CedarBridge Financial

05 | Model Risk Assessment / Model Card: CreditPath

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.

Model overview

Model name	CreditPath v3.2
Owner	Chief Credit Officer / Consumer Lending Analytics
Use case	Consumer installment-loan underwriting, pricing band recommendation, credit limit recommendation, and adverse-action reason generation.
Decision influence	Critical. The model produces a recommended approve/decline decision and pricing tier; underwriters review exceptions and policy-edge cases.
Model type	Gradient-boosted decision trees with calibrated probability of default and affordability overlay.
Users	Credit operations analysts, automated loan platform, fair-lending monitoring team.
Population	U.S. consumer applicants for unsecured installment loans.
Status	Conditional approval only. Production blocked until listed remediation is complete.

Intended use and business justification

CreditPath is intended to improve risk segmentation, reduce manual review time, support consistent underwriting, and align pricing to demonstrated repayment risk. The business justification is accepted only if CedarBridge can maintain accurate adverse-action reasons, fair-lending monitoring, model transparency sufficient for governance, and meaningful override/reconsideration controls.

Data and feature risk

Credit bureau variables	Delinquencies, utilization, inquiries, age of credit file	Predictive and commonly used, but can create disparate outcomes and thin-file impacts.	Permit with fair-lending monitoring and reason-code mapping.
Income and affordability	Verified income, debt-to-income, deposit cash-flow proxy	Useful for ability-to-repay and affordability; missingness may disadvantage gig workers.	Permit with missingness treatment and alternative documentation path.
Bank transaction data	Deposit volatility, overdraft frequency, inflow regularity	Potentially useful but can encode income instability and protected-class proxies.	Restrict; require validation and fairness sensitivity testing.
Geolocation / ZIP-level signals	Neighborhood or region indicators	High proxy risk for protected-class and redlining concerns.	Prohibit for decisioning unless specifically justified and legally approved.
Device or behavioral data	Typing speed, device, session patterns	Weak job-relevance equivalent for credit, high privacy/proxy concerns.	Prohibit for underwriting.

Validation summary

Conceptual soundness	Generally acceptable	Business-purpose document needs stronger explanation of affordability overlay and reject-inference limitations.	Remediate before full approval.
Data quality	Partially acceptable	Income verification missingness higher for nontraditional employment types.	Alternative documentation control required.
Performance	Acceptable in aggregate	AUC 0.78; calibration stable in prime band, weaker in thin-file segment.	Segment monitoring required.
Fair-lending/outcomes	Concern identified	Synthetic test shows higher decline rate for proxy-vulnerable group; requires further fair-lending review.	Production blocked pending remediation.
Explainability	Partially acceptable	Top-driver reason codes exist but reason accuracy QA not yet operational.	Adverse-action QA required.
Implementation	Not complete	Monitoring feeds and rollback playbook not yet tested.	No launch until completed.

Synthetic fair-lending and adverse-action test

Overall	54%	-	Baseline synthetic portfolio.
Group A	60%	1.00	Reference in this test.
Group B	52%	0.87	Monitor.
Group C	45%	0.75	Triggers investigation; not a legal conclusion by itself.
Thin-file applicants	38%	0.63	Requires alternative evidence path and segment validation.

Adverse-action reason controls

- Every declined application must receive specific and accurate principal reasons traceable to validated model drivers.
- Reason codes must be tested monthly against model inputs and reviewer notes.
- CreditPath cannot use features that produce reasons CedarBridge cannot explain to a consumer.
- Human overrides must not remove or replace accurate adverse-action reasons without documented justification.
- Complaint patterns tied to reason clarity are monitored as a KRI.

Human overrides and reconsideration

Underwriters may approve exceptions, send the file to manual review, or request alternative documentation when the model output conflicts with documented credit policy, accommodation needs, suspected data error, thin-file status, or customer-provided evidence. Decline recommendations cannot be finalized without the adverse-action reason engine, policy-rule confirmation, and an automated check for manual-review triggers.

Residual risk and approval decision

Fair-lending and proxy risk	High	Production blocked until alternative configuration and fair-lending review are complete.
Explainability / adverse-action reasons	High	Reason-code QA and complaint review must be operational before launch.
Model performance	Moderate	Acceptable with segment monitoring and calibration thresholds.
Operational resilience	Moderate	Rollback and manual fallback test required.
Vendor data dependency	Moderate	Data-feed SLA and incident notice must be executed.

Governance decision

Conditional approval for controlled parallel run only. CreditPath may not be used for final automated decline or pricing decisions until remediation items are closed, second-line validation signs off, Compliance approves adverse-action reason procedures, and the Executive Risk Committee accepts residual risk.

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.