JUPITER
CAPITAL MANAGEMENT

EXPERT ADVISORY & CONSULTING SERVICES

Digital Banking - AI Risk

AI

Effective AI risk management in digital banking is critical to protect customers, maintain regulatory compliance, and preserve financial stability and reputation; without it, models can produce biased credit decisions, miss fraud, expose sensitive data, or fail under stress, causing financial loss and legal exposure. Proactive controls — data governance, explainability, access controls, continuous monitoring, and incident response — reduce operational and regulatory risk while enabling safer innovation. By embedding risk management into the lifecycle of AI systems, banks can deliver better customer outcomes, sustain trust, and unlock the efficiency and personalization benefits of AI without creating systemic vulnerabilities.

RISK ASSESSMENT & GOVERNANCE

Inventory: catalogue AI uses (credit scoring, transaction monitoring, chatbots, personalization, underwriting, robo-advice), datasets, third-party models, cloud providers, and critical infra.
Threat modeling: list failure modes per use case (false credit denials, missed fraud, chatbot misinformation, biased pricing, model extraction).
Prioritization: rank by monetary loss, regulatory exposure (AML, fair lending), customer harm, and systemic concentration risk.
Risk register & SLAs: capture mitigations, owners, review cadence (monthly for high risk, quarterly for medium), performance SLAs, rollback criteria.
Committee & policy: risk/governance committee with risk, compliance, legal, tech, product, ops; policy on acceptable uses and human-in-the-loop thresholds.

DATA GOVERNANCE

Owners & stewardship: assign data owners for transaction logs, customer profiles, credit bureau feeds, call/chat transcripts.
Lineage, provenance & retention: maintain provenance for KYC/AML datasets; label synthetic data; set retention tied to regulations (logs, explanations, training snapshots).
Quality checks & minimization: automated validation (schema, missingness, anomaly detection); automated PII minimization scanners before training.
Privacy & consent: tokenization/pseudonymization, masking for testers, consent flags for marketing/personalization, consumer disclosure strategy.
Audit & immutable logs: access and transformation logs for audits and SARs; retention policies for evidence.

MODEL TRANSPARENCY, EXPLAINABILITY & DOCUMENTATION

Model cards & docs: purpose, inputs, outputs, limitations, training data snapshot, intended audience, CI/CD gating.
Decision rationale & disclosure: plain-language reasons for customer-impacting decisions and actionable remediation steps; disclosure points and consent where required.
Versioning & provenance logging: log model version, score, decision context, and human overrides for each decision; store evaluation artifacts.

BIAS DETECTION, FAIRNESS & ETHICS

Protected attributes & metrics: identify proxies; monitor group metrics (FPR/FNR, calibration, equal opportunity, outcome impact).
Routine audits & thresholds: scheduled fairness audits post-retrain or on drift; alert on defined disparity thresholds.
Mitigations & oversight: reweighting, reject-inference handling, threshold tuning, mandatory human review for borderline/high-impact cases.
Appeals & remediation: clear customer appeal channels, automated reconsideration workflows, remediation tracking.
Ethics policy & human oversight: explicit policy on unacceptable uses and escalation for sensitive decisions.

www.jupiter-cap.com Email: rod.boothe@jupiter-cap.com

JUPITER
CAPITAL MANAGEMENT

EXPERT ADVISORY & CONSULTING SERVICES

Digital Banking - AI Risk

AI

Effective AI risk management in digital banking is critical to protect customers, maintain regulatory compliance, and preserve financial stability and reputation; without it, models can produce biased credit decisions, miss fraud, expose sensitive data, or fail under stress, causing financial loss and legal exposure. Proactive controls — data governance, explainability, access controls, continuous monitoring, and incident response — reduce operational and regulatory risk while enabling safer innovation. By embedding risk management into the lifecycle of AI systems, banks can deliver better customer outcomes, sustain trust, and unlock the efficiency and personalization benefits of AI without creating systemic vulnerabilities.

SECURITY & ROBUSTNESS

Asset inventory & segregation: map models, endpoints (scoring API), and data stores; separate dev/test/staging from production; use synthetic/anonymized data in non-prod.

Access & secrets management: role-based access, least privilege, MFA, secrets rotation, and approval workflows.

Robustness testing: adversarial, poisoning, model extraction, and prompt-attack simulations; adversary modeling for top use cases.

Runtime monitoring & anomaly detection: monitor score distributions, feature drift, access spikes, and extraction attempts; automated alerts and containment triggers.

Business continuity & DR: RTO/RPO for model endpoints and data stores, fallback/runaway controls and failover plans.

THIRD-PARTY & VENDOR RISK

Due diligence: evaluate vendor models for explainability, data handling, compliance, robustness, and SLAs.

Contracts & rights: audit rights, data residency, liability clauses, model change notification, and exit/fallback terms.

Ongoing monitoring: performance comparison to inhouse baselines, concentration risk monitoring, and fallback plans.

PERFORMANCE MONITORING, DRIFT & MODEL LIFECYCLE

KPIs by use case: credit (approval rate, default rate by cohort), AML (detection rate, false positives), chatbots (escalation rate, CSAT).

Drift detection & retraining: monitor feature/outcome shifts; trigger retraining or investigation thresholds; document provenance for retraining data.

Feedback loops: feed investigator outcomes (fraud confirmations, appeals) into retraining pipelines with audit trails.

Revalidation & retirement cadence: monthly for high-impact models, quarterly for medium, annual for low; defined retirement process.

COMPLIANCE MONITORING & AUDITABILITY

Regulatory mapping: map models to applicable rules (fair lending, GLBA, AML/KYC, GDPR/CCPA) and create control checklists.

Evidence & retention: store training data snapshots, evaluation reports, model cards, and audit logs for regulatory requests.

Automated pre-deployment checks: data lineage, PII scanning, fairness thresholds, and legal/compliance gates in CI/CD.

Reporting templates: prebuilt regulator response templates and SAR/notice processes.

www.jupiter-cap.com Email: rod.boothe@jupiter-cap.com

JUPITER
CAPITAL MANAGEMENT

EXPERT ADVISORY & CONSULTING SERVICES

Digital Banking - AI Risk

AI

Effective AI risk management in digital banking is critical to protect customers, maintain regulatory compliance, and preserve financial stability and reputation; without it, models can produce biased credit decisions, miss fraud, expose sensitive data, or fail under stress, causing financial loss and legal exposure. Proactive controls — data governance, explainability, access controls, continuous monitoring, and incident response — reduce operational and regulatory risk while enabling safer innovation. By embedding risk management into the lifecycle of AI systems, banks can deliver better customer outcomes, sustain trust, and unlock the efficiency and personalization benefits of AI without creating systemic vulnerabilities.

STAKEHOLDER ENGAGEMENT & COMMUNICATION

Audience mapping: regulators, customers, internal teams (operations, product, risk, legal), and partners.
Clear external communication: short model summaries for customer-facing models and disclosure strategy.
Internal feedback loops: channels for frontline and customer complaints; escalate systemic issues to model owners and committee.
External review: periodic third-party audits for high-risk models and advisory input for sensitive systems.

INCIDENT RESPONSE & RECOVERY

Incident taxonomy & severity: define incidents (wrongful denial, major drift, data leak, adversarial attack) and severity levels.
Cross-functional team & roles: include fraud ops, legal, risk, IT, communications, product, and branch reps.
Runbook: detect → contain (rollback/throttle/fallback) → investigate → remediate → notify regulators/customers → post-mortem.
Drills & playbooks: tabletop exercises at least annually; include scenarios (AML evasion, major credit scoring error).

TRAINING, AWARENESS & CHANGE CONTROLS

Role-based training: short modules for executives, model owners, ops/frontline, devs, and vendors.
Hands-on simulations: credit appeals, fraud triage, chatbot escalation practice.
Change controls & approvals: CI/CD gates, change logs, mandatory signoffs for prod changes.
Metrics: track completions, competency checks, and tie to access privileges.

COST, CAPACITY & OPERATIONAL RESILIENCE

Cost governance: monitor cloud and inference costs, tagging, and budget alerts.
Capacity planning: load testing, performance SLAs, and throttling rules for spikes.
Operational metrics: MTTR, uptime, and response time SLAs for model endpoints.

www.jupiter-cap.com Email: rod.boothe@jupiter-cap.com