TAZI

EXECUTIVE BRIEF · MODEL RISK & AI GOVERNANCE · 2026

Prove Your AI Models and Agents Are Right, on Every Run

> SR 26-2 (April 2026) moves model risk management from calendar-based to risk-based oversight, and leaves generative and agentic AI to each institution's own framework. The second line now has to show, model by model and agent by agent, that decisions are right before they reach customers. TAZI gives the second line independent validation, ongoing monitoring and effective challenge for both, read-only, with a source for every number.

WHY OVERSIGHT IS CHANGING — AND THE TAZI ANSWER

Annual validation is no longer enough

SR 26-2 replaces "at least annually" with risk-based timing that varies with model purpose and changes. Monitoring has to keep up with every run.

AI agents sit outside the guidance

Generative and agentic AI are "not within the scope" of SR 26-2. The institution must review agents under its own GenAI framework, kept separate.

Independence must be shown

"The quality of validation ... depends on the rigor and effectiveness of the review", not on an org chart. Challenge has to be reproducible.

The TAZI answer: independent, reproducible second-line evidence

An independent explanation model re-learns the production prediction as named patterns, each measured against actual outcomes, with SHAP as corroboration. Drift, performance and fairness are checked on every run against thresholds. The Action Simulator reviews an AI agent's rules, outcomes and guardrails before customers are affected. A conversational agent answers the second line's questions in plain language on the first line's own outputs, without the developers. The second line reads; it never acts for the model.

TWO TASKS, KEPT SEPARATE: PROVEN ON ACCOUNT-TAKEOVER FRAUD

Live demo data: 350,167 scored transactions and 1,780 AI agent decisions.

Task 1: Predictive Model

REVIEWED UNDER SR 26-2

65× more fraud than random review in the top 1%

87.2% recall of actual fraud (3,022 of 3,467)

0.002 score PSI, March vs April: no material drift

Task 2: AI Agent

REVIEWED UNDER THE GENAI FRAMEWORK

179 legitimate customers found frozen by the agent's holds

71% of agent actions rest on a citable pattern

100% of 233 SIU proposals meet the documented gate

Committee Evidence

MATERIALITY · FINDINGS · OWNERS

7.2% of records with a data-quality failure the model's own metrics missed

8 of 20 explanation patterns citable as reasons

2 memos model and agent, with conditions

OUTCOMES

Ranking sound, precision falling

The top 10% of the queue holds 90% of frauds. April precision fell to 6.2% (19.6% in February) as fraud prevalence fell: re-test the threshold.

DATA QUALITY

A feed the metrics missed

25,072 of 350,167 records carry a placeholder amount. Compensating control: manual review override until the feed is fixed.

AGENT RELIABILITY

Confidence that does not discriminate

Stated probability 0.95 on 1,769 decisions; 3 of 15 reasoning texts contradict the record. Output reliability rated Red.

THE SECOND-LINE DIFFERENCE

6.4

False alerts per fraud caught at the 0.50 cut, located by score band and channel: the 0.50–0.70 band carries 10,517 false alerts for 132 frauds.

38% vs 9.4%

Agent actions that go against their own explanation pattern hit legitimate customers at 38%, versus 9.4% elsewhere. A human now confirms them.

© 2026, TAZI AI Systems Inc. All rights reserved. Do not copy or share without TAZI's approval.

ASK ANY MODEL OR AGENT IN PLAIN LANGUAGE

Six read-only questions, answered on the first line's own outputs; every number carries its source, scope and check.

Outcomes analysis

"How does the model perform against actual fraud, and what do false alerts cost?" — back-tested against outcomes and a random-review baseline.

Ongoing monitoring

"Any material drift, or data-quality exceptions?" — input and score drift (PSI), missing rates, thresholds on every run.

Conceptual soundness

"Which rules explain the predictions, and are they sound?" — named patterns, then SHAP; protected attributes and proxies checked.

Unintended consequences

"Where do false alerts fall on legitimate customers?" — by score band, segment and channel.

Effective challenge

"Why was this one transaction flagged?" — pattern name, then the drivers and each boundary.

AI agent oversight

"Does the agent follow its policy and the human gates?" — policy adherence, harm, pattern grounding, fairness with significance tests.

PATENTED & INDUSTRY PROVEN TECHNOLOGY

7 granted patents on continuously learning, human-understandable online machine learning · 33 Gartner mentions, 2019–2025, including Cool Vendor, Magic Quadrant and Hype Cycle.

THREE AGENTS ACROSS THE MODEL LIFECYCLE

01 · AT GO-LIVE

Solution Handoff Agent

Model performance, fairness verdict and data-readiness score. Supports pre-deployment validation and sign-off.

02 · IN PRODUCTION

Monitoring Agent

Health score, drift, live performance and fairness on each run. Supports ongoing monitoring and periodic review.

03 · ON DEMAND

Conversational Agent

Plain-language questions to any model or agent. Supports effective challenge without the developers.

ENTERPRISE READY, ON YOUR TERMS

Read-only access for the second line, which can only create its own solutions · role-based access control · drift monitoring and activity logs on every run · self-documented models · bring your own Python & R models · connects to SQL Server, Oracle, Snowflake, PostgreSQL, Databricks, BigQuery, Amazon S3, Azure Blob, Kafka and REST APIs · bring your own API keys and LLMs (BYOK/BYOM) · MCP connector for agents, no-code UI for humans · Slack, Teams, SMTP alerts · single server or Kubernetes · cloud, on-premise or hybrid.

ONE PLATFORM, THREE PATHS TO OVERSIGHT

The same explanations, monitoring and evidence, whether a model is built in-house, by a vendor or by the risk team.

1

Bring Your Own Models

Oversee the models already in use. Scored outputs or Python pickle code are brought in; production is not replaced.

2

Review Vendor Agents

Vendor agents are reviewed on equal terms: the same explanations and monitoring, challenged with the same evidence as in-house models.

3

Build Risk-Team Models

Challengers, explanation models of agent decisions, early-warning, conduct and customer-harm models, fairness segmentation.

★ The second line's own toolkit

Low barrier to entry

Reads existing model and agent outputs, read-only. Nothing is rebuilt; the production model stays.

Evidence the committee can use

Each memo states materiality, findings, owners, compensating controls and limitations, re-tested on each run. The committee decides.

Two tasks, kept separate

Predictive models under SR 26-2; AI agents under your own GenAI framework, each with its own memo and conditions.

EASY TO TRY

1 See the two demos

A 30-minute walkthrough: the six second-line questions on a predictive ATO fraud model, then policy adherence, customer harm and pattern grounding for an AI agent.

2 Start with a Proof of Value

Pick one model or agent in your inventory or waiting for validation. Share anonymized scored outputs and outcomes; receive the second-line answers and a committee memo in a scoped, time-boxed engagement.

TAZI supplies evidence for the institution's model risk management and GenAI governance; it supports, and does not replace, the institution's own validation. Figures: TAZI ATO demo, live data, pulled 2026-10-23.

TAZI

Contact: info@tazi.ai
San Francisco (HQ) · Istanbul (R&D) · Charlotte (East Coast) · Boston (Northeast)

© 2026, TAZI AI Systems Inc. All rights reserved. Do not copy or share without TAZI's approval.