CedarBridge Financial

03 | AI/Model Risk Classification Methodology

Financial-services AI governance engagement

Engagement type	Simulated portfolio case study
Organization	CedarBridge Financial, fictional mid-sized consumer lender and digital bank
Scope	Enterprise AI and model governance across credit, fraud, customer service, marketing, and operations
Primary scenario	CreditPath consumer credit underwriting model and ClearSignal fraud-AI vendor

This artifact uses fictional systems and synthetic data. It is designed to show governance judgment, model-risk discipline, consumer-protection awareness, and operational evidence rather than to represent any real company or vendor.

Purpose

This methodology provides a repeatable way to classify CedarBridge AI systems and models as prohibited, critical, high, moderate, or low risk. The classification determines governance depth, validation requirements, approval authority, and monitoring cadence.

Risk-tier definitions

Prohibited	Use is inconsistent with law, policy, consumer protection, or CedarBridge risk appetite.	Do not deploy; disable, investigate, and remediate.
Critical	Direct or material influence over consumer eligibility, pricing, account restrictions, regulatory reporting, capital, AML, or financial risk.	Formal assessment, independent validation, committee approval, enhanced monitoring, board visibility.
High	Material customer, compliance, operational, or financial impact but not critical enterprise exposure.	Assessment, second-line review, monitoring, documented control owner.
Moderate	Supports decisions or interactions but human users verify outputs before action.	Privacy/security review, acceptable-use limits, periodic review.
Low	Internal support with no regulated data, customer decisioning, or external effect.	Approved tool controls and training.

Scoring factors

Consumer impact	Internal only	Indirect support	Material customer treatment	Eligibility/pricing/account restriction
Decision influence	Drafting only	Advisory	Strong recommendation	Automated or near-automatic decision
Data sensitivity	Public/internal	Business confidential	Consumer data/NPI	Sensitive data or high aggregation
Explainability	Clear rules	Interpretable model	Complex but explainable	Black-box or weak reasons
Regulatory exposure	Low	Operational risk	Consumer compliance	Credit/fair lending/AML/regulatory reporting
Scale	Small pilot	Team use	Business unit	Enterprise or high volume
Vendor dependency	None	Low	Material vendor	Critical third-party dependency
Recourse availability	Not needed	Internal correction	Customer support path	Formal customer reconsideration required

Classification rules

- Any prohibited-use trigger overrides the score and blocks deployment.
- Any model that directly influences credit approval, pricing, credit limit, adverse-action reasons, AML alerts, or account restrictions is at least High and may be Critical.
- A system using nonpublic financial information in a customer-facing process is at least Moderate.
- Any vendor model with limited transparency and material customer impact requires enhanced third-party review before approval.
- Risk tier must be reassessed after material changes, new data sources, expanded populations, new jurisdictions, or drift/incident events.

Prohibited-use screen

Credit decision cannot generate specific and accurate adverse-action reasons	Opaque decline model with generic "does not meet model standards" output	Reject or redesign before use
Protected-class inference or prohibited proxy use	Model infers race, religion, disability, or immigration status for decisioning	Disable and escalate
Deceptive customer-facing AI	Chatbot implies it is a human banker or hides AI use in regulated servicing interaction	Disallow until disclosure and controls are implemented
Unapproved public GenAI with consumer data	Staff pastes loan file or complaint into public chatbot	Incident response and training remediation

Required governance by tier

Inventory record	Yes	Yes	Yes	Yes
AI/model risk assessment	Optional	Targeted	Required	Required
Independent validation	No	As needed	Required for material models	Required before production and annually
Fair-lending/consumer review	No	As needed	Required if consumer impact	Required
Executive risk approval	No	No	Sometimes	Required
Board visibility	No	No	Exception-based	Semiannual and incident-based
Monitoring frequency	Annual	Quarterly/Semiannual	Monthly/Quarterly	Monthly or more frequent

Reference basis

This simulated package is informed by the 2026 Revised Guidance on Model Risk Management issued by the Federal Reserve, OCC, and FDIC, NIST AI RMF 1.0, the CFPB circular on adverse-action notices for complex credit algorithms, CFPB guidance on specificity of adverse-action reasons, and the 2023 interagency guidance on third-party relationships. These references are used as design inputs for a portfolio artifact, not as legal advice.