

The judgements auditors test in an expected credit loss model

Expected credit loss is a framework of judgements that happens to end in a calculation. The models that struggle in audit tend to share the same three weaknesses.

June 10, 2026 · 11 min read

soniassociates.in/insight-ecl-models-audit-scrutiny

AT A GLANCE

- Expected credit loss failures in audit come back to three causes: segmentation that ignores how customers actually differ, forward-looking adjustments asserted rather than derived, and judgements exercised but never written down.
- For companies on the simplified approach, a provision matrix is only as good as its ageing data and its segments. One blended loss rate across very different customers is the most common shortcut, and the easiest to spot.
- For lenders and NBFCs, discipline over staging matters more than model sophistication. Auditors test the boundary between Stage 1 and Stage 2 harder than they test the probability of default.
- Management overlays are legitimate, and the standard allows for them. An overlay without a quantified basis and a condition for release is a warning sign, not a cushion.

IFRS 9, and Ind AS 109 in India, replaced the incurred-loss model, in which provisions waited for evidence of default, with an expected-loss model, in which every receivable carries a probability-weighted estimate of future loss from the day it is recognised. That shift moved credit provisioning from observation to prediction. Prediction is judgement, however many decimal places the spreadsheet shows.

Having designed and reviewed these models, including one for a listed IT services company with a receivables book approaching a billion dollars, we see a consistent pattern. Models rarely fail because the mathematics is wrong. They fail because the judgements beneath the mathematics were never made explicit, never supported, or never revisited, and auditors, along with the regulators who review auditors, have learned exactly where to look.

Segmentation: the judgement made before any number exists

Every expected credit loss model starts from an implicit claim: *these exposures are similar enough to share loss statistics*. That claim is segmentation, and it is where most corporate models are weakest. A provision matrix that applies one set of ageing-bucket loss rates to an entire receivables book is saying that a government department, a multinational, and a two-year-old distributor all behave the same way on credit. No credit manager believes that, and the model should not either.

Defensible segments follow the way risk actually differs in the business, usually through some combination of these:


- **Customer type:** government and public sector, large corporate, SME, and related parties, which need individual assessment rather than a matrix.
- **Geography,** especially where receivables span countries with different enforcement realities and currency risk.
- **Product or contract structure.** Project receivables billed on milestones age differently from recurring subscription billing.
- **Security or credit enhancement.** Advances, bank guarantees, and credit insurance change the loss given default, not the probability of default.

The test we apply is simple. If the credit team describes two groups of customers differently in conversation, the model should describe them differently in numbers. Where history is too thin to calculate rates by segment, the standard allows reasonable and supportable proxies, but the choice of proxy is itself a judgement that needs documenting.

EXHIBIT 1

A single blended loss rate over-provides for one part of the book and under-provides for another

Expected loss rates by ageing bucket for two customer segments in one receivables book, percent



Illustrative figures. As the mix between segments shifts, a blended matrix misstates the provision even when its historical arithmetic is right.

Source: IFRS 9 Financial Instruments, paragraph B5.5.35 on provision matrices (mirrored in Ind AS 109); SA&A analysis

Soni Aditya & Associates

The provision matrix: honest history, then honest adjustment

The mechanics for trade receivables under the simplified approach are well known. Build flow rates by ageing bucket from historical data, derive loss rates, and apply them to the current ageing. What separates a good matrix from a weak one is less well known:

- 1. Integrity of the flow rates.** Loss rates should come from cohort migration, meaning how much of each bucket rolls forward instead of being collected, over enough periods to be stable and with the write-off policy applied consistently. A history that never writes anything off shows zero loss, and proves only that write-offs are being avoided.
- 2. The definition of default.** The standard includes a rebuttable presumption at 90 days past due. Rebutting it is allowed, and many business-to-business books really do collect at 180 days, but the rebuttal needs evidence in the form of actual collection curves, not commercial optimism.
- 3. The forward-looking adjustment.** Historical rates describe the past, and expected credit loss prices the future. The adjustment should link named economic or portfolio indicators, such as interest rates, sector stress, or changes in customer concentration, to a direction and size of adjustment. A policy note saying “management considered forward-looking information and concluded no adjustment is required”, with nothing behind it, is among the weaknesses auditors challenge most often.

Historical loss rates describe the past. Expected credit loss prices the future. What bridges the two is a judgement, and it has to be written down.

For lenders, staging is where the audit happens


In NBFCs and other lenders using the general approach, audit attention gathers at one boundary: the line between Stage 1, with twelve-month expected loss, and Stage 2, with lifetime expected loss. Crossing that line on a significant increase in credit risk can multiply a provision several times over, which makes the criteria that define it the most consequential judgement in the whole model.

EXHIBIT 2

The stage boundaries are where the provision multiplies, and where auditors concentrate

The general approach to impairment under IFRS 9

Deterioration: rebuttable presumptions at 30 days past due for a significant increase in credit risk, and at 90 days for default. Qualitative signs should move exposures sooner.



Cure: exposures move back once the criteria no longer apply, usually after a defined probation period of sustained performance.

Simplified. Trade receivables without a significant financing component are measured under the simplified approach, at lifetime expected credit loss.

Source: IFRS 9 Financial Instruments, section 5.5 and paragraphs 5.5.11 and B5.5.37 (mirrored in Ind AS 109)

Soni Aditya & Associates

Strong frameworks for identifying a significant increase in credit risk have three things in common:

- **Several indicators, with days past due as only one of them.** The 30-day presumption is a backstop, not a definition. Requests to restructure, covenant breaches, a worsening bureau score, and sector watchlists all signal deterioration before the payment record does.

- **Symmetry.** Exposures must be able to cure back to Stage 1 under stated conditions, and the cure criteria, usually a probation period of sustained performance, should be as explicit as the criteria for deterioration.
- **Stability testing.** If small changes in parameters swing large populations between stages, the boundary is arbitrary and the auditor will treat it that way. Sensitivity analysis on staging criteria is cheap insurance.

On the parameters themselves (probability of default, loss given default, and exposure at default), sophistication should follow materiality. A detailed vintage-based term structure suits a large retail book; a well-supported mapping to external ratings may be perfectly adequate for a concentrated corporate book. What is never adequate is a set of parameters whose origin nobody can explain.

Overlays: legitimate, if engineered


Sooner or later every real model meets a situation its data has never seen: a pandemic, a sector shock, a regulatory intervention. Management overlays exist for exactly that, and auditors accept them when they are engineered, with a stated trigger, a quantified basis however approximate, a defined population, and a release condition that says when and how the overlay unwinds. What draws challenge is the opposite: a round number, held for years, drifting between prudence and profit smoothing depending on the quarter. An overlay that cannot say when it will end is not really a provision. It is a reserve, and reserves are what the standard was written to eliminate.

The reconciliation of the loss allowance that IFRS 7 requires is the best single test of whether a model is behaving. When the allowance moves, the reconciliation shows whether it moved because the portfolio grew, because exposures deteriorated or recovered, or because someone changed the model. A board should be able to read it and understand why this year's provision differs from last year's.

EXHIBIT 3

A loss allowance reconciliation shows whether the provision moved for the right reasons

Illustrative movement in a lender's expected credit loss allowance over one year, ₹ crore



Illustrative figures. Stage transfers, cures, and overlay changes are where auditors will ask for the reasoning behind the number.

Source: IFRS 7 Financial Instruments: Disclosures, paragraph 35H; IFRS 9 (mirrored in Ind AS 107 and Ind AS 109); SA&A analysis

Soni Aditya & Associates

The documentation that ends audit debates

The deliverable that separates a smooth audit from a contentious one is a methodology paper: a living document, typically fifteen to twenty-five pages, that sets out every judgement in the chain. It covers the basis for segmentation, where the data comes from, the definition of default, how flow rates are calculated, the forward-looking framework, the staging criteria, how overlays are governed, and the results of validation. Each judgement carries its reasoning and a date for review.

The paper does something less obvious than satisfy auditors. It forces the judgements to be made consciously. Most weak models are weak because their key choices were made without thought by whoever built the first spreadsheet, and have been inherited unexamined ever since. Writing the methodology down turns a spreadsheet into a policy that survives staff turnover, a change of auditor, and the next crisis the data has not seen.

Expected credit loss is where the financial statements show how well a company understands its own receivables. Companies that take that view stop asking what provision the auditor will accept, and start asking what the portfolio is telling them. The audit usually follows.

ABOUT THE AUTHOR

[Aditya Soni](#) is the founding partner of Soni Aditya & Associates and leads its financial reporting and technical accounting work. He is a chartered accountant and was previously in the Assurance practice of Ernst & Young in Pune.

This article reflects the author's understanding of the relevant laws and standards at the date of publication. It is general in nature and is not a substitute for advice on specific facts.

Soni Aditya & Associates, Chartered Accountants, Pune

aditya@soniassociates.in · +91 93400 75288 · [Read online](#)