
STRENGTHENING FINANCIAL INSTITUTIONS

# CECL: The Anatomy of a Defensible Reserve

Released April 2026

## Introduction: Why Some CECL Reserves Are Defensible - and Others Are Not

Most Current Expected Credit Loss (CECL) frameworks satisfy the accounting standard, but many struggle to answer a simpler question: *why did the reserve change?* When the drivers of the allowance are difficult to trace, governance becomes harder, reliance on qualitative adjustments increases, and the narrative around credit risk weakens. A defensible reserve is the outcome of a framework designed to translate evolving credit risk into an explainable estimate of expected loss. This paper outlines four structural pillars that determine whether a CECL implementation produces opaque results or a reserve that can be clearly interpreted, governed, and defended.

Institutions that design their CECL frameworks with this architecture in mind gain a clearer view of how risk develops within their portfolios and a stronger foundation for explaining that risk when it matters most.

**KEY TAKEAWAY**

Defensible CECL reserves come from structure: the right inputs, the right granularity, the right conditional logic, and a model design that preserves the drivers of loss.

**CLIENT-FOCUSED SOLUTIONS**

Since 2003, Wilary Winn has provided independent, objective, fee-based advice to financial institutions and now serves more than 600 clients across the country.

**Our main service lines include:**

> ASSET LIABILITY MANAGEMENT (ALM)

> CURRENT EXPECTED CREDIT LOSS (CECL)

> MERGERS & ACQUISITIONS (M&A)

> VALUATION OF LOAN SERVICING

> FAIR VALUE DETERMINATIONS

## Compliance Without Clarity

Since CECL became effective, most financial institutions have successfully implemented the processes necessary to comply with the accounting standard. Models have been developed, methodologies documented, and allowance calculations incorporated into regular financial reporting. Institutions can demonstrate that their frameworks consider historical performance, incorporate reasonable and supportable forecasts, and apply a defined methodology for estimating expected credit losses. From a compliance perspective, the transition has largely been achieved.

Yet many CECL implementations reveal a deeper issue: while the model produces an allowance figure, the drivers of that figure are not always clear. Reserves move from quarter to quarter for many reasons: portfolio balances grow or contract; borrower performance evolves; macroeconomic forecasts change. These developments should naturally influence the estimate of expected credit loss. The challenge arises when the connection between those developments and the resulting reserve movement becomes difficult to trace.

In some cases, model outputs appear stable even when underlying portfolio characteristics change. In others, reserves shift materially while the explanation remains broad or qualitative. Institutions often rely on overlays or adjustments to reconcile results with expectations, but those adjustments sit outside the structural logic of the model itself. When this occurs, the framework may still satisfy the technical requirements of CECL. However, it lacks analytical transparency.

This distinction is important because CECL ultimately sits within a governance environment. The allowance must be explainable to leadership. Auditors must evaluate whether the methodology is conceptually sound. Regulators must determine whether the framework appropriately reflects the risk embedded in the institution's loan portfolio. Across these contexts, one question consistently emerges: *why did the reserve change?*

A defensible CECL framework should allow that question to be answered clearly and consistently. Reserve movement should be traceable to identifiable drivers such as portfolio growth, changes in borrower credit characteristics, shifts in macroeconomic forecasts, or differences in the timing of expected loss. When those drivers can be isolated and measured, the allowance becomes easier to interpret and easier to defend. When they cannot, explanations tend to rely more heavily on judgement and less on structural analysis. Explainability, therefore, is a function of model design.

The remainder of this paper examines the structural elements that determine whether a CECL framework produces a reserve that is opaque or defensible. Four design decisions in particular shape that outcome: the inputs used to represent risk, the granularity used to preserve risk distinctions, the conditional logic that allows assumptions to respond to changing conditions, and the structural model used to project expected loss through time. Together, these pillars determine whether a CECL implementation merely complies with the standard or provides a clear and defensible view of credit risk.

## Pillar I: The Right Inputs

### WHAT DATA ACTUALLY DRIVES LOSSES

A CECL framework can only be as informative as the inputs used to estimate expected loss. Regardless of the modeling approach, the allowance ultimately reflects the data used to represent borrower behavior, collateral performance, and economic conditions. When those inputs capture the drivers of credit risk, the reserve becomes responsive to real changes in the portfolio. When they do not, the model may appear stable but offer limited insight into how risk is evolving.

Many institutions possess extensive loan-level data, but the presence of data alone does not guarantee that the model captures meaningful risk signals. Loan files often contain dozens of fields, historical reports contain years of charge-off experience, and macroeconomic datasets offer countless indicators. The challenge is not simply collecting information; it is identifying which variables materially influence credit performance and structuring them so the model can interpret their effects.

In practice, the inputs that most consistently drive expected loss fall into four broad categories.

**Credit quality** remains one of the strongest predictors of borrower performance. Indicators such as credit scores, internal risk ratings, debt service coverage ratios, and delinquency status provide direct insight into the likelihood that a borrower will default. Monitoring how these characteristics evolve over time can reveal early shifts in portfolio risk.

**Collateral position** is equally important for secured lending. Loan-to-value ratios, updated property valuations, and collateral type influence both borrower incentives and potential recovery outcomes. Changes in collateral coverage can significantly alter expected severity, particularly during periods of economic stress or declining asset values.

**Exposure path** determines how risk unfolds over the life of a loan. Amortization schedules, remaining term, and prepayment behavior influence the balance that remains outstanding when losses occur. A framework that does not account for this evolving exposure risks applying loss assumptions to balances that may no longer exist.

**Macroeconomic conditions** provide the broader context in which borrowers operate. Unemployment trends, housing price movements, and interest rate environments affect both borrower performance and collateral recovery potential. CECL requires institutions to consider reasonable and supportable forecasts of these conditions, making macroeconomic inputs an essential component of forward-looking loss estimation.

| Input Category | Examples of Variables | Why It Matters |
|---|---|---|
| Credit Quality | Credit score, internal risk rating, debt service coverage ratio, delinquency status | Indicates the borrower's ability and willingness to repay, directly influencing default probability. |
| Collateral Position | Loan-to-value ratio, updated property values, collateral type | Determines potential recovery in the event of default and influences borrower incentives. |
| Exposure Path | Loan balance, amortization schedule, remaining term, prepayment behavior | Defines how exposure evolves over time and the balance outstanding when losses occur. |
| Macroeconomic Conditions | Unemployment rate, housing price index, interest rates, sector conditions | Captures the economic environment that influences borrower performance and collateral values. |

Table 1: Key Inputs That Drive Expected Credit Loss

*Table 1 illustrates that CECL inputs must capture the drivers of borrower performance and recovery.*

Together, these inputs form the foundation of a defensible CECL framework. The key is not the quantity of variables but their ability to capture the drivers of credit performance. When inputs are properly structured, the model begins to reflect observable changes in the portfolio. Shifts in borrower credit characteristics influence expected default behavior. Changes in collateral values alter recovery expectations. Updated economic forecasts affect projected performance paths. The reserve becomes interpretable because its drivers are visible.

When inputs are misaligned or incomplete, the opposite occurs. Historical loss averages may be applied broadly across the portfolio without regard to underlying credit attributes. Collateral values may remain static even as market conditions change. Macroeconomic forecasts may influence the allowance only through high-level overlays rather than integrated assumptions. In these cases, the model's ability to explain reserve movements becomes limited.

> Inputs determine what the model can see.

Selecting the right inputs therefore establishes the foundation for explainability. Once those inputs are identified, the next challenge is determining how they should be organized within the model. Even the most informative data can lose its value if it is aggregated too broadly. The second pillar addresses this challenge: the level of granularity used to preserve meaningful distinctions within the portfolio.

## Pillar II: The Right Granularity

### WHEN AGGREGATION DESTROYS INFORMATION

Identifying the right inputs is only the first step in building a defensible CECL framework. Equally important is how those inputs are organized within the model. The level of granularity used to segment the portfolio determines whether meaningful distinctions in risk are preserved or averaged away. Granularity, in this sense, is about preserving information.

Many CECL frameworks rely on broad portfolio groupings that combine loans with materially different credit characteristics into a single segment. This approach simplifies modeling and often produces reserves based on stable historical loss averages. However, it also compresses important differences in borrower quality, collateral position, and loan structure into a single assumption. When those differences are averaged together, the model loses its ability to respond proportionally to changes in portfolio composition.

Consider a portfolio segment that includes borrowers with both strong and weak credit profiles. If the model applies a single loss assumption across the entire group, growth in higher-risk loans may not significantly affect the modeled loss rate. Similarly, improvements in borrower credit quality may not reduce expected loss in a measurable way. The reserve becomes less sensitive to the underlying drivers of risk. Aggregation creates stability, but it can also obscure causality.

**Table 2: What Happens When Risk Is Aggregated**

| Example Portfolio Segment | Key Risk Attributes | True Loss | Aggregated Loss | Distortion |
|---|---|---|---|---|
| **Prime** | Credit Score > 760, LTV < 60% | 0.10% | 0.55% | Overstated |
| **Standard** | Credit Score 680 - 760, LTV 60% - 80% | 0.35% | 0.55% | Slightly Overstated |
| **Higher Risk** | Credit Score <680, LTV > 80% | 1.20% | 0.55% | Understated |

*Table 2 illustrates that aggregation produces a single average loss rate that ails to reflect the true risk profile of individual loans.*

A more effective approach segments the portfolio according to the attributes most closely linked to credit performance. For consumer and residential lending, this often involves stratifying loans by combinations of credit score and loan-to-value ratio. For commercial portfolios, segmentation may incorporate debt service coverage, leverage, and industry exposure. Delinquency status and origination channel can also provide meaningful distinctions in borrower behavior.

| Table 3: How Portfolio Segmentation Preserves Risk Distinctions | | | | |
|---|---|---|---|---|
| **Example Portfolio Segment** | **Key Risk Attributes** | **Expected Default Behavior** | **Expected Severity** | **Why Segmentation Matters** |
| **Prime Mortgages** | Credit Score > 760, LTV < 60% | Very Low | Very Low | Strong Borrower credit and significant collateral cushion reduce both default probability and loss severity. |
| **Standard Mortgages** | Credit Score 680 - 760, LTV 60% - 80% | Low | Moderate | Moderate credit strength and collateral position produce stable but non-zero loss expectations. |
| **Higher Risk Mortgages** | Credit Score <680, LTV > 80% | Elevated | Higher | Lower borrower credit quality and limited equity increase both default probability and potential loss severity. |
| **Investor CRE Loans** | DSCR 1.25-1.5, Stabilized Property | Moderate | Moderate | Income-producing collateral reduces risk but performance remains tied to property cash flow and economic cycles. |
| **Higher Risk CRE Loans** | DSCR <1.25, Higher Leverage | Elevated | Higher | Weaker cash flow coverage and higher leverage increase sensitivity to economic conditions and refinancing risk. |

*Table 3 illustrates that meaningful segmentation preserves differences in borrower risk, collateral position, and economic sensitivity that would otherwise be lost through*

These segmentation approaches allow the model to differentiate between loans that exhibit materially different risk characteristics. Conditional assumptions can then be applied to each segment, allowing expected performance to evolve as borrower attributes change. In this way, granularity preserves the relationship between inputs and outcomes.

Granularity decisions also influence the choice between loan-level and cohort-level modeling. Real estate loans (both residential and commercial) benefit from loan-level analysis as each contains unique collateral characteristics or large exposure. Individual loans may possess attributes that materially influence expected loss, making it important to evaluate them separately.

Other portfolios may be more appropriately modeled using stratified cohorts. High-volume consumer lending products, for example, often contain large numbers of loans with similar structural characteristics. In these cases, segmentation by key risk attributes can preserve relevant distinctions while maintaining statistical stability.

The objective is meaningful segmentation over maximum segmentation. Excessive granularity introduces its own challenges. Very narrow segments can lead to sparse historical data, making it difficult to estimate reliable performance assumptions. Small sample sizes may produce volatile estimates and unstable model behavior. A defensible CECL framework therefore balances informational detail with statistical credibility.

When segmentation captures the true drivers of credit performance without fragmenting the dataset, the model gains explanatory power. Changes in portfolio composition translate into measurable shifts in expected performance. Growth in higher-risk segments increases modeled loss. Improvements in borrower credit attributes reduce it.

The reserve becomes responsive because the distinctions that matter have been preserved. Granularity therefore serves a critical role in the architecture of a CECL framework. It determines whether the model captures the shape of risk within the portfolio or reduces it to a single average.

Granularity determines what the model can distinguish.

Once those risk distinctions are preserved, the next challenge is determining how the model responds to them. The third pillar addresses this issue by examining the logic used to translate borrower attributes and economic conditions into expected performance.

## Pillar III: The Right Conditional Logic

### ASSUMPTIONS MUST RESPOND TO RISK

Once the appropriate inputs have been identified and the portfolio has been segmented in a way that preserves meaningful risk distinctions, the next design question emerges: *how should the model respond to those differences?*

This is where conditional logic becomes essential. A forward-looking credit framework must translate borrower attributes, collateral position, and economic conditions into expected performance. In many CECL implementations, however, core assumptions remain largely static. Default probabilities, loss severity, or prepayment speeds may be applied uniformly across large segments of the portfolio, regardless of variations in borrower characteristics or economic outlook.

While this approach can produce stable results, it often requires additional adjustments to reflect changing conditions. Qualitative overlays are frequently introduced to compensate for shifts in credit quality, collateral values, or macroeconomic expectations that the underlying model does not fully capture. The result is a framework that relies increasingly on judgement rather than structure.

Conditional logic addresses this limitation by allowing key performance assumptions to vary in response to observable risk drivers. Instead of applying a single default rate across an entire portfolio segment, the model can assign default rate assumptions conditional on attributes such as credit score, leverage, debt service coverage, or delinquency status. As those characteristics evolve, the projected default behavior evolves with them.

| Table 4: How Conditional Assumptions Reflect Borrower Risk | | | |
|---|---|---|---|
| Borrower Profile | Key Risk Attributes | Economic Environment | Conditional Default Rate |
| **Strong Borrower** | Credit Score 780, LTV 60% | Stable Economy | 0.40% |
| **Strong Borrower** | Credit Score 780, LTV 60% | Recession Scenario | 0.90% |
| **Higher-Risk Borrower** | Credit Score 640, LTV 90% | Stable Economy | 2.80% |
| **Higher-Risk Borrower** | Credit Score 640, LTV 90% | Recession Scenario | 5.10% |

*Table 4 illustrates that conditional assumptions allow CECL models to respond proportionally to differences in borrower credit quality and collateral position.*

A similar approach applies to loss severity. For secured lending, recovery expectations are closely tied to collateral coverage. As loan-to-value ratios change or property values shift with market conditions, the expected loss severity should adjust accordingly. A model that incorporates these relationships directly is able to respond to changes in collateral risk without relying on broad qualitative adjustments.

Prepayment behavior should also be modeled conditionally. Borrowers often respond to changes in interest rate incentives, housing turnover conditions, or loan seasoning. By allowing prepayment assumptions to vary with these factors, the model can better reflect how exposure evolves over time.

Macroeconomic variables represent another important dimension of conditional logic. Economic forecasts influence borrower performance through multiple channels. Rising unemployment may increase default probability. Declining housing prices may affect both borrower incentives and collateral recovery values. Changes in interest rates may alter refinancing behavior. When these relationships are incorporated directly into the model, the framework becomes responsive to the broader environment in which borrowers operate.

The key advantage of conditional logic is that it allows the model to respond proportionally to changes in risk. As borrower attributes deteriorate, default expectations increase. As collateral overage improves, severity assumptions decline. As macroeconomic forecasts shift, projected performance paths adjust. The reserve changes because the underlying drivers of risk have changed.

In contrast, static assumptions often produce the opposite dynamic. The model itself remains largely unchanged while the surrounding environment evolves. To bridge the gap between model output and observed conditions, institutions introduce qualitative adjustments. Over time these adjustments can accumulate and make the allowance increasingly difficult to interpret. Conditional logic reduces the need for these adjustments by embedding the response to risk directly within the structure of the model.

This does not eliminate judgement. Credit modeling will always require informed assumptions and periodic recalibration. However, when those assumptions are tied to observable drivers of performance, the framework becomes more transparent and mode stable. The allowance moves because borrower behavior, collateral conditions, or economic expectations have changed – not because an adjustment was applied after the fact. In this way, conditional logic reinforces the explainability of the reserve. It ensures that changes in the portfolio and the economic environment translate naturally into changes in expected loss.

Conditional logic determines how the model responds.

The final pillar addresses how those expectations are translated into the allowance itself. Even when inputs are appropriate, segmentation preserves risk distinctions, and assumptions respond to changing conditions, the modeling framework must still determine how losses unfold through time. The structure of the model ultimately determines whether those dynamics are preserved or compressed.

## Pillar IV: The Right Structure

### WHY DISCOUNTED CASH FLOW PRESERVES CAUSALITY

Even when the correct inputs are identified, meaningful segmentation is established, and assumptions respond to borrower attributes and economic conditions, one critical design decision remains: *the structure used to translate those expectations into an allowance.*

The structure of the model determines whether the drivers of credit risk remain visible or become compressed into a single estimate. At its core, expected credit loss is a time-dependent event. Borrowers make scheduled payments over time. Some loans prepay; others default. Collateral may be liquidated, and recoveries may occur months or years after default. The magnitude and timing of these events determine the economic impact of credit loss. A modeling framework must therefore preserve three elements simultaneously: exposure, probability, and time.

Discounted cash flow (DCF) modeling provides a structure that maintains these relationships explicitly.

A DCF-based CECL framework begins with the contractual cash flows of the loan or portfolio segment. Scheduled principal and interest payments define the baseline exposure path. From that baseline, the model incorporates the behavioral assumptions that determine how those cash flows change over time. Conditional prepayment alters the remaining balance trajectory. Conditional default determines the volume of loans that cease performing. Severity assumptions estimate the portion of exposure that will ultimately be lost, and liquidation timelines determine when recoveries occur. The final step discounts the expected cash flows to present value.

By modeling these elements through time, a DCF structure preserves the path through which loss occurs. Exposure evolves through amortization and prepayment. Defaults occur at varying levels along the exposure path. Recoveries arrive with their own timing. The allowance reflects the present value of those expected outcomes.

This structure provides several important advantages.

First, it preserves causality. Changes in borrower attributes, collateral conditions, or macroeconomic forecasts influence default probabilities, severity expectations, or prepayment behavior. Those changes flow through the model's structure and ultimately affect the timing and magnitude of expected loss. The resulting reserve movement can be traced back to specific drivers.

Second, the framework maintains sensitivity to timing. Two portfolios may have similar lifetime loss percentages but very different loss timing. Defaults that occur earlier in a loan's life typically have a greater present value impact than those occurring later. A structure that preserves timing captures this difference directly.

Third, discounted cash flow modeling integrates naturally with forward-looking macroeconomic scenarios. Changes in unemployment, housing prices, or interest rates influence borrower behavior and collateral performance. Because these variables affect the timing and probability of events within the model, their impact on expected loss can be evaluated consistently across different economic paths.

Finally, the structure provides transparency. Each assumption – prepayment, default, severity, liquidation timing – represents a distinct component of expected loss. When a reserve movement occurs, the institution can evaluate which component contributed to the change. This modular structure enhances explainability.

Alternative modeling approaches often compress these relationships into a single lifetime loss rate. While such methods may satisfy minimum compliance requirements, they reduce the ability to distinguish between exposure changes, timing effects, and behavioral shifts. Loss becomes an average rather than a path.

A defensible CECL framework benefits from a structure that preserves the path. Discounted cash flow modeling accomplishes this by aligning the mechanics of the model with the economic reality of credit loss. Exposure evolves through time, borrower behavior alters that exposure, and expected loss reflects the present value of the resulting shortfall.

When the structural framework maintains these relationships, the allowance becomes easier to interpret, easier to govern, and easier to defend.

Structure determines how the model explains.

| Table 5: Discounted Cash Flow Overview | |
|---|---|
| **Aspect** | **Description** |
| **Methodology** | Projects future cash flows and discounts them to present value to estimate credit losses. |
| **Macroeconomic Considerations** | Fully integrates past events, current conditions, and forward-looking macroeconomic forecasts. |
| **Credit Loss Calculation** | Models default probability and loss severity separately, enhancing accuracy and granularity. |
| **Data Granularity** | Analyzes loans individually or in detailed cohorts, incorporating updated borrower credit and collateral data. |
| **Model Complexity** | Prospective and dynamic, incorporating detailed loan-level attributes and changing conditions. |
| **Prepayments** | Prepayments are modeled directly based on borrower incentives, market interest rates, and updated loan information. |
| **Use Cases** | Can be used for multiple purposes beyond reserve estimation, including ALM, stress testing, and loan pricing. |
| **Adjustments** | Typically requires fewer adjustments due to its granularity and incorporation of current and forecasted conditions. |
| **Predictive Power** | Highly predictive, adjusting dynamically to changes in borrower creditworthiness and economic forecasts. |

*Table 5 highlights the advantages and functionality of a discounted cash flow model.*

With the four pillars in place – inputs, granularity, conditional logic, and structure – the CECL framework gains the elements necessary to produce an explainable reserve. The final section considers what this architecture looks like in practice and how institutions can recognize when a CECL framework has achieved structural integrity.

## What a Defensible Reserve Looks Like

When the four pillars of a CECL framework align – appropriate inputs, meaningful granularity, conditional assumptions, and a coherent structural model – the allowance begins to exhibit a set of recognizable characteristics. These characteristics are less about the size of the reserve and more about the clarity with which it can be interpreted.

A defensible reserve is explainable.

Changes in the allowance can be traced to identifiable drivers. Growth in the portfolio increases exposure. Shifts in borrower credit quality influence projected default behavior. Changes in collateral coverage alter expected recovery outcomes. Updated macroeconomic forecasts affect projected performance paths. Each of these drivers flows through the model in a way that allows the institution to understand why the reserve moved.

In practical terms, this means the allowance can be decomposed into its underlying components. Quarter-over-quarter changes can be attributed to factors such as portfolio growth, mix shifts, credit migration, macroeconomic forecast updates, or differences in the timing of expected loss. The reserve becomes easier to interpret because the relationships between inputs, assumptions, and outcomes remain visible.

A defensible reserve is also responsive.

Since the model incorporates the drivers of borrower behavior and economic conditions, it reacts proportionally as those conditions evolve. Deterioration in borrower credit characteristics increases expected default behavior. Improvements in collateral values reduce expected severity. Changes in macroeconomic forecasts influence projected performance through modeled relationships. The allowance adjusts because the underlying drivers of risk have changed.

At the same time, the *framework* maintains stability. Stability does not mean the reserve remains constant. It means the model responds to changes in risk in a predictable and coherent way. When assumptions are tied to observable drivers and the structure of the model preserves exposure and timing, reserve movement tends to reflect genuine changes in portfolio conditions rather than unexplained volatility. These qualities ultimately support stronger governance.

The allowance can be explained to boards of directors with confidence. Auditors can evaluate the methodology with a clear understanding of how assumptions influence outcomes. Regulators can assess whether the framework appropriately reflects the risk embedded in the institution's loan portfolio. In each case, the allowance is supported by a structure that connects observable risk drivers to the resulting estimate of expected loss.

CECL is often viewed primarily as an accounting requirement, but its design has broader implications. At its best, the framework functions as a forward-looking system for understanding how credit risk evolves across a portfolio. When the model preserves the drivers of loss and the structure through which those losses unfold, the allowance becomes an analytical reflection of credit risk.

A defensible reserve is therefore defined by the structural integrity of the framework that produces it. Institutions that design their CECL models with this architecture in mind gain a clearer view of how risk develops within their portfolios and a stronger foundation for explaining that risk when it matters most.

> If your institution is ready to move beyond opaque reserve calculations and toward a CECL framework that produces explainable results, Wilary Winn can help you implement that architecture quickly and defensibly. Our approach applies the same structural principles described in this paper – a discounted cash flow foundation, meaningful segmentation that preserves risk distinctions, conditional assumptions that respond to borrower attributes and macroeconomic conditions, and a framework designed to translate those drivers into an allowance that can be clearly decomposed and explained. Contact Wilary Winn today to implement a framework for a defensible CECL reserve for your institution.