Creating an Artificial Intelligence Framework in Investment Operations

Post-event report | 20 May | Franklin Templeton, Morrison Street, Edinburgh

Executive summary

The event brought together investment operations participants and subject matter experts to explore how firms can move AI from pilots to production safely. The discussion was highly engaged and reinforced that the challenge is not whether firms should adopt AI, but how they can adopt it with confidence, clear accountability and auditable evidence.

Most attending firms described themselves as early in their AI journey, with access to enterprise tools such as Copilot and GPT-like services. Only two firms considered themselves further along, typically through Technology-led deployment rather than citizen development. Only one firm described a pre-approval governance structure for AI deployments.

Headline messages

- AI governance should enable innovation rather than block it.
- Senior management pressure for automation and cost reduction is ahead of formal governance maturity.
- AI risk is generally not yet integrated as a standalone component of governance structures.
- The immediate need is a practical toolkit: ownership, inventory, risk classification, due diligence, deployment readiness and ongoing monitoring.

Session snapshot

Session element	Summary
Purpose	Consensus on practical AI deployment and oversight in investment operations.
Method	Chatham House-style discussion using maturity positioning, facilitator prompts, expert challenge and control checklists.
Expected outputs	AI Governance & Oversight Principles plus an Operations Toolkit for AI Deployment & Governance.

Event summary visual based on anonymised notes.

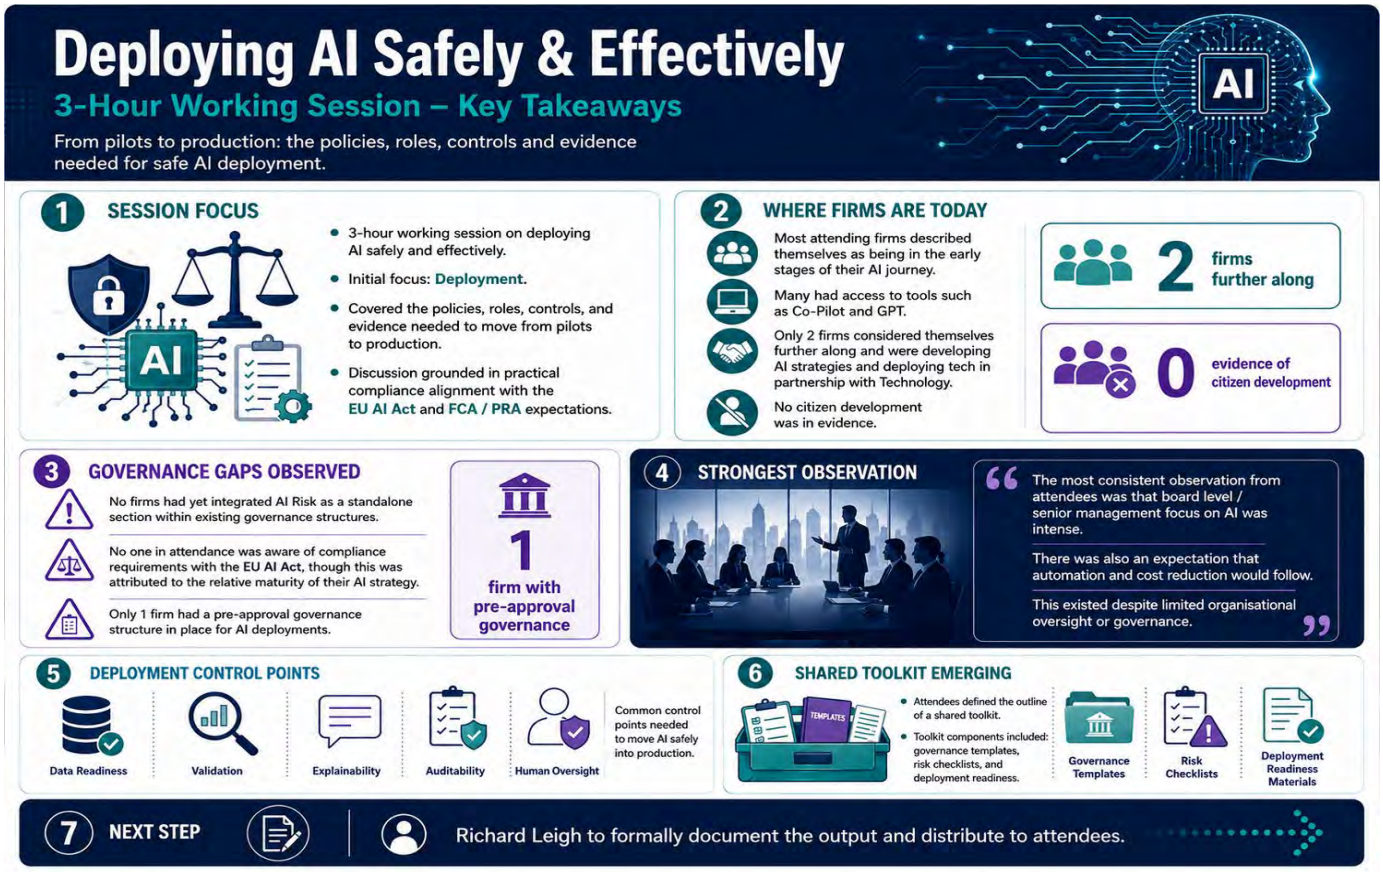

1. Discussion design, technology baseline and maturity

How the event was structured

The agenda moved from context to framework design. It started by establishing why AI matters now, then asked each participant to position their organisation on a maturity curve before moving into ownership, risk classification, due diligence and governance. The event was deliberately framed as a working session rather than a lecture: the aim was to capture what firms know, what they do not yet know, and where shared best practice is needed.

What the maturity discussion showed

- Use is moving from fixed-data automation into variable-data and judgement-support cases.
- Most firms are exploratory or controlled/developing rather than integrated/optimised.
- Some firms have local activity but limited cross-divisional consistency, creating uneven levels of progress and engagement.
- There was no evidence of citizen development in the attending group; production AI activity was being progressed with Technology.
- Awareness of EU AI Act obligations was limited, possibly because current use cases are still largely low-risk productivity or experimentation.

Maturity level	Observed position	Practical implication
Exploratory / ad hoc	Most attendees	Potential for fragmented pilots, unclear ownership and inconsistent approvals.
Controlled / developing	Two firms further along	Emerging AI strategy, Technology partnership and more formal approval routes.
Integrated / optimised	Not evidenced	Target state: Board visibility, AI Council, inventory, SMCR accountability, monitoring and assurance.

Technology baseline: governance expectations rise as systems become more probabilistic and autonomous.

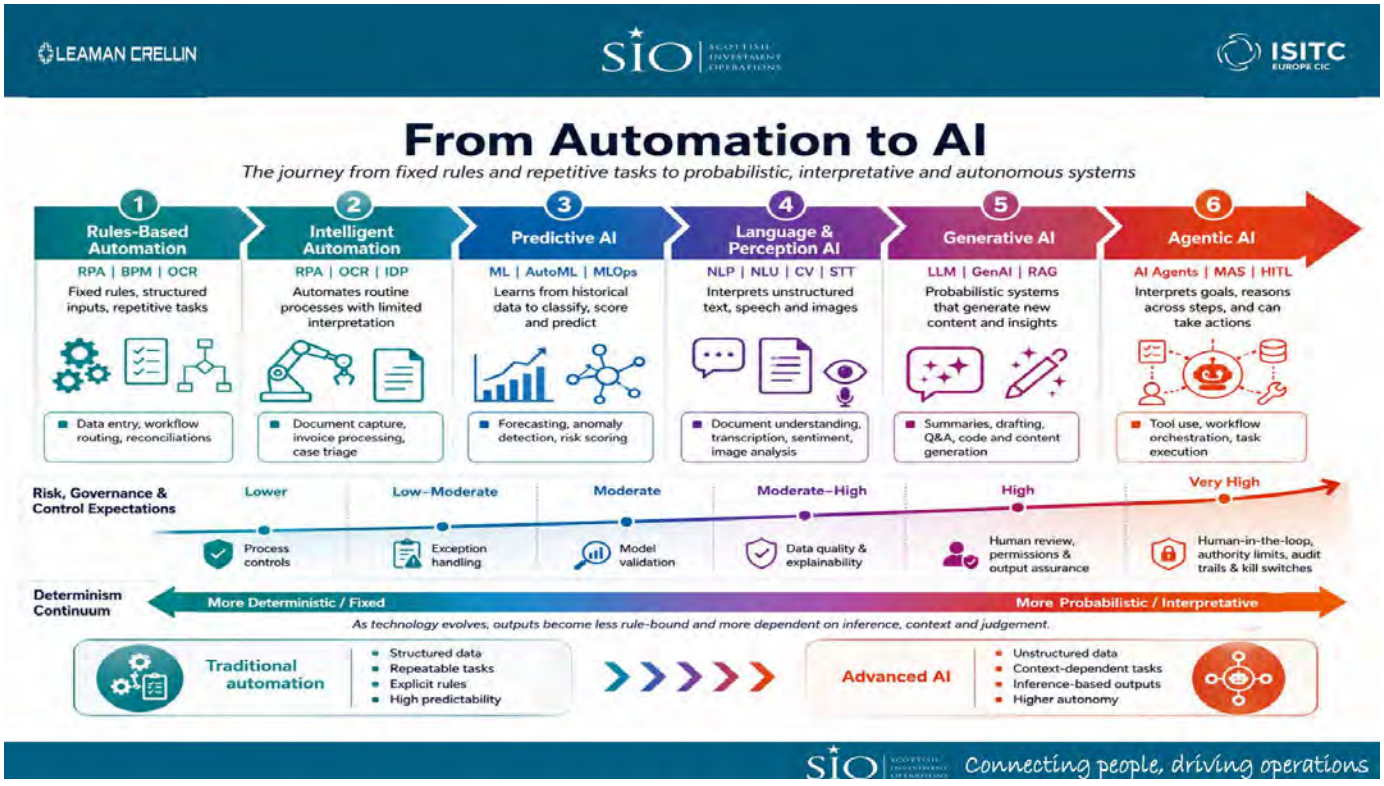

Interpretation

The baseline slide helped distinguish rules-based automation from probabilistic AI and agentic workflows. This matters because the control model changes as systems move from deterministic, repeatable processes to outputs based on inference, context and judgement. The roundtable identified autonomy as a spectrum rather than a binary: the key question is when additional challenge, human oversight and evidence are required.

In operations, the most realistic early use cases include summarisation, exception triage, regulatory intelligence, client reporting support, settlement assistance and data quality analysis. These may start as productivity aids, but the risk profile changes if they become embedded in live processes or important business services.

2. Governance, accountability and risk classification

Governance and accountability

The strongest observation from the discussion was the tension between senior management focus on AI and the relative immaturity of organisational oversight. Attendees described senior and Board-level attention as intense, with an expectation that automation and cost reduction would follow. However, few firms had put in place the governance infrastructure needed to support those expectations safely.

Area	What good should include
Board and senior management	AI strategy, risk appetite, challenge and periodic MI on AI risk exposure.
SMCR accountability	Named accountability for AI risk, technical challenge and business outcomes.
AI Council / forum	Use-case approval, technical review, risk sign-off, escalation and re-approval.
Ownership model	Clear split between business owner, model/tool owner, data owner and control owner.
Existing frameworks	AI embedded into enterprise risk, operational resilience, compliance, vendor, data and model risk processes.

Key discussion point

The operating model does not need to be identical across firms. What matters is that the firm can explain who approved the use case, what risks were assessed, what controls were implemented, who owns the outcome, and how live performance will be monitored.

Risk classification framework: proportionality should determine the governance burden.

Risk Classification Framework

Risk level	Category	Example use cases	Minimum controls required
High	Consumer-facing autonomous decision	Real-time investment recommendations; credit eligibility; automated advice suitability assessment.	Board approval; named SMF accountability; independent validation; human override; Consumer Duty impact assessment; audit trail.
High	Financial crime and surveillance	AI transaction monitoring; market surveillance; fraud detection with automated block or flag capability.	Named SMF accountability; independent validation; false positive rate management; MLRO oversight; full audit trail.
High	Operational AI in important business services	AI embedded in settlement processing, client reporting or regulatory transaction reporting.	Operational resilience mapping; continuity planning; third-party CTP assessment; kill switch; SMF24 or equivalent oversight.
Medium	Internal decisioning support	AI drafting credit memos; summarising regulatory intelligence; generating compliance reports.	Documented human review; accuracy monitoring; named business owner; periodic validation.
Medium	Customer-facing non-advisory	AI chatbot for general enquiries; complaint routing; personalised marketing.	Consumer Duty review; bias testing; escalation to human agents; vulnerability monitoring.
Low	Internal productivity	Meeting summaries; internal communications; non-production code review.	Acceptable use policy; data classification controls; periodic review.

Connecting people, driving operations

Risk inventory and classification

Participants discussed whether firms should maintain a standalone AI inventory or extend existing model, system, vendor and risk inventories. Either approach can work, provided the inventory is complete, owned and maintained. Minimum fields should include purpose, owner, data processed, decision impact, customer data, regulated activity, third-party tools and shadow AI.

- High-risk use cases include consumer-facing autonomous decisions, financial crime or surveillance AI, and AI embedded in important business services.
- Medium-risk use cases include internal decision support and customer-facing non-advisory AI.
- Low-risk tools still require acceptable use and data controls, especially to prevent confidential or customer data leakage.
- Classification must be dynamic and revisited at least annually, and after material change, incident or expansion of scope.

3. Deployment controls, ongoing governance and recommended outputs

Deployment due diligence

The deployment discussion focused on how firms can evidence that an AI system is ready for live use. The common control points were data readiness, validation, explainability, auditability and human oversight. The group recognised that firms should not rely solely on vendor claims for higher-risk systems, and that due diligence outputs should be retained so decisions can be reconstructed and explained.

Control point	Evidence expected
Data readiness	Source, lineage, permitted use, quality, completeness, timeliness and controls preventing restricted data input.
Validation	Accuracy, limitations, error rates, bias, behaviour under poor data, and independent validation for higher-risk systems.
Explainability	Clear description of what the system does, what it cannot do, and how outputs should be used.
Auditability	Approval records, testing results, prompts/inputs where appropriate, overrides, exceptions and decision trail.
Human oversight	Defined human-in/on-the-loop arrangements with authority to review, override, escalate or stop use.

Ongoing governance

AI governance must continue after deployment. Monitoring should cover performance, overrides, escalations, data quality, incidents and near-misses. Risk classifications should be reviewed at least annually and whenever the model, vendor, data, usage or regulatory context changes.

Risk	Detection approach
Drift	Compare live performance with baseline; monitor input data, output quality, exceptions, overrides and user corrections.
Bias	Test outputs across populations, products, geographies and scenarios; check proxy variables and feedback loops.
Hallucination	Use source checking, approved reference material, retrieval/source-grounding and human review for higher-risk outputs.

Recommended post-event toolkit

The session should now be converted into a practical set of reusable materials that firms can adapt. The toolkit should be proportionate and should support safe adoption, rather than creating unnecessary friction.

Output	What it should contain
Governance & Oversight Principles	Board and senior management accountability, SMCR alignment, proportionate governance and lifecycle monitoring.
Deployment readiness checklist	Use-case owner, purpose, data, risk tier, approvals, validation, explainability, oversight and monitoring.
Risk framework	Risk tiers, examples, minimum controls, change triggers and reclassification criteria.
Governance template	AI Council or approval forum terms of reference, membership, decision rights, escalation and MI cadence.
AI Council charter	Scope, responsibilities, relationship with existing risk committees, and incident escalation routes.
Follow-up survey	Re-test attendee maturity later in 2026 and identify areas where SIO / ISITC Europe can support collective progress.

Immediate actions

- Richard Leigh to document the toolkit outputs and circulate the post-event report to attendees.
- Use the report to validate whether attendees agree with the themes, gaps and proposed toolkit structure.
- Continue exploring a shared sandbox / innovation hub for cross-firm operational pain points, rather than purely individual firm use cases.