BankAtlas

The AUC-Validated Bank Risk Model

Methodology and Evidence — A Plain-Language Overview

Executive Summary

BankAtlas turns the full U.S. banking call-report record into a validated, queryable risk model. At its core are two complementary lenses on every bank, every quarter:

- **A composite grade (A–E)** that summarizes a bank's overall risk across six dimensions — asset quality, earnings, capital, funding and liquidity, concentration, and an overall composite. Crucially, the grade is weighted by what has actually predicted bank distress in the historical record, not by analyst judgment.
- **A Run-Vulnerability Index (RVI)** that isolates funding and liquidity fragility — the kind of run risk a healthy-looking composite grade can miss. A bank can carry a clean grade and still light up on the RVI.

This document explains, in plain language, how the grade is built and validated, and shows the historical evidence that lower grades are followed by materially higher failure rates. It is an overview intended for prospective clients; it describes what the model measures and how it is validated, not the proprietary weighting and construction details.

The Two-Lens Framework

Most bank-risk tools give you a single read. BankAtlas deliberately separates two questions that behave very differently in a crisis:

- **"Is this bank fundamentally sound?"** — answered by the composite grade, which integrates capital, asset quality, earnings, funding, and concentration into one event-validated score.
- **"Can this bank survive a bad week?"** — answered by the Run-Vulnerability Index, which focuses narrowly on funding structure, deposit reliance, and liquidity buffers.

The two are designed to disagree. A bank can post a solid composite grade while its funding structure quietly deteriorates — exactly the pattern that surprised observers in 2023. Used together, the grade catches broad fundamental weakness and the RVI catches funding fragility, so the platform flags both kinds of risk rather than averaging them into a single number that hides the more acute one. As a practical rule, a sustained high RVI warrants attention even when the composite grade looks reassuring.

How the Grade Is Validated

The defining feature of the BankAtlas composite is that it is empirically validated against real outcomes. In plain terms:

- **Metrics earn their weight.** Each candidate risk metric is tested for how well it actually separates banks that later experienced distress from those that did not, across the historical record. Metrics with stronger demonstrated power to anticipate distress carry more weight in the composite; weak or noisy metrics carry little or none.

- **Comparisons are peer-relative.** Every bank is scored against peers of similar size, so a community bank is judged against community banks — not money-center institutions. Size cohorts are defined consistently across time (under $1B, $1–10B, $10–100B, and over $100B in assets).
- **Grades reflect standing within cohort.** The composite is converted into a within-cohort percentile and mapped to an A–E letter grade, with A representing the strongest banks in the cohort and E the weakest.

The result is a grade that is reproducible, refreshed quarterly, and — most importantly — anchored to what has historically preceded bank distress, rather than to opinion. The specific weighting and construction methodology is proprietary.

The Evidence: Lower Grades Fail More

If the grade is meaningful, weaker grades should be followed by higher subsequent failure and enforcement rates. They are — strongly and consistently. The table below summarizes outcomes across roughly 597,000 rated bank-quarters from 2003 through early 2026.

Grade	12-mo failure	24-mo failure	12-mo enforcement
A	0.03%	0.08%	0.16%
B	0.04%	0.11%	0.16%
C	0.05%	0.22%	0.27%
D	0.16%	0.56%	0.73%
E	**2.85%**	**4.82%**	2.62%

Within twelve months, E-rated banks failed at roughly 86 times the rate of A- and B-rated banks. The separation is monotonic — each step down the grade scale carries higher subsequent risk — and it persists across both 12- and 24-month horizons.

This is the core claim of the model, and it is the single most important thing to understand about a BankAtlas grade: it is not a label, it is a calibrated probability statement grounded in two decades of outcomes.

Verified Against the 2023 Regional-Bank Crisis

The 2023 regional-bank failures are a natural test of the two-lens framework. Both lenses were examined as of the quarters preceding each receivership, using the same model that runs today.

First Republic Bank — caught by the grade

- **Composite grade:** E in 2021, deteriorating to D and holding there through its May 2023 receivership — a bottom-tier grade for roughly two years ahead of failure.
- **Run-Vulnerability Index:** elevated through 2022 and climbing to the 100th percentile of its peer cohort by the first quarter of 2023.

First Republic illustrates the grade working as intended: a persistently weak composite flagged the bank well before its funding situation became acute, and the RVI then confirmed the late-stage liquidity stress.

Silicon Valley Bank — caught by the second lens

- **Composite grade:** B — a healthy-looking grade right up to its March 2023 failure. On the composite alone, SVB did not appear to be in distress.

- **Run-Vulnerability Index:** elevated — in the upper range of its peer cohort — signalling funding fragility that the headline grade did not capture.

SVB is the case the two-lens framework is built for. The composite grade said the bank looked sound; the Run-Vulnerability Index disagreed, pointing to the funding and liquidity structure that ultimately drove the failure. A single blended score would have averaged these signals together and muted the warning.

Taken together, the two cases make the argument for two lenses: First Republic was caught by the grade, SVB only by the Run-Vulnerability Index. Monitoring both flags both kinds of fragility.

Governance and Refresh

- **Quarterly refresh.** Grades and the Run-Vulnerability Index update each quarter as new call-report and UBPR data is filed.
- **Coverage-aware.** Where a bank's data is too sparse to support a reliable grade, the grade is withheld rather than reported with false precision.
- **Peer-stable.** Size cohorts and peer comparisons are held consistent across time so that grades are comparable from quarter to quarter and bank to bank.

What the Model Does — and Does Not — Claim

- **It is an early-warning lens, not a guarantee.** Lower grades carry materially higher historical failure rates, but most low-graded banks do not fail; the grade ranks relative risk, it does not predict any individual outcome with certainty.
- **Percentiles are relative.** Grades reflect standing within a peer cohort, so a broadly weakening cohort can mask absolute deterioration. The metrics behind the grade remain available for that reason.
- **Relationships can drift.** The link between metrics and outcomes is periodically re-estimated as banking conditions evolve, so the model is maintained rather than fixed.
