

Grounded AI with Validated Actions

Everyone is shipping agents. Almost no one can put them in production. For regulated, high-stakes decisions, an agent that is not grounded in your data, cannot explain itself, and acts without validated actions is a liability, not a product. The bottleneck isn't capability — it's trust: enterprises deploy only the agents they can explain, audit and govern. That gate is exactly where TAZI wins.

WHY AGENT PILOTS STALL — AND THE TAZI ANSWER

Ungrounded agents hallucinate

LLM-only agents guess. In finance, a wrong-but-confident, inconsistent answer is unshippable.

Automation tools skip the decision

Process-automation agents automate the doing, not the judgment the doing depends on.

No audit trail = no deployment

Compliance, drift and human oversight are afterthoughts, so pilots never reach production.

The TAZI answer: grounded agents

Every decision sits on your data and explainable machine learning. Actions are proposed by agents from your dynamic population of data, and validated by expert-panel and focus-group agents against your organization, business KPIs and focus population. Your team sees how models and agents work through easy-to-use conversational agents and web interfaces — with questions tailored to each specific problem you need to solve.

ONE PLATFORM, THREE PATHS TO PRODUCTION

TAZI is a platform, not a point solution: it ships with pre-built use cases, and your teams build their own on the same governed foundation.

1

Bring Your Own Models

Onboard the models you already have. TAZI wraps them with explanations, drift monitoring, validation and auto-generated compliance documentation — compliant and production-ready.

2

Deploy Pre-Built Use Cases

Out-of-the-box explainable agents — fraud detection, retention & growth, voice of customer — grounded in your data and live in days, not quarters.

3

Build New Use Cases

Create models from scratch for use cases that don't exist yet. Define the risk, build, validate and deploy — your experts in the loop at every gate.

Low barrier to entry

One use case at a time, data cleaned and improved as we go. Agents can fill in and identify data problems. A digestible process.

Your team's life gets easier

Agents answer questions in plain language, every alert arrives with a reason and a validated next step, and compliance documentation writes itself.

The market advantage is built in

TAZI's field-proven explainable-ML, documentation, production, multi-agent IP is already there. You deploy it; competitors have to build it.

PRE-BUILT GROUNDED AGENTS: DEPLOY IN DAYS

Grounded in your data, coached by your people, documented by default.

Fraud Detection Agent

ACCOUNT TAKEOVER (ATO) · AML

+37% investigation efficiency

\$480K more fraud caught / year

\$1M → \$235K false-negative cost

Growth Agent

BANK RETAIN & GROW · WEALTH RETAIN & GROW

15% attrition reduction

20% save rate · \$8M NPV

32× ROI (up to) in the first year

Voice of Customer Agent

COMPLAINTS · VOC · CALL-CENTER QA

100% severity accuracy

1 min to classify a month of complaints

42 Level-3 cases recovered from Level 1

THE VALIDATED-ACTION DIFFERENCE

The same predictions, with the Expert-Panel + Focus-Group validation layer on top:

94% → 42%

Needless manager alerts in wealth retention; 27% of the book cleared as “no action”; corrupted-data cases caught with zero false alerts.

20 → 0

False SAR exposures on legitimate customers in the top ATO tier: full action ladder, 100% auditable rule trail.

98 → 4

Top-tier AML SAR escalations (96% fewer, all corroborated); “source of funds” tipping-off phrasing removed 99% → 0.4%.

SPOTLIGHT: ATO FRAUD AGENT

Explainable ML on session, device, transaction and payee-velocity data, retrained on investigator feedback, scores account-takeover risk per session and transaction — with a reason code and validated actions on every alert.

+37% / 53%

Efficiency gain (test / hold-out); \$480K more fraud caught per year, \$2.5M on hold-out.

50% → ~80%

True positives; false-negative cost down from \$1M to \$235K.

~72 POV hours

From prepared data to a compliant, A/B-tested, production-ready model.

ENTERPRISE READY, ON YOUR TERMS

Start small, together. Every engagement opens with a joint discovery phase — we scope the first use case with you, and each win funds the next bite of the apple.

Explainable and validated by design · continuous learning keeps models current without full retrains · drift monitoring on every deployed model · auditable activity logs and auto-generated model cards and compliance documentation · cloud, on-prem or hybrid · bring your own keys and LLMs · multitenant across business units · real-time via a single REST API.



T A Z I

Contact: info@tazi.ai

San Francisco (HQ) · Istanbul (R&D) · Charlotte (East Coast) · Boston (Northeast)