



STRENGTHENING FINANCIAL INSTITUTIONS

Defensible Default Assumptions for Commercial Real Estate: Evidence from Loan-Level Data

Released August 2026

Introduction

Commercial real estate default assumptions carry more weight than almost any other input in a reserve analysis or stress result, yet most institutions cannot trace theirs back to supporting data. Internal default history is too thin, peer averages describe other institutions' portfolios and vendor assumptions arrive without a detailed view of the loans behind them. This paper describes a different approach: deriving default assumptions from the loan-level disclosures required under Regulation AB II, a public record of how a large population of commercial properties has actually performed. The paper delivers:

- A default grid built from public, loan-level performance data
- Forward-looking application consistent with CECL, ALM, and stress testing

A default assumption is only as defensible as the data to which it can be traced.

KEY TAKEAWAY

Default assumptions developed from public, loan-level performance data provide greater accuracy, transparency, segment-level relevance, and a forward-looking analytical foundation. Peer averages and opaque vendor assumptions do not.

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

Executive Summary

Commercial real estate default is a low-frequency, high-severity event, which makes it both the most consequential assumption in a commercial credit model and the hardest to support with an institution's own history. Most institutions bridge the gap with peer averages, vendor default inputs, or broad charge-off history, and each of these substitutes fails the same test: when a reviewer asks where the number comes from and why it fits this portfolio, the answer does not trace back to representative data.

This paper presents an alternative built on the asset-level disclosures required under Regulation AB II. Since 2016, publicly issued commercial mortgage securitizers have filed standardized loan-level data with the SEC, covering origination characteristics such as property type, debt service coverage, and loan-to-value, along with ongoing payment performance through liquidation. The record now spans over nine years and includes over 23,000 loans with aggregate original balances of \$396 billion, a body of observed default behavior that is far larger than any single institution's experience and, unlike vendor supplied input data, fully public.

The paper describes the data and its boundaries, how the assumptions are constructed, what the record shows, and how the assumptions apply on a forward-looking basis consistent with CECL. Rather than averaging default experience within cohorts, we analyze the relationship between default and its two most consistent drivers, debt service coverage and loan-to-value, directly from the loan-level records.

The conclusion is practical. An institution that grounds its default assumptions in loan-level performance data can answer a question that peer medians and vendor-supplied data cannot: *What is the assumption, which loans support it, and why are those loans comparable to our portfolio?*

The Problem: Defending Commercial Default Assumptions

THIN DATA, BORROWED AVERAGES, AND ACCUMULATED JUDGMENT

Every commercial real estate reserve, and every stress result that depends on one, rests on an assumption about how often loans default. That number is small, is difficult to observe directly, and moves the answer more than almost any other input. Yet when an examiner, an auditor, or a board member asks a simple question about where that default rate comes from, many institutions struggle to give an answer that traces back to data.

The difficulty begins with the data an institution has on hand. A commercial real estate loan default is a low-frequency event. A well-underwritten portfolio can go years with few defaults, or none at all, particularly during the benign part of a cycle. This is a good outcome for the institution, but it leaves the modeler with too few events to estimate a credible default rate for any specific portfolio segment. An institution might hold hundreds of commercial real estate loans spread across office, retail, multifamily, and industrial properties and still lack enough realized defaults to say, with any confidence, how an office loan behaves differently from an industrial one. The internal record, on its own, cannot support the distinctions that matter most.

When internal data runs short, institutions turn to substitutes. Each substitute weakens the accuracy and defensibility of the result.

Peer data, drawn from call report averages or industry aggregates, reflects a national blend of institutions and property types that rarely matches any single portfolio. It answers the question of how commercial loans have performed on average, not the question of how the individual attributes of a specific institution's loan portfolio affects performance.

Vendor default assumptions offer convenience, but they often arrive without a clear view of the population from which they were derived. When the source loans are not disclosed, the institution cannot demonstrate that the assumption reflects its own portfolio segments or explain the number when someone asks.

Historical charge-off averages, applied broadly across a portfolio, compress real differences in property type, loan structure, and debt service coverage into a single rate. **The assumption is stable, but only because it has been averaged past the point of meaning.**

Each of these approaches can satisfy the mechanics of a reserve or a stress calculation. The problem surfaces later, in the governance environment where the assumption has to be defended. Auditors must judge whether the methodology is conceptually sound. Regulators must determine whether it reflects the risk actually embedded in the portfolio. Boards must understand the number well enough to stand behind it. Across all three, the same question returns: *Where does the default rate come from, and why is it the right one for this portfolio?*

An assumption that cannot be traced to representative data invites a familiar cycle. Qualitative overlays accumulate to bridge the gap between the modeled rate and observed conditions, and the explanation shifts from evidence toward judgment. Over time the default assumption becomes harder to interpret, and the institution carries a number it cannot fully account for.

A default assumption is only as defensible as the data to which it can be traced.

This matters more for commercial real estate than for almost any other segment of a loan portfolio. Commercial real estate concentrations are high at many community banks and credit unions, the individual loans are large, and a default, when it occurs, tends to be severe. Recent conditions have added to the pressure, as loans underwritten in a low-rate environment reach maturity and refinance at materially higher interest rates, and as some property types work through structural shifts in demand. In this environment, supervisory attention to how institutions support their commercial real estate assumptions has increased, and a default rate that rests on a peer median or an undisclosed vendor population is increasingly difficult to defend.

The way out of this problem is not more judgment layered on top of thin data. It is a larger and more granular body of observed default behavior, drawn from commercial loans whose characteristics can be matched to the institution's own. The data exists. Under Regulation AB II, publicly issued securitizers must provide loan-level disclosure for their commercial mortgage holdings, creating a detailed and public record of how a large population of commercial properties has actually performed over time. The next section describes that data and why it provides a stronger foundation for defensible default assumptions.

The Data: Loan-Level Disclosure Under Regulation AB II

A PUBLIC, STANDARDIZED, LOAN-LEVEL RECORD

Regulation AB II is a set of SEC disclosure rules governing publicly offered asset-backed securities.¹ The Commission adopted the rules in September 2014, and asset-level disclosure became required for offerings made on or after November 23, 2016. Issuers file the data on Form ABS-EE as Exhibit 102.² The filing is made at issuance alongside the prospectus and then updated with each periodic distribution report, producing a continuous public record of each loan from origination through payoff, modification, or liquidation.

The disclosures are extensive. For each loan in a covered securitization, the filings report the characteristics that drive credit performance: property type, geographic location, origination date and balance, interest rate and amortization structure, maturity date, debt service coverage, property value, and occupancy. They then track performance over the life of the loan: payment status, delinquency, transfer to special servicing, modification, workout strategy, and, where a loan fails, the liquidation proceeds and realized loss.

The result is a public performance history for a large population of commercial mortgages. The dataset assembled for this analysis from EDGAR filings covers over 23,000 loans across 369 securitization trusts, with an aggregate original balance of \$396 billion, with monthly performance observed from December of 2016 through June of 2026.³ It spans the pandemic shock of 2020 and 2021, the 525 basis point increase in the federal funds target range between March 2022 and July 2023, and the repricing of office and other property types that followed. Default assumptions estimated over this period reflect performance through genuine stress, rather than only the benign years that distort historical averages.

Three properties of this data distinguish it from the substitutes described above:

1. It is loan-level. Peer averages describe portfolios; this data describes loans. Default behavior can be measured within combinations of property type, debt service coverage, and loan-to-value, which is precisely the granularity at which an institution underwrites and at which risk varies.
2. It is standardized. Every issuer reports the same fields under the same definitions, which means observations from hundreds of securitizations can be combined into a single analytical population without the reconciliation problems that plague pooled bank data.
3. It is public. Every filing sits on the SEC's EDGAR system, available to any analyst, auditor, or examiner who wants to verify it. An assumption derived from this data does not ask the reviewer to trust a vendor's undisclosed population. It invites the reviewer to check the work.

¹ Asset-Backed Securities Disclosure and Registration, SEC Release No. 33-9638, 79 Fed. Reg. 57184 (September 24, 2014).

² Required content is specified in Schedule AL, Item 1125 of Regulation AB, 17 C.F.R. § 229.1125.

³ Filed as Exhibit 102 to Form ABS-EE, available through the SEC's EDGAR system at www.sec.gov/edgar.

The value of this data is not only its size. It is that every number derived from it can be traced back to a public record.

Population and Representativeness

WHERE THE DATA FITS, AND WHERE IT DOES NOT

The usefulness of any external dataset depends on how closely its loans resemble the loans the assumptions will be applied to. For the securitized commercial mortgage population, the resemblance is close in some respects and limited in others. The differences are identifiable, and each can be addressed directly.

The loans in public commercial mortgage securitizations share a recognizable profile. They are secured by income-producing, non-owner-occupied properties. They are generally larger than the typical community financial institution loan. They are predominantly non-recourse, underwritten to the property's cash flow rather than to a guarantor. They are serviced by third parties under pooling agreements rather than managed by the lender that originated them. The population comprises over 23,000 loans totaling roughly \$396 billion, with an average balance of approximately \$18 million and median balances by property type ranging from \$2.5 million for cooperative housing to \$23.7 million for office. Office, retail, multifamily, lodging, mixed use, and industrial collateral together account for roughly 87% of total balances.

A community financial institution's commercial real estate portfolio differs on several of these dimensions. Recourse to guarantors is common. The lender holds the relationship, monitors the borrower directly, and can intervene early when a credit deteriorates. Loan sizes are smaller, and the property mix leans further toward smaller multifamily, mixed-use, and local retail and industrial properties.

These differences have directional effects. Non-recourse borrowers can surrender an underwater property at limited personal cost, so the securitized population might default more readily than otherwise comparable recourse loans. Third-party servicers might resolve distress differently than a relationship lender by being less accommodating to the borrower. Where an effect runs in a known direction, it can be stated and reflected in how the assumption is applied.

Owner-occupied commercial real estate is essentially absent from this data. Securitized commercial mortgages are investor properties. **An owner-occupied portfolio, which behaves more like a business loan secured by real estate than like an investment property, requires a different foundation, and this paper does not extend to it.**

Within its boundaries, the data permits matching on characteristics rather than reliance on general resemblance. Every loan carries its property type, coverage, and value, so default behavior is measured against the same attributes an institution tracks on its own credits. A CRE-concentrated institution holding stabilized multifamily and retail properties does not need the average of the securitized universe. It needs to know how default behavior varies with coverage and value among properties of those types, and the data is granular enough to answer exactly that.

Representativeness is therefore a matching exercise rather than a yes-or-no judgment, and loan-level detail is what makes the matching possible.

Constructing the Assumptions

DEFAULT AS A FUNCTION OF COVERAGE AND VALUE

The objective is a default assumption that can be stated for any loan, responds to the condition of that loan, and rests on observed experience.

The first requirement is a clear definition of default. A loan is treated as defaulted if it is ninety or more days delinquent, if it has been disposed of or liquidated, if it has been resolved through a discounted payoff, or if it has a recorded realized loss. The definition captures loans that have entered the loss process, and it pairs consistently with how severity is measured, so default frequency and loss given default describe the same event.

The more consequential choice is what to do with the data once default is defined. Institutions conventionally estimate commercial default rates by pooling. Loans are grouped by portfolio segment, historical charge-off experience is averaged over a lookback period, and the average is applied to the segment's current balance, supplemented with peer data where internal history is thin. More developed frameworks assign default rates through the internal risk rating, where each grade carries a rate and a loan's assumption changes only when its grade changes. Both approaches attach the assumption to the group rather than the loan. A pooled average does not distinguish a loan at 1.50 coverage from one at 1.05 in the same segment, and a rating-based rate responds to deteriorating coverage only when the credit is downgraded, a step that is discrete, lagged, and dependent on judgment. Institutions that go further and tabulate default rates within combinations of characteristics encounter a different limitation: default is rare, cells multiply quickly, and the cells carrying the most weight are often supported by a small number of events.

We take a different approach. Rather than averaging within segments, we estimate the relationship between default and its drivers directly from the loan-level records, using each loan's months of exposure so that the result is a rate per unit of time rather than a count of events. The two drivers are debt service coverage ratio and loan-to-value, the characteristics that most consistently separate loans that default from loans that do not. The output is a continuous relationship: any combination of coverage and value returns a default rate, including combinations that no loan in the data happens to occupy exactly.

Three properties follow from estimating a relationship rather than tabulating averages:

1. Every loan informs every estimate. A sparsely populated region of the relationship takes its value from the behavior of the full population rather than from the small number of loans that occupy it.
2. The economic ordering holds throughout. Default rises as coverage weakens and as loan-to-value increases, across the full range of both drivers, without smoothing or reordering after the fact.

3. The assumption responds continuously. A loan moving from 1.30x coverage to 1.20x coverage receives a different rate, rather than the same rate until a cohort boundary is crossed.

As the condition of a loan changes, the assumption changes with it.

The relationship is estimated separately for each property type because coverage and value do not carry identical weight in office, retail, multifamily, industrial, and hospitality lending. It is also estimated once across the entire population, producing an all-property relationship.

The volume of experience behind each property type varies, and the estimation accounts for this directly. Where a property type is well represented, the assumption rests on that type's own estimated relationship. Where the data behind a type is thinner, the assumption draws on the type's own experience to the extent that experience supports it, and on the relationship estimated across the full population for the remainder. The weighting is determined by the data rather than assigned by judgment. The major property types each carry between 88% and 97% of their own experience, while thinly represented types, such as cooperative housing and healthcare, rely principally on the broader population until their own record accumulates.

We also checked whether weak coverage and high leverage compound, meaning whether a loan is more dangerous when both are present than the two effects would suggest on their own. The data did not support treating them that way, so each driver is applied independently.

The estimated relationships are applied without adjustment; the starting point is what the data supports. Where an institution's circumstances warrant judgment, it belongs in a distinct and visible layer above the estimated relationship rather than embedded within it, so a reviewer can always see where the evidence ends and judgment begins.

Estimated Default Behavior

DEFAULT RATES ACROSS COVERAGE AND VALUE

The exhibits below present fitted annual default rates at representative combinations of debt service coverage and loan-to-value. The underlying relationship is continuous and returns a rate for any combination; the grid format is a presentation convenience, not the structure of the assumption. The row and column totals are also fitted values, each conditioned on a single driver, and the overall total is the fitted rate for the population as a whole.

Exhibit 1: All-Property Fitted Annual Default Rates, DSCR x LTV						
All-Property	Below 1.00x	1.00x - 1.25x	1.25x - 1.50x	1.50x - 2.00x	2.00x+	Total
0-40%	0.76%	0.34%	0.25%	0.15%	0.10%	0.11%
40-60%	1.82%	0.82%	0.60%	0.37%	0.23%	0.27%
60-80%	3.24%	1.47%	1.07%	0.66%	0.41%	0.48%
80-100%	5.73%	2.62%	1.91%	1.18%	0.73%	0.86%
100%+	13.20%	6.16%	4.51%	2.82%	1.75%	2.05%
Total	3.74%	1.70%	1.24%	0.77%	0.47%	0.56%

Exhibit one presents the all-property relationship. The economic ordering holds in both directions across the full surface: within every coverage band, default rises as loan-to-value increases, and within every loan-to-value band, default falls as coverage strengthens. The magnitude of the differences is the central finding. A loan with coverage above 2.00x and loan-to-value below 40% carries a fitted annual default rate of 0.10%. A loan with coverage below 1.00x and loan-to-value above 100% carries 13.20%. Between the strongest and weakest combinations lies a factor of more than one hundred.

The totals illustrate what a single average conceals. The fitted rate for the population as a whole is 0.56%, which is the order of figure an institution applies when it reserves its commercial portfolio using one blended assumption. Loans with strong coverage and low loan-to-value default at a fraction of that rate, while loans with weak coverage and high loan-to-value default at several multiples of it. A reserve built on the average is therefore overstated for the majority of a well-underwritten portfolio and understated for the loans which carry the greatest risk.

Exhibit two presents the office relationship, blended as described in the prior section. Property type continues to matter after coverage and value are considered. Measured across its population, office defaults at 0.75% annually against 0.56% for all property types. The shape differs as well: office rates at strong combinations of coverage and value run above the all-property surface, while rates at the weakest combinations run below it, a flatter relationship consistent with a sector stress that has affected office credits broadly rather than only the weakest among them.

Exhibit 2: Office Fitted Annual Default Rates, DSCR x LTV						
Office	Below 1.00x	1.00x - 1.25x	1.25x - 1.50x	1.50x - 2.00x	2.00x+	Total
0-40%	0.56%	0.32%	0.25%	0.18%	0.13%	0.14%
40-60%	1.37%	0.78%	0.62%	0.44%	0.32%	0.35%
60-80%	2.47%	1.41%	1.13%	0.80%	0.57%	0.64%
80-100%	4.46%	2.55%	2.04%	1.46%	1.04%	1.17%
100%+	10.60%	6.16%	4.95%	3.55%	2.55%	2.85%
Total	2.87%	1.64%	1.31%	0.93%	0.67%	0.75%

Office currently carries the highest incidence of delinquency in the population, and those credits have not yet resolved into realized default. As they do, the additional events enter the estimation directly, and the office relationship updates on observed evidence rather than through an adjustment applied after the fact.

The exhibits present rates as of the measurement date. Their application over a projection period is the subject of the next section.

Forward-Looking Application

FROM MEASURED EXPERIENCE TO PROJECTED DEFAULT

A relationship estimated on historical data becomes a forward-looking assumption through two mechanisms:

1. The condition of the loan itself. The input assumption is not a portfolio average applied uniformly. Each loan receives the rate implied by its own coverage and value, so as a borrower's operating performance changes, the assumption changes with it. A property whose coverage has fallen since origination carries a higher rate at the next measurement automatically, because the rate is a function of the loan's own condition rather than a manual adjustment applied on top of it. This is the response to credit migration that blended assumptions cannot produce.
2. The economic environment. Default frequency does not hold constant through a cycle, and an assumption fixed at its historical level will understate risk entering a downturn and overstate it coming out. The starting rate implied by each loan's characteristics is therefore scaled over the projection period by a factor that follows the economic forecast, driven by projected unemployment and specified by property type, since property types do not share the same sensitivity to economic conditions.

The assumption therefore moves for two distinct reasons: the condition of the loan changed, or the economic outlook changed. Each can be isolated, which allows reserve movement to be decomposed and explained. For CECL, the structure maps onto the standard's requirements without modification.⁴ Over the reasonable and supportable forecast period, rates follow the institution's economic forecast through the scaling factor. Beyond that horizon, the factor reverts and each loan settles toward the rate implied by its own characteristics under long-run conditions, a reversion target grounded in the observed experience of comparable properties.

The same assumptions serve the institution's other risk modeling frameworks. In asset liability management and stress testing, an adverse economic path scales the same underlying relationships, so the default assumptions behind a stress result and those behind the allowance are consistent by construction. One body of evidence, applied across CECL, ALM, and concentration analysis, replaces the patchwork of unrelated assumptions that examiners rightfully increasingly question.

Each element of the assumption can be supported directly. The rate applied to each loan follows from that loan's coverage and value, the relationship behind the rate was estimated from publicly disclosed commercial mortgage performance, and the forward path follows the institution's own economic forecast. At no point does the explanation rest on an undisclosed population or an unsupported judgment.

Conclusion

Commercial real estate default assumptions have long rested on a weak data foundation: too few internal observations, substitutes that describe other institutions' portfolios, and vendor numbers that cannot be inspected. Regulation AB II changed what is possible. A public, standardized, loan-level record of commercial mortgage performance now exists, deep enough to estimate how default behavior varies with the characteristics institutions already track, and recent enough to reflect performance through stressed conditions.

⁴ FASB ASC 326-20 addresses the reasonable and supportable forecast period and reversion to historical loss information.

The data supports an assumption that responds when a loan's coverage deteriorates, that can be explained by reference to the characteristics that produced it, and that rests on evidence a reviewer can verify. For institutions with significant commercial real estate concentrations, the value of assumptions that behave this way compounds across every framework that uses them: the allowance, the stress results, and the conversations with auditors and examiners that follow both.

The approach described in this paper aligns with how Wilary Winn builds assumptions: estimated from observed performance, responsive to the condition of each loan, and supported by evidence a reviewer can verify. A reserve is only as defensible as the assumptions that produce it, and the same is true for every analysis built on those assumptions. Contact us to discuss how this approach would apply to your institution.