Classification: External



Governing Agentic AI in Financial Services

Practical Methods for Navigating Responsible Innovation

Dr. Swapnil Govind Dambe

Trust by design for ethical, agent-led transformation

Classification: External

Points to Know



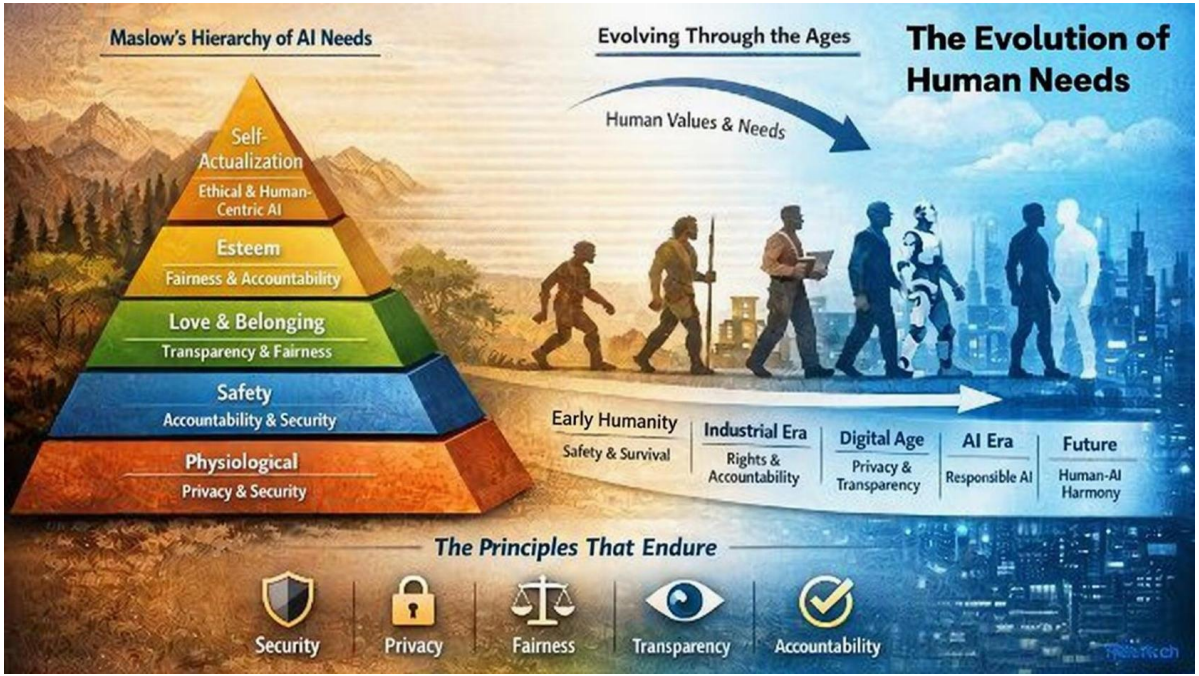
The storyline:

- **Start with trust**
  Technology evolves, human needs endure
- **See the shift**
  From branch to agent, and who must now be governed
- **Govern by design: Human centric way**
  Five lifecycle phases, scaled with FINOS AIGF
- **Lead with ethics**
  Five principles for better client outcomes

The storyline: build trust, see the shift, understand the risk, govern by design, then lead with ethics.

Classification: External

Technology evolves. Five human needs endure.



Five human needs have outlasted every revolution.

Tools keep changing. People still expect the same: **safe, fair and trusted service.**

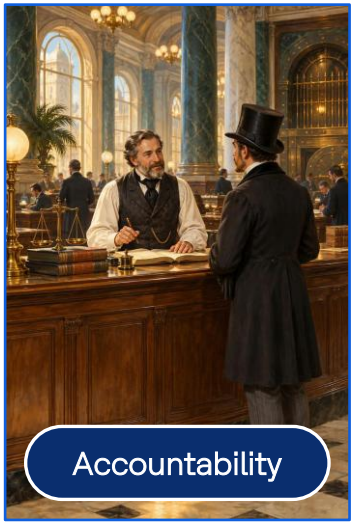
No agent should move money without earning all five.

Security. Privacy. Fairness. Transparency. Accountability. The five needs every agent must earn.

Kick-start: every era asked one question: can I trust it? In financial services, governance is how agents earn that trust.

Classification: External

From branch to agent: how financial services evolved



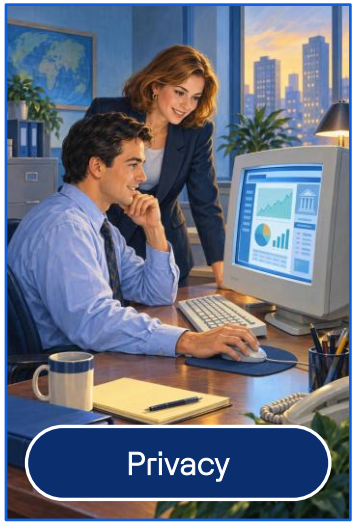
Face-to-face relationships

Know who answers to you



Phone-based networks

Verify who is on the line



Digital transformation

Protect data going online



Data-driven and mobile

Treat every client equally



Agentic and human-led

See and explain every action

1800s–1900s | Mid-20th century | 1990s–2000s | 2010s–today | Now and next

Each era put one human need in focus. Agentic must earn all five, and governance is how.

Classification: External

From branches to apps to agents, would you let one move a client's money?

Every era gave customers more reach. Agentic gives them action.

Agentic governance is not about controlling technology. **It's about empowering people around agents.**

Not governing a model, but governing how people:

01
Build agents **safely**

02
Question and **verify outputs**

03
Make the **final call**

04
Own the **outcome**

Classification: External

Why agent governance is essential

Dimension	Why it is essential	If left ungoverned
Agent behaviour	Agents act, not just answer. Intent, authority and tool use must stay within set limits.	Unauthorised or harmful actions, unfair outcomes, client harm and no clear accountability.
Orchestration	Many agents hand off work. Roles, permissions and handoffs need one accountable owner.	Access creep, duplicate payments and cascading errors that nobody can trace or stop.
Cost to value	Every step consumes tokens, tool calls and credits, whether or not it helps the client.	Runaway spend, paying for activity instead of outcomes, and a business case that quietly erodes.
Performance KPIs	Trust has to be measured: safety, accuracy, takeover speed and cost per outcome.	Silent drift, false confidence, a board blind to risk, and findings found by regulators first.

Ungoverned agents fail quietly, at machine speed.

Why now: agents now act as digital workers, and the EU AI Act, ISO/IEC 42001 and NIST AI RMF demand demonstrable control.

Classification: External
Classification: Internal

2026 OPEN SOURCE IN FINANCE FORUM New York

Human-Centric Responsible AI Governance for Trusted Enterprise Adoption

The goal is to ensure Agentic AI systems are trustworthy, transparent, fair, and aligned with human values, while mitigating risks like bias, misuse, or unintended harm.

Core AI Governance Aspect	Dimensions	Intent (Purpose)	Responsible AI Tenets	Coverage	Key Outcomes
Strategic (Vision, Ethics, and Oversight)	**Human-Centric Governance & Ethics**	Establish independent ethical oversight & governance bodies	**Accountability & transparency:** board ownership and traceable approvals. Fairness, privacy & security: ethical policy oversight.	AI Ethics & Oversight Board AI Approval & Model Risk Board	Ethical, value-aligned AI with robust, inclusive oversight
Functional (Processes and Lifecycle Management)	**Human-Centric AI Life Cycle Governance**	Embed human-centric checks & accountability across full AI lifecycle	**Fairness & transparency:** bias checks and explainability. Accountability, privacy & security: lifecycle owners and approval controls.	**Use Case Intake → Model Dev → Validation → Deployment → Monitoring**	Fair, safe, continuously governed AI lifecycle with transparency
Technical (Risk Mitigation and Security)	**Human-Centric AI Risk, Security & Compliance**	Identify, assess & mitigate risks; enforce security, privacy & global standards	**Security & privacy:** access, encryption and consent controls. Fairness, transparency & accountability: testing, model lineage and audit evidence.	Risk Engine & Tuning Policy/Compliance (NIST, ISO 42001, etc.) Privacy, Consent & Impact Assessments	Lower risk exposure, secure & auditable models, strong regulatory alignment
Operational (Enablement and Culture)	**Operational Humen Centric AI Culture Enablement, Tools & KPI Cockpit**	Foster a responsible AI culture through ethics, collaboration, and continuous learning	**Accountability & transparency:** human review, training and KPI reporting. Fairness, privacy & security: monitored use and incident response.	Role-based ethical training, AI literacy, Human-in-the-Loop tools, RAI dashboards, and collaboration platforms	Ethically empowered teams with measurable impact on fairness, safety, and societal value, fostering continuous learning and innovation

Human at the center
Every control is designed so people feel safe, secure and in charge. **ALL DIMENSIONS**

Human in command
Leaders set AI purpose, risk appetite and limits, and can stop any system. **STRATEGIC**

Human in the loop
People review and approve high-impact decisions before they take effect. **FUNCTIONAL + TECHNICAL**

Human on the loop
People supervise live AI and agents and can intervene or override.: **OPERATIONAL**

HCLTech

Classification: External

From empathy to execution: ethical by design

1 Assess	2 Design	3 Validate	4 Operate	5 Evolve
The question we ask Who is affected, and is an agent the right answer?	**The question we ask** What may the agent see, say and do? When does a human step in?	**The question we ask** Is it fair, safe and resilient to misuse and runaway loops?	**The question we ask** Can we trace every action and weigh its cost against value?	**The question we ask** Has autonomy grown? Re-earn trust, or retire access.
Key Agent Governance activities Risk-tier the use case; set autonomy level **Agent metrics:** Baseline KPIs and success targets for the agent	**Key Agent Governance activities** Define roles, least-privilege access and approval gates **Agent metrics:** Thresholds for accuracy, escalation, cost	**Key Agent Governance activities** Test fairness, misuse and loop limits before go-live **Agent metrics:** Pre-launch task success and error rates vs. targets	**Key Agent Governance activities** Log every action; enforce spend, retry and takeover caps **Agent metrics:** Live success rate, latency, cost per task	**Key Agent Governance activities** Review KPIs; re-earn trust or retire access **Agent metrics:** Trend scorecard raises or cuts autonomy
Who answers Business owner	**Who answers** Engineering lead	**Who answers** Independent risk	**Who answers** Service owner	**Who answers** Governance lead

Full lifecycle coverage: strategy to enterprise scale. Governance you can prove, not just promise.

Classification: External

Scaling with confidence: FINOS AIGF in practice

Practice	What it means in action
01 Industrialise	Reuse and operationalize proven FINOS AIGF risk and control patterns instead of starting blank.
02 Verticalise	Turn local policy into runtime controls: permissions, approval gates and action logs.
03 Optimise	Optimise for outcomes, not activity: judge cost per successful outcome, not tokens consumed.
04 Democratise	Democratize with responsibility: share approved tools, usage limits and non-confidential lessons.

Classification: External

Five principles for better outcomes

Beyond governing agents: using them ethically to improve outcomes and client experiences.

Accountability	Transparency	Fairness	Privacy	Security
Governing agents A named owner for every agent action and decision.	**Governing agents** Every action is logged, traceable and explainable.	**Governing agents** Bias tested before go-live and monitored after.	**Governing agents** Data used only for its purpose and consent.	**Governing agents** Least-privilege access and hard spend limits.
Better experience Clients know who stands behind every outcome.	**Better experience** Plain-language reasons build client confidence.	**Better experience** Consistent, equal treatment for every client.	**Better experience** Clients share more because their data is safe.	**Better experience** Money and data stay protected at machine speed.

Govern with control. Use with ethics. Deliver trusted financial services.

Why it matters: the principles that govern agents also shape how clients experience them.

Classification: External



Thank you

Questions and conversation welcome

Governing agents so people can trust them

Dr. Swapnil Govind Dambe

Trust by design for ethical, agent-led transformation