

## — FINANCIAL SERVICES · VERTICAL 01

# For institutions where "usually right" is not a release criterion.

Cohorte's programs for banks, insurers, asset managers, and fintechs. Mapped to SR 11-7, SS1/23, DORA, and the AI Act. Built on the founder's published research in conformal prediction and self-consistency.

## FOR

CIOs, CROs, Heads of AI,  
L&D Directors

## REGIMES COVERED

SR 11-7 · SS1/23 · DORA ·  
AI Act · ISO 42001

## CLOSEST REFERENCE

PwC AI Factory · 60+  
shipped systems

## DIRECT LINE

[charafeddine@cohorte.co](mailto:charafeddine@cohorte.co)

## 01 WHY FS NEEDS DIFFERENT TRAINING

# In a bank, an AI output is a regulated decision.

A credit decision, a fraud flag, an underwriting note, a suitability assessment. Each is governed. Each has a regulator. Each has a model-risk function asking: what does this system do when it is wrong, and how would you know?

**SR 11-7 / SS1/23****The supervisors named the gap.**

Both regimes require documented evaluation, verification, and ongoing monitoring of any model influencing a business decision. Most "AI training" teaches the model. They do not teach independent validation, ongoing monitoring, or governance.

**DORA & AI ACT****Two regimes, one operational obligation.**

DORA (in force since January 2025) covers ICT third-party risk. The EU AI Act puts most credit-scoring and insurance pricing into the High-Risk Article 6 category, triggering Articles 12 (logging), 14 (human oversight), 15 (accuracy / robustness).

**YOUR PEERS CHOSE ALREADY****The market is moving.**

JPMorgan deploys LLM Suite to ~200,000 staff under strict access controls. Goldman has GS AI Platform with internal evaluation gates. The question is no longer "should we train on AI." It is "how do we train on AI without breaking model risk."

## 02 WHY COHORTE

# The methodology your model-risk team is going to ask for.

Taught by the practitioner who published it. Four named primitives. Each with a published source. Each with an open-source reference implementation.

**CONFORMAL PREDICTION****Provably right within stated bounds.**

The mathematical primitive that converts a model output into a statistically rigorous confidence interval. The closest analogue to the VaR / ES discipline already familiar to a market-risk team.

**Source:** The conformal-calibration method, from our reliability-certification paper. [teams.cohorte.co/research](https://teams.cohorte.co/research).

**SELF-CONSISTENCY****Detecting confabulation.**

If the same prompt run five times produces five different answers, the model is not reasoning. The verification gate that catches this before output reaches a customer-facing channel or a regulator-facing log.

**Source:** The self-consistency method, from the same paper. [teams.cohorte.co/research](https://teams.cohorte.co/research).

**EXPLOITATION SURFACE****The red team your CISO asked for.**

Prompt injection. Indirect injection through retrieval corpora. Tool-call hijacking. The systematic taxonomy your security team uses to scope penetration testing of an LLM application.

**Source:** Our 10,000-trial exploitation-surface taxonomy. [teams.cohorte.co/research](https://teams.cohorte.co/research).

**REFERENCE STACK****Open-source. Inspectable. Yours.**

TrustGate (verification gates), Guardrails (policy enforcement), Agent-Auth (authorization), Agent-Monitor (observability). Six public repositories. Your engineers audit the layer before deploying it.

**Where:** [github.com/Cohorte-ai](https://github.com/Cohorte-ai) · MIT / Apache 2.0.

03 USE CASES WE TRAIN AGAINST

# The AI work financial institutions are actually shipping.

Where each sits. Risk category. Verification primitive. The honest matrix for the AI workflows your team will own.

USE CASE	WHERE IT SITS	RISK CATEGORY	VERIFICATION PRIMITIVE
<b>Customer service copilot</b>	Retail bank, insurer	Conduct & mis-selling	Self-consistency, response guardrails, escalation gates
<b>Underwriting assistant</b>	Life / P&C insurer	AI Act high-risk · fairness	Conformal bounds, adversarial fairness tests
<b>Credit memo drafting</b>	Commercial / corporate bank	Credit risk · disclosure	Source attribution, citation-grounded generation, audit logs
<b>KYC / AML alert triage</b>	Compliance, financial crime	AML & sanctions	Decision logs, human-in-the-loop gates, explanation generation
<b>Suitability assessment</b>	Wealth, retail brokerage	MIFID II suitability	Rule-based gates, output classification, regulator-facing logs
<b>Fraud / transaction monitoring</b>	Payments, retail bank	SR 11-7 / SS1/23	Conformal anomaly bounds, false-positive tracking
<b>Equity research / sell-side notes</b>	Asset manager, sell-side	Disclosure · market abuse	Source attribution, content classifier, publication gate
<b>Code generation for trading</b>	Quant, market risk	Model risk	Static analysis, regression evaluation
<b>Internal knowledge agent</b>	Bank-wide	Conduct · IP leakage	Access control, DLP gates, query auditing

## 04 SAMPLE CURRICULUM · 12-WEEK FS BOOTCAMP

# Module by module. Tuned to your stack in scoping.

Six modules across the twelve weeks. The actual curriculum is tuned to your stack and use cases in the pre-bootcamp scoping call with Charafeddine.

MODULE	TOPIC	WHAT YOUR TEAM PRODUCES
01	<b>Scoping &amp; operating discipline.</b> Process First. The Thinking Stack. LUMEN scoping brief drafted live for one real FS workflow.	Mission charter and LUMEN scoping brief, <b>signed off by sponsor.</b>
02	<b>AI Engineering Foundations.</b> The Three V's (Vibes, Variance, Vendor). Accountable Development Lifecycle. Prompt architecture, multi-model patterns, MCP tool use.	Working <b>system v0.1</b> against your stack and your data.
03	<b>Trust &amp; Verification.</b> The Confidence Problem. Self-consistency sampling. Conformal prediction. TrustGate integration.	Verification gates wired in. <b>Statistical guarantees</b> on output with documented bounds.
04	<b>Accountable Agents · 4-Layer Architecture.</b> Platform Protocol. Agent-Auth (least privilege). Guardrails as architectural constraints. Agent-Monitor.	Agent-Auth and Guardrails layered in. <b>Least-privilege patterns</b> documented for model-risk review.
05	<b>Model Risk Governance.</b> SR 11-7, SS1/23, DORA mapping. AI Act Articles 9, 10, 12, 14, 15. ISO 42001 clauses 6-10. Immutable audit-trail design.	Governance log, <b>regulatory mapping by section number</b> , audit-trail design defensible to model-risk function.
06	<b>Capstone &amp; Sponsor Brief.</b> Red-team workshop. Founder critique. Sponsor demo to leadership. Operating brief written, reviewed, signed.	Production-grade FS system. Sponsor brief that <b>survives model-risk review without amendments.</b>

05 THE FS PORTFOLIO

# Same operating layer. Vertical-specific tuning.

Four programs. One entry-level, two depth options, one transformational. Pricing fixed; scope negotiable. Vertical-specific tuning bends the exercises and the regulatory annex.

**ENTRY · DE-RISKER**

## FS Pilot

€8K-€12K · 4 weeks · 2-3 people

One workflow, taken to proof: a working, verified prototype, with the LUMEN brief and verification design behind it. Half-fee refund clause. 100% credit toward follow-on within 90 days.

**SINGLE-TEAM DEPTH**

## FS Team Bootcamp

€4,200 / seat · 12 weeks · 6-15 seats

Private cohort. Your team, your systems. Pre-bootcamp scoping with Charafeddine. 12 months Engine Room per graduate. Capstone ships at week 12. Sample curriculum above.

**FIRM-WIDE BREADTH**

## FS Curriculum License

€12,000 / year · up to 25 seats

Five role-based learning paths. Completion dashboards. SSO. Quarterly Q&A with Charafeddine. 12 months Engine Room access for all seats. Beyond 25 seats: scoped per engagement.

**TRANSFORMATIONAL**

## FS AI Readiness Program

from €35,000 · 3 to 6 months

Two-day on-site assessment. Training. Monthly 2h founder advisory. Written AI Operating Model playbook at close. The deepest commitment, scoped per engagement.

**Funding (honest).** Cohorte is not Qualiopi-certified today (2026 roadmap, in progress). Direct CPF and direct OPCO not available right now. Standard path: direct L&D budget. For OPCO routing: via Qualiopi-certified partner. Email [teams@cohorte.co](mailto:teams@cohorte.co).

## 06 FS OBJECTIONS WE ALWAYS HEAR

# The honest answers. Before the call.

The four objections an enterprise FS buyer raises in every conversation. Answered here so the call can spend its time on the work.

**"Our model-risk team has not approved any LLM-based tool yet."**

Good. Cohorte's curriculum is what you give them so they can. The methodology is the technical layer underneath SR 11-7 §III independent-validation expectations. The Bootcamp produces a governance log and regulatory mapping cited by section number — designed to survive model-risk review without amendments.

**"Procurement / vendor risk has a 60-day onboarding."**

The Pilot is the way around this: fixed-price, fixed-scope, often clears procurement under a faster threshold. We have run Pilots while full vendor onboarding is in progress.

**"Why a 6-person company, not a Big-Four consultancy?"**

Big-Four sells a strategy deck. Cohorte sells a trained team and a shipped artifact. Several clients use both. We are the layer that comes after the deck. Named reference: Patrick Monteiro, CIO of PwC France & Maghreb. Reference call arranged after a mutual NDA.

**"AI training is not in this year's budget."**

Cohorte is not Qualiopi-certified today (2026 roadmap, in progress), so direct CPF and direct OPCO are not available right now. Standard path is direct L&D budget; OPCO-routed engagements go via a Qualiopi-certified partner. The Pilot at €8K-€12K typically fits inside discretionary spend mid-year.

**DIRECT LINE · NO SALES GATE**

## Talk to Charafeddine directly.

Sixty-minute discovery. No deck. We help you scope the right entry point: Pilot to de-risk, Bootcamp to install, License to scale, AI Readiness to transform.

**[charafeddine@cohorte.co](mailto:charafeddine@cohorte.co)**

---

**Cohorte SAS** · Paris & Rabat · Founded 2022  
[teams.cohorte.co](https://teams.cohorte.co) · [cohorte.co/teams](https://cohorte.co/teams)