advantages, change the role of relationship managers and shift some firms from price-setting towards responding to agent-driven comparison and execution. These possibilities were not presented as inevitable or immediate. They illustrate why agentic AI should be treated as an operating-model question: who the agent represents, which objective it optimises, what data and tools it may access, and where responsibility remains with a human or legal entity.

NEAR-TERM OPPORTUNITY

The strongest practical case was bounded agency: systems acting inside clearly defined internal workflows, with restricted permissions, deterministic control points and accountable human owners. Fully autonomous customer-facing deployment remained a much higher bar.

Part 2 RISKS, LIMITATIONS AND CONTROL

The central architectural challenge

A recurring concern was the mismatch between probabilistic models and deterministic financial workflows. Large language models are effective at interpreting unstructured information and producing plausible language, but financial calculations, policy thresholds, reconciliations, eligibility rules and regulated approvals often require reproducible outcomes and near-zero tolerance for error.

The problem becomes more difficult in multi-agent chains. An error introduced by one agent can be reformulated, amplified or acted upon by another. The governance challenge is end-to-end reliability, not the average accuracy of an individual model call.

Key risks and practical control directions

RISK	WHY IT MATTERS	CONTROL DIRECTION
Compounding errors	Small failure rates can accumulate across multi-step or agent-to-agent workflows.	Limit chain length; validate structured outputs; test end-to-end; route uncertainty and exceptions to humans.
Auditability and lineage	Dynamic context and tool calls can make decisions difficult to reproduce or explain.	Log prompts, data, tool calls, versions, approvals and state transitions; preserve deterministic records.
Accountability and ownership	It may be unclear who owns an agent, its actions and its failures across business and technology teams.	Assign named business owners, decision rights, risk acceptance and incident responsibilities.
Incentive misalignment	An agent may optimise the stated objective in a way that conflicts with customer outcomes, conduct or institutional intent.	Define objective hierarchies, constraints, customer duties and stop conditions; test adversarial scenarios.
Consumer and data risk	Autonomous actions can amplify privacy, bias, suitability, disclosure and data-residency concerns.	Use least-privilege access, consent and data controls, outcome monitoring and human review for material decisions.
Operational and cost loops	Retries, external API changes or recursive tasks can create latency, rate-limit failures and uncontrolled spend.	Set budgets, timeouts, transaction limits, circuit breakers, fallback modes and real-time monitoring.

A practical design principle

The discussion pointed towards hybrid architectures: use probabilistic models for semantic translation, unstructured context and drafting; pass structured outputs through deterministic rules, calculations and schema checks; and retain human escalation for exceptions or consequential decisions. In simple terms, the model may propose, but controlled systems and accountable people must determine what is executed.

Key observations

- Agentic AI is not only a model-risk issue. It is an operational-resilience, identity, permissions, workflow and governance issue.
- Controls need to be designed at the process and system level, including interactions between agents, data sources and external tools.
- Human-in-the-loop should be defined precisely: who reviews, at which decision boundary, with what information and authority?
- Institutions should test failure modes, incentives and economic behaviour, not only benchmark output quality.

PART 2 SUMMARY

Agentic AI could create substantial value by coordinating complex work, but autonomy increases the consequence of design mistakes. Responsible adoption requires bounded permissions, deterministic validation, traceable actions, aligned incentives, economic controls and clear human accountability.



THE STRONGEST CROSS-CUTTING FINDING

Talent, skills and education

MAIN FINDING

The near-term constraint on AI adoption is not only access to technology. It is the current lack of practical education and understanding among professionals about what AI can do, where it fails, how it can be applied to a real workflow, and how people should work effectively alongside it.

1 AI fluency must extend beyond technical specialists

Financial institutions need specialist engineering, data, model-risk and security talent, but most value will be realised by domain professionals who can identify the right problem, provide context, judge output and redesign work. This is less about teaching everyone to code than teaching professionals to exercise informed judgement with AI.

3 Organisations must protect the development of junior expertise

Automation creates a talent paradox. Entry-level work often contains repetitive tasks that AI can perform, but those tasks have traditionally helped junior professionals learn the domain, recognise patterns and develop judgement. Graduate programmes therefore need redesign rather than reduction: rotations, supervised AI-enabled work and deliberate exposure to exceptions preserve progression towards expertise.

2 Behavioural change and learning-by-doing are essential

Employees need time to learn how to harness new tools in service of their role. Adoption cannot be achieved through access and policy alone. Safe sandboxes, guided use-case labs and manager support let professionals experiment while developing good habits. Training is most effective when it is connected to a real workflow, a measurable outcome and a clear control boundary.

4 The future profile is hybrid and adaptable

The talent market was competing aggressively for advanced technical specialists, yet participants cautioned against over-investing in a narrow skill set that may become obsolete as tools evolve. Institutions were also looking beyond traditional profiles, recognising the value of psychology, communication, creativity, systems thinking and customer understanding. The durable profile is a “supercharged professional”: someone who combines domain expertise, sound judgement and AI-enabled ways of working.

A practical capability agenda

Role-based AI literacy

Define what leaders, managers, risk teams, technologists, relationship teams and operations professionals each need to know and do.

Redesigned graduate pathways

Preserve apprenticeship, domain exposure and judgement-building while using AI to increase capacity and personalise development.

Use-case and workflow labs

Train people on real processes: mapping tasks, selecting control points, measuring value and reviewing failures.

Skills diagnostics and continuous learning

Assess current capability, target gaps and refresh learning as models, tools and roles change.

Leadership and governance education

Equip senior leaders to set risk appetite, challenge investment cases, understand agentic operating models and sponsor behavioural change.

Cross-sector talent partnerships

Industry, regulators, government and education providers should share practical examples, build pipelines and reduce duplication in scarce expertise.

The implication is clear: human capability must be treated as core AI infrastructure. Investment in models and platforms will not generate sustained value if the workforce cannot select appropriate use cases, use tools critically, redesign processes and remain accountable for outcomes.

LOCAL IMPLEMENTATION IN A GLOBAL AI SYSTEM

Conclusion and future discussion

Although the planned third discussion area was not held as a standalone segment, geopolitical and cross-border issues surfaced throughout the roundtable. Participants noted that AI was a global technology deployed through local institutions, laws, data environments, languages, market structures and customer expectations. This created a persistent tension between the desire for global scale and the need for local control.

The geopolitical dimension of AI

Model and cloud ecosystems were concentrated among a small number of global providers, while financial institutions and regulators remained accountable within national and regional frameworks. Neutrality, strategic dependence and possible technology fragmentation therefore sat alongside data residency, consumer protection, cybersecurity and operational resilience.

The discussion also highlighted uneven capacity. Large institutions and major markets can invest in compute, specialist talent, testing and governance; smaller firms and fragmented regional businesses may face much higher relative costs. If left unaddressed, this gap could shape which institutions and jurisdictions capture the productivity gains from AI and which remain dependent on external platforms.

What organisations can do

Set global principles, implement locally Common standards for accountability, data, testing and customer outcomes, with controls adapted to local legal and market conditions.	Build architectural and provider resilience Modular designs, clear interfaces, model evaluation and contingency plans to reduce lock-in and support substitution where feasible.	Map cross-border data and decision rights Know where data, models, agents and approvals sit, and which actions may cross jurisdictions or trigger local obligations.	Invest in local capability Develop leaders, domain experts, risk teams and technical talent in each market rather than relying only on a central AI function.
--	--	--	---

What regulators and policymakers can do

- Expand joint testing, regulatory exchange and evidence-sharing across jurisdictions, building on the UK–Singapore partnership model.
- Use common taxonomies and assurance concepts so firms can compare use cases, controls and maturity without assuming identical rules.
- Clarify expectations for accountability, audit trails, human oversight and operational resilience as systems become more agentic.
- Support talent pipelines, access to experimentation infrastructure and practical education for firms with fewer resources.
- Monitor provider concentration, competition, neutrality and systemic dependencies while preserving room for responsible innovation.

A PROPOSED NEXT STEP

Follow-up roundtable: Scaling AI Across Borders — Governance, Sovereignty and Talent

A second discussion could examine which controls and operating principles travel across markets; how institutions should manage model and cloud concentration; how cross-border testing can coexist with data and consumer protections; and what shared talent and education infrastructure is needed for responsible scale.

The first roundtable demonstrated the value of bringing industry, regulators, policymakers and capability leaders into the same conversation. The next phase should move from identifying tensions to developing practical options: portable control patterns, cross-border testing mechanisms, shared evidence standards and workforce models that help institutions adopt AI responsibly across different markets.

CONVENING THE DIALOGUE

About the co-organisers



The British Chamber of Commerce Singapore is an independent membership organisation that connects and represents a broad UK–Singapore business community. It supports business growth, trade, knowledge-sharing and policy dialogue, and convenes sector committees and senior forums across the bilateral and ASEAN ecosystem.



Standard Chartered is a global bank connecting corporate, institutional and affluent clients to growth opportunities across Asia, Africa and the Middle East. Singapore is a major centre for its regional and global business leadership, technology operations and innovation activity.



CFTE is a global capability platform helping financial institutions, regulators and governments build the human capabilities required for the AI era. Since 2017 it has supported more than 200,000 professionals across over 130 countries and worked with more than 100 leading organisations. Its work focuses on developing “Supercharged Professionals” who combine domain expertise, sound judgement and AI-enabled ways of working.

Acknowledgements

The co-organisers thank the roundtable co-chairs, the Acting British High Commissioner to Singapore, and all participants and observers for contributing practical experience and candid perspectives.

The discussion brought together leaders from financial institutions, asset management, technology, transformation, regulation, policy and the wider ecosystem.

Method and source note

Prepared from the roundtable briefing note, participant preparation materials, the FCA AI in Finance Observatory methodology, the supplied discussion transcript and meeting summary, and official background information from the co-organisers.

Transcript language was not quoted directly; findings were consolidated and paraphrased to preserve confidentiality. Co-organiser information accessed 3 August 2026.

Prepared under the Chatham House Rule. This report is an anonymised synthesis and does not constitute regulatory guidance or policy.

