CO-HOSTED BY

British
High Commission
Singapore

CFTE

From Copilots to AI Agents: Making Agentic AI Work in Banking

Where AI agents can create value, what makes them bankable, and how institutions can prepare.



CONTEXT AND SCOPE

About the roundtable

Why this conversation matters now

AI in banking is moving from copilots that assist people with individual tasks towards AI agents that can take action, coordinate multiple steps and interact with tools and systems. Although adoption remains at an early stage, financial institutions are beginning to explore where these capabilities can create practical value across banking workflows.

This closed-door roundtable brought together banking leaders and regulators from Singapore and the UK to share current use cases, emerging ideas and lessons from experimentation. The discussion focused on what is working, what remains difficult and what distinguishes credible deployment from experimentation.

Regulators and financial institutions face a shared question: how can agentic AI be made bankable? Participants considered the controls, governance and human oversight required, how international regulation may need to evolve, and the talent and organisational capabilities institutions will need for an agentic future.

The roundtable combined practical banking experience with UK and Singapore regulatory perspectives, while exploring which capabilities banks must build internally and where trusted external partners can accelerate responsible adoption including through technology, implementation, governance and talent development.

CORE FRAMING

The question is no longer only whether an AI agent can work, but where it can operate safely, under what controls, and with which trusted partners.

HOSTED BY British High Commission Singapore and CFTE	**FORMAT** Closed-door leadership roundtable	**VENUE** His Majesty's Trade Commissioner's Residence
OPENING REMARKS Martin Kent, His Majesty's Trade Commissioner for Asia Pacific	**CO-CHAIRS** Tram Anh Nguyen Martin Kent	**PARTICIPANTS** Senior leaders from the banking sector and the FCA

CONTEXT AND SCOPE

Purpose of the roundtable

A practical senior-level discussion on what it would take to deploy agentic AI responsibly in banking. The session was designed to move beyond general discussion of agentic AI and develop a more concrete view of where agents could realistically be used, what would make those deployments acceptable in a regulated environment, and what capabilities institutions need to build.

Identify realistic deployment opportunities within banking

Participants shared practical examples of agentic AI use cases across banking and financial services, exploring where agents are already being deployed or seriously considered and what separates credible deployment from experimentation.

Define the controls and regulatory conditions that make AI agents bankable

The discussion examined what banks and regulators need in order to trust AI systems that can take actions and interact with institutional systems, including how international regulation may need to evolve.

Building readiness: talent, operating models and trusted partners

Participants explored the institutional capabilities needed to deploy agentic AI responsibly, including which capabilities banks should build internally and how roles and skills may change.

Confidentiality and reporting scope

The discussion was conducted under the Chatham House Rule. Findings are anonymised, aggregated and non-attributable, and are presented as thematic synthesis rather than individual viewpoints. No theme should be read as the formal position of the FCA, any participating bank, the British High Commission Singapore, CFTE or any individual participant.

Co-chairs and opening remarks



CO-CHAIR

Tram Anh Nguyen

Co-Founder, CFTE
Chairwoman, Global Women in AI



CO-CHAIR

Martin Kent

His Majesty's Trade Commissioner for Asia Pacific,
British High Commission Singapore

OPENING REMARKS · MARTIN KENT

Martin delivered the opening remarks as host and co-chair, placing the discussion within the wider UK-Singapore relationship in financial services and technology. He highlighted the UK's push to accelerate responsible AI adoption, including the Treasury-backed AI Adoption Plan for Financial Services, and noted the UK's strength across both financial services and the AI ecosystem.

He also stressed the importance of cross-border collaboration, pointing to the UK–Singapore Financial Dialogue and the FCA-MAS AI-in-Finance Partnership as examples of how regulators and industry can share insights, test new approaches and support responsible adoption across markets.

No single institution, regulator or country has all the answers. The future of AI in financial services will be shaped through collaboration, shared learning and the exchange of ideas across borders.

CONTRIBUTORS

At the table



Megan Otway
FCA Financial Services Attaché
Financial Conduct Authority



Olivier Crespin
Chief Analytics Officer
DBS Bank



Dr. Jingyuan Zhao
Head, Group Data Office
OCBC



Alvin Eng
Head of UOB Innovation Group
UOB



Kim Hong
Executive Director, Deputy CDO and Head of Data Platforms
Bank of Singapore



Takashi Yamaguchi
Managing Director and General Manager, Governance Planning Department, Asia Pacific Division
SMBC



Ravindra Kumar
Chief Information Officer
Maybank

OBSERVERS



Huy Nguyen Trieu
Co-Founder
CFTE



Catriona Lindsey-Renton
Head of Financial Policy
British High Commission Singapore



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



Mansi Mehta
Head of Business Development & Partnerships, APAC
CFTE



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 an 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		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 roundtable examined what changes when banks move from AI copilots that assist people towards systems capable of coordinating work and taking action. The central questions were where meaningful value may emerge, what controls make agentic systems credible in regulated banking, and what institutional capabilities are required for responsible deployment. Throughout the discussion, the challenge was moving beyond familiarity with AI towards repeatable institutional value: clearer ways to determine where autonomy is appropriate, what evidence is required, how risk and accountability should be managed, and how institutions and regulators can learn together as deployment develops.

01 Familiarity with AI is not the same as institutional value

Confidence and proficiency with AI tools are increasing, but access and usage do not by themselves demonstrate repeatable business value. The stronger test is whether AI changes a real workflow and improves outcomes that can be measured, governed and sustained.

02 The strongest conversation starts with the workflow

Broad claims about "using AI" gave way to workflow-level questions: what an agent would actually do, which steps it performs, which systems it interacts with, where the human sits and what measurable value the change should create.

03 Moving from assistance to action raises the control threshold

Once AI can coordinate steps, use tools or initiate actions, mistakes can become operational events rather than simply poor outputs. Permissions, accountability, testing, auditability, monitoring, escalation and resilience become materially more important.

04 Experimentation is producing lessons as well as successes

Failed approaches were discussed alongside promising use cases, underlining the value of narrowing the workflow, defining human responsibility early and treating failed assumptions as evidence for better design.

05 The challenge is increasingly to scale the framework around AI

The emphasis was not on maximising autonomy, but on reusable approaches for deciding how much autonomy is appropriate, what controls apply, what evidence is required and when humans must remain decisive.

06 Agentic AI is an institutional-readiness challenge

Responsible deployment depends on domain ownership, AI and data capability, governance judgement, workflow redesign, clear decision rights and people who can challenge and supervise increasingly capable systems.

OVERALL SYNTHESIS

Agentic AI is not simply a more powerful copilot. Allowing AI to act inside banking workflows changes the value, control and organisational questions that institutions must solve. The next challenge is not only scaling agents, but scaling the confidence, controls and evidence that allow institutions to use them responsibly.

Part 1 FROM EXPERIMENTATION TO DEPLOYMENT

From copilots to AI agents: where is real value emerging?

From AI proficiency to actual adoption

Familiarity with AI was clearly increasing, but participants drew a distinction between tool proficiency and repeatable business value.

STEP 01
AI access
Employees have access to AI tools.

STEP 02
AI use
They use them in parts of their daily work.

STEP 03
Workflow change
AI becomes embedded in how a defined process is carried out.

STEP 04
Measurable value
The institution can show improved outcomes, capacity, risk management, service or economics.

Availability is not the same as adoption. Adoption is not the same as measurable value.

The workflow test

Moderation focused the discussion on real workflows rather than generic AI use cases. Three questions were put to every candidate example.

1
What does the agent actually do?
Which steps it performs, which systems it interacts with, what action it can initiate and what information it needs.

2
Where does the human sit?
Who verifies, approves, handles exceptions, intervenes, and who remains accountable for the outcome.

3
What measurable value should it create?
Time, capacity, customer responsiveness, reduced manual effort, control effectiveness or better decision support.

What makes a workflow more credible for bounded agency?

Repeatability
Identifiable stages and recurring patterns, not entirely bespoke work.

Clear data and context
Reliable access to information without uncontrolled data exposure.

Bounded authority
A clear definition of what the agent may and may not do.

Reversibility
Errors can be detected, stopped or corrected without disproportionate harm.

Human escalation
Exceptions can be routed to someone with the authority to intervene.

The strongest opportunities are not necessarily where AI appears most impressive. They are where a clearly defined workflow allows an agent to coordinate useful work within understandable boundaries and produce measurable value.

Part 1 IMPLEMENTATION LESSONS

What banks are learning from experimentation

One of the most valuable features of the discussion was the willingness to speak not only about promising applications but also about approaches that had not worked as expected. Early deployment experience provides evidence about where an AI system's boundaries actually sit: a disappointing pilot may expose an unclear objective, poor workflow design, inadequate context, an unsuitable human-AI division of labour, or controls added too late.

What experimentation was teaching institutions

01 Start narrower than instinct suggests

A broad assistant may impress in demonstration but becomes harder to govern across multiple systems, contexts and decision types. Narrowing the mandate makes both performance and accountability clearer.

02 Design the human role at the beginning

"Human in the loop" is not a sufficient control unless the institution defines who intervenes, when, with what information, and with what authority.

03 Integration is part of the use case

An agent that performs well in isolation may fail to produce value once introduced into a real process with legacy systems, changing rules and existing controls.

04 Exceptions reveal the real operating model

The normal path is usually easier to automate. Edge cases and unusual circumstances show whether the process can withstand autonomy and where judgement must stay human.

05 Measure value and risk together

A technically successful agent is not necessarily a scalable one. Institutions need to assess business outcomes alongside error rates, escalation volume, control effectiveness, operational resilience and total cost.

FAILURE AS GOVERNANCE EVIDENCE

Experimentation can also be treated as governance evidence.

What assumption was wrong?

Was the workflow less standardised than expected?

Where did control become difficult?

Was authority too broad, context incomplete, escalation unclear?

What would have to change?

Could scope, architecture, data or the human role be redesigned?

Institutions should treat experimentation as a way to discover the safe operating boundary of an AI system, and not simply as a demonstration that the technology works.

Part 2 FROM ASSISTANCE TO ACTION

Making AI agents bankable

Moving from copilots to agents changes the institutional challenge. A copilot may generate information for a person to consider. An agent may coordinate several steps, access tools, retrieve and update information, communicate with systems and initiate actions. As autonomy grows, continuous human review may become impractical. The control environment has to move from reviewing individual outputs towards governing the authority, behaviour and consequences of the system as a whole.

Authority

What is the agent permitted to do?

Systems, tools, actions it may initiate, limits and prohibited actions.

Identity and permissions

How does the institution know who is acting?

Identifiable roles and least-privilege access matched to the mandate.

Accountability

Who remains responsible?

Clear ownership of the workflow, exceptions, failures and final decisions.

Human oversight

Where must a person intervene?

Decision boundaries, reviewer authority and stop or escalate conditions.

Assurance and auditability

What evidence is required?

Pre-deployment testing and preserved records of inputs, actions and approvals.

Resilience

What happens when the agent fails?

Fallback modes, stop conditions, escalation routes and transaction limits.

A practical architecture for controlled autonomy

01

AI interprets and proposes

Probabilistic models for language, context and judgement support.

02

Controlled systems validate

Deterministic rules, calculations, schemas and policy checks.

03

Humans retain decision rights

Exceptions escalate to accountable people, not an undefined loop.

04

Monitoring closes the loop

Log actions, evaluate outcomes and feed evidence into the next iteration.

As autonomy increases, governance needs to move closer to the point at which the system acts.

Part 2 REGULATORY DIRECTION

From principles to operational safeguards

Banks wanted greater clarity on risk, compliance and the evidence required before granting greater autonomy. This does not necessarily mean a wholly new regulatory regime for every AI application. The emerging direction in both the UK and Singapore combines existing regulatory principles, practical testing, industry collaboration and increasingly operational approaches to AI risk.

UNITED KINGDOM

Preparing for more autonomous finance

The FCA's Mills Review, published in July 2026, examines how AI could reshape retail financial services through 2030 and beyond, monitoring the transition towards more autonomous models, scaling the FCA's AI Lab and enabling the foundations for agentic finance. The FCA continues to describe its approach as principles-based and outcomes-focused, relying on existing frameworks rather than a separate AI rulebook, while AI Live Testing and the Supercharged Sandbox let firms and the regulator learn from systems under realistic conditions.

- Existing accountability remains relevant — AI does not remove responsibilities firms already hold.
- Testing becomes increasingly important — confidence comes from evidence of behaviour, not design descriptions.
- Regulation learns from deployment through live testing and industry engagement.

SINGAPORE

Moving towards operational AI risk management

Singapore's approach has developed through complementary initiatives. MAS has proposed Guidelines on AI Risk Management for financial institutions, still at consultation stage, covering governance, AI inventories, lifecycle controls, human oversight, third-party risk and monitoring. Project MindForge produced an industry-developed AI Risk Management Toolkit to help institutions operationalise risk management across traditional AI, generative AI and emerging agentic systems.

In July 2026, MAS and industry published Safeguards for Agentic Finance at Runtime (SAFR) under BuildFin.ai — real-time safeguards intended to keep agent behaviour within predefined mandates, policies and risk boundaries once continuous human intervention is no longer practical.

The common direction

FROM	TOWARDS
AI governance primarily as policy	**Governance embedded in the operating system**
Reviewing model outputs	**Governing actions and permissions**
Point-in-time approval	**Continuous monitoring and evidence**
Generic "human oversight"	**Explicit decision and escalation boundaries**
Institution-only learning	**Structured learning between firms, regulators and the ecosystem**

In practice, AI governance is becoming more operational: defining who or what may act, under which mandate, with what evidence, and what happens when the system reaches its boundary.

Part 3 INSTITUTIONAL READINESS

Building readiness: talent, operating models and trusted partners

The final section focused on institutional readiness rather than the technology itself. The issue became whether institutions themselves are ready to own and govern increasingly capable systems, and to redesign workflows, roles, decision rights, controls and relationships with external partners. The key distinction is between capabilities the institution must own and capabilities partners can responsibly accelerate.

01 What capabilities must a bank build and retain internally?

Domain ownership
People who understand the underlying banking process and remain accountable for the business outcome.

AI and data capability
The ability to evaluate, integrate, monitor and challenge AI systems.

Governance and risk judgement
The ability to determine what agents are permitted to do and where limits apply.

Workflow redesign
The ability to redesign processes around AI rather than add AI to the existing process.

Leadership and decision rights
Senior leaders able to set objectives, risk appetite, accountability and escalation structures.

02 What can responsibly be entrusted to a partner?

Technology capability
Can the partner deliver reliable technology appropriate for institutional deployment?

Financial-services expertise
Does it understand banking workflows, controls and the operating environment?

Security
Can it meet the institution's security and resilience expectations?

Regulatory understanding
Does it understand the regulatory and governance context the technology will operate in?

Implementation capability
Can it move beyond proof of concept to integration and operating deployment?

Talent development
Can it strengthen the institution's own capability rather than create permanent dependency?

THE PARTNER PRINCIPLE

External partners can accelerate technology, implementation and capability building. But business ownership, risk appetite and accountability remain institutional responsibilities.

WHAT REMAINS OPEN

- Which banking workflows are mature enough for bounded agency today?
- How much autonomy is appropriate for different processes?
- What evidence should be required before moving from testing to deployment?
- Which decisions must remain explicitly human-owned?
- How should banks and regulators learn from deployments as the technology evolves?

Responsible adoption depends not only on whether an agent works, but whether the institution around it is capable of owning, challenging, governing and using it.

THE STRONGEST CROSS-CUTTING FINDING

Scale the framework, not only the agent

The roundtable suggested that the immediate challenge for regulated financial institutions may not be how quickly they can maximise AI autonomy. A more pressing question is whether they can develop a repeatable institutional framework that makes different levels of autonomy understandable, governable and progressively deployable. Scale begins not by giving agents broader authority, but by making the control framework reusable across different workflows.

Five elements of a reusable agentic framework

1 Classify the level of agency

Distinguish assistance, recommendation, delegated execution and more autonomous action rather than treating all AI systems alike.

2 Define the mandate and decision rights

Specify which actions the agent may take, which remain prohibited and which require explicit approval.

3 Set an evidence threshold for deployment

Determine what testing, assurance and outcome evidence must exist before moving to production.

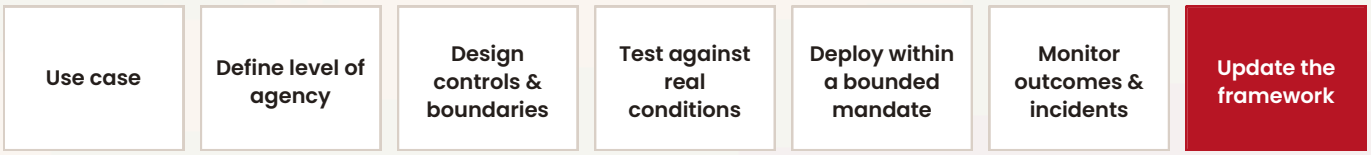
4 Operate controls at runtime

Apply permissions, policy checks, limits, monitoring and escalation when the system acts, not only at approval.

5 Learn systematically from failure and exceptions

Feed incidents, unexpected behaviour, overrides and unsuccessful experiments back into system design and governance.

From pilot to reusable practice



A bank does not need to reinvent its governance approach for every new agent. If institutions can agree internally, and increasingly learn across the sector, on common ways to classify autonomy, define authority, capture evidence and escalate exceptions, future deployments can begin from a stronger control baseline.

The scalable asset may therefore be not only the AI agent itself, but the institutional framework that determines when and how an agent can be trusted.

LOCAL IMPLEMENTATION IN A GLOBAL AI SYSTEM

From shared discussion to shared practice

The discussion brought together banking perspectives from Singapore and the wider Asian region alongside UK regulatory experience. That diversity reinforced a practical reality: AI technologies are global, but financial institutions remain accountable within local regulatory, legal, data and customer environments. The goal is therefore unlikely to be one identical AI rulebook across markets. A more practical aim is portable principles, shared evidence and common control concepts that institutions can apply locally.

What financial institutions can do

- **Use common language for use cases and autonomy** — describe AI systems consistently so risk and maturity can be compared across units and jurisdictions.
- **Map cross-border data and decision rights** — know where data, models, agents and approvals sit and when a workflow may trigger obligations elsewhere.
- **Share lessons, not only success stories** — failures, overrides and abandoned approaches produce valuable evidence about safe operating boundaries.
- **Build capability locally** — ensure domain experts, risk teams and technologists in each market can understand and challenge AI.

What regulators and the wider ecosystem can do

- **Expand joint testing and evidence-sharing** — create mechanisms for institutions and regulators to learn from controlled deployment.
- **Develop interoperable assurance concepts** — shared language for autonomy, evidence and auditability reduces duplication even where regulation differs.
- **Translate principles into practice** — frameworks matter when the professionals who redesign and operate workflows can apply them.
- **Build capability alongside governance** — skills development should advance with technological and regulatory change, not follow it.

COMMON LANGUAGE AS INFRASTRUCTURE

The FCA AI in Finance Observatory taxonomy, developed with CFTE and used as a shared reference for this roundtable, illustrates one practical form of collaboration. Its role is not to determine whether an application is compliant or low risk; it provides a common way to describe what the use case is, where it operates, why it matters and what is known about its implementation.

CFTE's work across industry, regulators and education places it in a distinctive position to help translate emerging AI principles into practical capability. Regulators can set expectations and create environments for responsible experimentation; financial institutions contribute evidence from real deployment; education and capability partners turn those lessons into shared frameworks, practical skills and repeatable ways of working.

The next phase of agentic finance will depend on combining practical deployment evidence, credible controls, institutional capability and cross-border learning.




RESIDENCE OF
HIS MAJESTY'S
TRADE COMMISSIONER




Ravindra Kuma
Alvin Eng


Ravindra Kumar
Dr Jingxian Zhao


Tram Anh Nguyen


Catriona Lindsey-Renton


Amy Ross


RESIDENCE OF
HIS MAJESTY'S
TRADE COMMISSIONER

CONVENING THE DIALOGUE

About the organisers

British High Commission Singapore

The British High Commission Singapore represents the United Kingdom in Singapore and supports the bilateral relationship across government, trade, investment, policy and wider areas of cooperation. Its work in financial services includes the annual UK-Singapore Financial Dialogue and support for responsible AI adoption across markets.

CFTE

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 more than 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, all participants and observers for contributing their experience and perspectives.
The session brought together banking, regulatory and capability perspectives to explore how agentic AI may develop in regulated financial services.

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 10 September 2026.

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