

Underwriting High-Yield GCC Credit: Pricing the Risk Behind the Coupon

A disciplined framework for translating a double-digit coupon into the probability of default, loss given default and capital that justify it, so the lender prices the risk rather than chases the yield.

Chennakeshav Adya

Independent Researcher

June 2026

ABSTRACT

A high coupon is a question, not an answer. The companion question to where do Gulf double-digit coupons come from is whether, on any given transaction, the coupon adequately prices the risk it carries. This paper sets out a disciplined underwriting and pricing framework for high-yield Gulf real-asset credit, built around the two quantities that determine whether a coupon is sufficient: the probability of default and the loss given default. It develops a structured method for estimating each, for combining them into an expected loss and an unexpected-loss capital charge, and for deriving the coupon a rational lender should require given a target return on the capital at risk. The framework is deliberately conservative and transparent: every required coupon can be traced to an assumption about default, recovery, cost and return, and the assumptions are stress-tested. The central argument is that underwriting, not yield, should drive the lending decision. A lender who estimates the loss independently and prices to a target return on capital will accept some high-coupon transactions and decline others that quote the same coupon, because the coupon required to compensate for the risk differs across transactions even when the headline does not. The analysis is illustrated with modelled figures, a loss-given-default decomposition, a risk-based pricing curve, an underwriting scorecard, and sensitivity and scenario analysis. It is intended for credit and special-situations funds underwriting Gulf real-asset transactions and for allocators assessing the underwriting discipline of the managers they back. All figures are modelled and illustrative, not forecasts, and are calibrated to market structure rather than to any individual transaction.

JEL Classification: G21, G23, G24, G32, G17

Keywords: underwriting, credit risk, probability of default, loss given default, expected loss, risk-based pricing, GCC, United Arab Emirates, private credit, recovery

1. INTRODUCTION

The previous question, where the Gulf's double-digit coupons come from, has a structural answer: they are predominantly compensation for illiquidity, complexity and the supply of scarce capital, with a smaller credit-spread component for expected loss. That answer is reassuring in aggregate but useless on a single transaction, because it describes the market average rather than the deal in front of the lender. The coupon on a specific transaction may adequately price its specific risk, or it may not, and the difference between the two is the difference between a sound loan and a loss waiting to happen. This paper is about that difference. It asks not where the coupon comes from but whether, on any given deal, the coupon is enough.

The discipline the paper advocates inverts the instinct of yield-driven investing. A yield-driven lender starts with the coupon and asks whether it is attractive; an underwriting-driven lender starts with the risk and asks what coupon would be required to compensate for it, then compares that required coupon with the one on offer. If the offered coupon exceeds the required coupon, the transaction creates value; if it falls short, it destroys value however high it looks. The required coupon is the output of an underwriting process, and the quality of that process determines the quality of the lending book. The coupon on offer is a fact; the coupon required is a judgement, and the lender's edge lies in making that judgement well.

Two quantities govern the required coupon: the probability that the borrower defaults, and the fraction of the exposure the lender loses if it does. The product of the two is the expected loss, the through-the-cycle average cost of credit that the coupon must first cover before it earns anything. Beyond expected loss, the lender must hold capital against the possibility that losses exceed the average in a bad year, the unexpected loss, and must earn a return on that capital. The required coupon is therefore the sum of the funding cost, the expected loss, the cost of the capital held against unexpected loss, the operating cost of underwriting and servicing, and the target margin. This paper builds that sum component by component and shows how it varies across transactions.

The stakes are highest precisely where the coupon is highest. The transactions that quote twenty per cent are, by the logic of the companion paper, those with the largest credit-spread component, the most subordinated claims on the weakest collateral and sponsors. These are the transactions where underwriting matters most, because the coupon is genuinely compensation for a credit risk that must be estimated accurately, and where the cost of getting it wrong is greatest. A disciplined framework is not a bureaucratic overlay on these transactions; it is the difference between earning the credit spread and funding the loss it was meant to compensate.

The remainder of the paper is structured as follows. Section 2 reviews the literature on credit underwriting, default and recovery modelling, and risk-based pricing, and develops the propositions the analysis examines. Section 3 sets out the data and the stylised underwriting and pricing methodology. Section 4 presents the results: estimation of default and recovery, the expected-loss matrix, the loss-given-default decomposition, the risk-based pricing curve, and an underwriting scorecard. Section 5 stresses the framework through sensitivity and scenario analysis. Section 6 discusses implementation. Section 7 concludes.

2. LITERATURE REVIEW AND PROPOSITIONS

The framework draws on four bodies of work: the theory and practice of credit underwriting; the modelling of probability of default; the modelling of loss given default and recovery; and the theory of risk-based pricing and economic capital. Each contributes a component to the pricing equation developed in Section 3.

2.1 Credit Underwriting and the Five Cs

The practitioner tradition of credit underwriting, distilled in the long-standing framework of the five Cs, character, capacity, capital, collateral and conditions, remains the backbone of any serious credit assessment. The framework predates formal default modelling and complements it: where statistical models estimate the probability of an event, the five Cs interrogate the borrower-specific factors that drive that probability. In a market such as the Gulf, where transaction-level data are scarce and each deal is idiosyncratic, the qualitative discipline of the five Cs carries proportionately more weight, because there is no large historical dataset on which a purely

statistical model could be trained. The underwriting scorecard developed later in the paper is, in effect, a structured, quantified version of the five Cs adapted to Gulf real-asset lending.

2.2 Modelling the Probability of Default

The academic literature on default probability runs from the structural models of Merton (1974), which derive default from the evolution of a borrower's asset value relative to its liabilities, to the reduced-form models that treat default as a statistical event calibrated to observed data, and to the accounting-ratio models associated with Altman (1968), whose Z-score remains a widely used predictor of corporate distress. For private, collateral-heavy, project-like Gulf transactions, none of these transfers cleanly: there is no traded equity to drive a structural model, no deep default dataset to fit a reduced-form one, and the borrowers are often special-purpose rather than operating companies. The practical response, adopted here, is a structured judgemental estimate of default probability anchored to sponsor quality, asset quality and transaction structure, disciplined by the logic of the formal models rather than mechanically applying them.

2.3 Loss Given Default and Recovery

If default probability is the more uncertain input, loss given default is the more controllable one, because it is driven by the collateral and structure the lender itself negotiates. The literature on recovery, including the work of Altman and co-authors on realised recoveries, emphasises that loss given default is highly variable, depends on seniority and security, and is systematically worse in downturns, precisely when defaults cluster. For real-asset lending the recovery analysis is concrete: it turns on the value of the underlying property, the haircut to that value in a forced sale, the costs and time of enforcement, and the lender's rank in the security structure. Because the lender controls the security package, it controls, within limits, its own loss given default, which is why structuring for recovery is as much a part of underwriting as estimating default.

2.4 Risk-Based Pricing and Economic Capital

The fourth strand is the theory of risk-based pricing, which holds that the price of credit should reflect the risk it carries, covering expected loss, the cost of capital held against unexpected loss, and a target return. The economic-capital framework, developed in the risk-management literature and embedded in bank regulation, distinguishes expected loss, which is a cost to be priced into the coupon, from unexpected loss, which is a risk to be capitalised. A rational lender prices each transaction to earn a target return on the capital it must hold against that transaction's unexpected loss. This is the engine of the pricing framework: the required coupon is whatever delivers the target return on capital after covering expected loss and costs.

2.5 Synthesis and Propositions

Together these literatures imply that the coupon a lender should require is a computable quantity, built from estimates of default, recovery, cost and a target return on capital, and that the offered coupon should be accepted only when it exceeds this required coupon. They further imply that because required coupons vary across transactions with their risk, two transactions quoting the same coupon can have opposite value to the lender. These implications motivate the following propositions.

Proposition 1. The coupon a rational lender requires on a Gulf real-asset transaction is a computable function of the probability of default, the loss given default, the cost of capital and a target return, and can be estimated transaction by transaction.

Proposition 2. Loss given default, governed by collateral and structure, is more controllable by the lender than the probability of default, so structuring for recovery is the highest-leverage underwriting activity.

Proposition 3. Required coupons vary widely across transactions of superficially similar headline yield, so the offered coupon alone is uninformative about whether a transaction creates value.

Proposition 4. A structured underwriting scorecard that quantifies sponsor, asset and structure predicts realised loss better than the coupon does, and should therefore drive the lending decision.

Proposition 5. Because losses cluster in downturns, the capital a lender holds against unexpected loss, and the discipline of pricing to a return on that capital, are what make a high-yield book survivable through a cycle.

2.6 How the Propositions Connect

The propositions describe an underwriting philosophy. Proposition 1 makes the required coupon computable; Proposition 2 locates the lender's greatest leverage in controlling recovery; Proposition 3 warns that the offered coupon is uninformative on its own; Proposition 4 asserts that disciplined underwriting beats yield as a predictor of loss; and Proposition 5 insists that surviving the cycle requires capitalising and pricing for the bad year, not the average one. If they hold, the lending decision should be driven by an underwriting process that estimates loss, prices to a return on capital, and structures for recovery, with the coupon as a check rather than a driver. The sections that follow build and test that process.

2.7 The Special Character of Real-Asset Underwriting

Underwriting a real-asset loan differs in an important respect from underwriting a corporate one. A corporate lender is exposed to the ongoing performance of a business, a stream of future cash flows whose continuation depends on management, markets and competition. A real-asset lender is exposed, in the first instance, to the value of a tangible thing, land, a building, a block of units, that exists whether or not the borrower prospers. This makes recovery analysis more concrete and, in principle, more reliable: the lender can value the collateral, assess its saleability and structure its claim over it, rather than forecasting the fortunes of an enterprise. The flip side is that real-asset values are cyclical and can fall sharply in a downturn, and that the collateral, however valuable on paper, must actually be realised through an enforcement process that takes time and costs money.

This special character shapes the underwriting emphasis. For real-asset credit, the recovery side of the expected-loss calculation is both more important and more controllable than for corporate credit, because the collateral is the primary source of repayment in stress and the lender negotiates its claim over it. The default side, by contrast, depends heavily on factors outside the lender's control: the property cycle, the sponsor's other commitments, the availability of refinancing. The framework reflects this by treating recovery as the variable the lender should work hardest to control through structure, and default as the variable it should assess most carefully but cannot fully manage.

2.8 Information Asymmetry and the Sponsor Relationship

A persistent feature of private lending is information asymmetry: the sponsor knows more about the asset, the plan and its own circumstances than the lender does. The literature on lending relationships shows that this asymmetry is mitigated by repeated dealing, by reputation and by the alignment that equity contribution and personal guarantees create. In a young market with limited public information, the relationship dimension carries more weight than in a transparent one, because the lender cannot fall back on disclosed financials, ratings or a deep secondary market for price discovery. The sponsor's track record, reputation and willingness to stand behind a transaction become primary inputs to the default assessment.

This has a direct underwriting implication. The lender should price information asymmetry as a risk and reduce it through structure: meaningful sponsor equity that aligns incentives, guarantees that put the sponsor's own resources at risk, reporting obligations that narrow the information gap over the life of the loan, and a preference for sponsors with whom the lender has dealt before. A transaction with a strong, aligned, well-known sponsor warrants a lower default estimate than an identical transaction with an unknown one, and the difference is not sentiment but a rational response to the information the lender does and does not have.

2.9 Why Capital, Not Just Loss, Must Be Priced

A subtle but consequential point in the framework is the distinction between pricing for expected loss and pricing for capital. A lender that charges a coupon sufficient to cover only its funding cost and its average loss will appear profitable in normal years and will fail in a bad one, because it has held no compensation for the capital that absorbs losses above the average. The capital is real, it must be raised, held idle against the possibility of a bad year, and remunerated, and the coupon must pay for it. Omitting the capital charge is the single most common way that high-yield lenders flatter their economics, because expected loss is visible and modest in good years while the cost of capital is easy to ignore until it is needed.

Pricing for capital also changes the relative attractiveness of transactions. Two transactions with the same expected loss can require very different capital if their unexpected losses differ, that is, if one is far more likely than the other to lose much more than average in a stressed year. The transaction with the fatter tail requires more capital and therefore a higher coupon, even though its expected loss is identical. A framework that prices only expected loss would treat the two as equivalent and would systematically under-price the riskier one. Pricing for capital corrects this, and it is why the unexpected-loss buffer, not just the expected loss, belongs in the required-coupon equation.

3. DATA AND METHODOLOGY

The study is a structured framework rather than a fitted statistical model, for the reason given above: Gulf real-asset credit lacks the large, clean default dataset that statistical estimation requires. The methodology is therefore a transparent, assumption-driven pricing engine whose every output can be traced to a stated input, and whose inputs are calibrated to the structure of the market rather than to proprietary data.

3.1 Data and Inputs

Qualitative inputs, the structure of Gulf security packages, enforcement processes and the behaviour of sponsors and assets, are drawn from the established characteristics of the market. Quantitative inputs, default probabilities by tier, recovery components, costs and target returns,

are set at the centre of plausible ranges and clearly labelled as illustrative. Single figures should be read as midpoints, and Section 5 varies them across their ranges. No figure is a forecast for any specific transaction.

3.2 The Pricing Equation

The core of the methodology is the required-coupon equation. The required coupon equals the funding cost plus the expected loss plus the cost of capital held against unexpected loss plus operating cost plus target margin. Expected loss is the probability of default multiplied by the loss given default. The capital charge is the unexpected loss, the amount by which losses could exceed the expected in a stressed year, multiplied by the lender's target return on capital. The equation makes explicit that the coupon must do four jobs in sequence: return the lender's funding, cover the average loss, pay for the capital that absorbs the bad-year loss, and leave a margin. A coupon that covers only the first two is not profitable once capital is priced; a coupon that covers none is a subsidy to the borrower.

3.3 Estimating the Probability of Default

Default probability is estimated through a structured assessment of three factors: sponsor quality, the borrower's track record, financial strength and alignment; asset quality, the saleability, location and completion status of the underlying property; and transaction structure, the leverage, tenor and covenants. Each factor is scored, and the scores map to a default-probability range. The approach is judgemental but disciplined: it forces the underwriter to assess the drivers of default explicitly rather than inferring them from the coupon, and it produces an estimate independent of the price on offer, which is essential if the estimate is to serve as a check on that price.

3.4 Estimating Loss Given Default

Loss given default is built from the collateral upward. Starting from the gross value of the security, the framework applies a valuation haircut to reflect uncertainty in the appraisal, a forced-sale discount to reflect the difference between an orderly and a distressed sale, enforcement costs, legal, agency and carrying, and a time discount to reflect the months or years between default and realisation. What remains is the net recovery; one minus the net recovery, relative to the claim, is the loss given default. Because each of these deductions is influenced by the security the lender negotiates and the enforceability of its claim, the loss given default is, within limits, a variable the lender can manage rather than a fixed property of the borrower.

3.5 Base-Case Assumptions

Table 1 sets out the central assumptions. They are conservative where conservatism matters: default probabilities are set at the higher end of plausible ranges and recoveries at the lower end, so that the required coupons the framework produces are if anything too high, and a transaction that clears the framework clears it with margin to spare. The full parameter set and the ranges used in the robustness analysis appear in Appendix A.

Table 1. Base-Case Assumptions for the Underwriting Framework

Parameter	Base-case assumption
Funding cost	5.0% (dollar funding for a non-bank lender)
Probability of default by tier	2%-12% (Tier 1 to Tier 5)
Loss given default by tier	15%-60%

Valuation haircut	10%-15% of appraised value
Forced-sale discount	10%-20%
Enforcement cost	5%-10% of claim
Enforcement time	6-30 months
Capital held vs unexpected loss	2x-3x expected loss
Target return on capital	15%
Operating / servicing cost	0.8% of assets

Indicative, illustrative assumptions used to drive the pricing framework; not forecasts and not drawn from any single transaction.


4. RESULTS AND DISCUSSION

This section builds the pricing framework from its components: the expected-loss matrix, the loss-given-default decomposition, the risk-based required-coupon curve, the underwriting scorecard, and the relationship between underwriting quality and realised loss.

4.1 The Expected-Loss Matrix

Figure 1 presents expected loss as a matrix over probability of default and loss given default. Reading the matrix makes the first discipline of underwriting concrete: expected loss is the product of two quantities, and a low value of either keeps the expected loss modest even when the other is high. A transaction with a high default probability but strong collateral, low loss given default, can have a lower expected loss than one with a moderate default probability but weak collateral. The matrix warns against assessing either dimension in isolation, and it locates the dangerous transactions in the bottom-right, high default probability and high loss given default together, where expected loss escalates fastest.

Figure 1. The Expected-Loss Matrix: PD times LGD




Expected loss is the product of default probability and loss given default; strong collateral can offset high default risk.

The matrix also frames the lender's task. Default probability is largely a property of the borrower and the environment, which the lender assesses but does not control; loss given default is largely a property of the collateral and structure, which the lender negotiates. Moving a transaction leftward in the matrix, reducing default probability, is mostly a matter of selection; moving it upward to a lower loss given default is a matter of structuring. The lender's leverage, per Proposition 2, lies disproportionately in the second.

4.2 Decomposing Loss Given Default

Figure 2 decomposes loss given default for a representative secured transaction, building from the gross collateral value down to the net recovery through the valuation haircut, the forced-sale discount, enforcement costs and the time discount. The decomposition reveals how a loan that looks comfortably over-collateralised on day one can deliver a disappointing recovery in default, because each deduction erodes the cushion, and the deductions compound. A nominal loan-to-value that looks safe can translate into a material loss once the haircuts, discounts, costs and delays of a real enforcement are applied.

Figure 2. Building Loss Given Default: From Collateral to Net Recovery



Each deduction, haircut, sale discount, enforcement cost and time, erodes the collateral cushion; they compound.


The practical lesson is that the headline loan-to-value overstates the protection. A lender that underwrites to the appraised value, without applying the haircuts and discounts that a real enforcement imposes, will systematically underestimate its loss given default and so under-price its credit. The disciplined lender underwrites to the net recovery, not the gross collateral, and structures the security to minimise each deduction: faster enforcement, clearer title, first-ranking charges and cash control all raise net recovery, and each is a lever the lender can pull at origination.

4.3 Recovery and Loan-to-Value

Figure 3 (presented after the recovery analysis) and Figure 7 together show how net recovery falls as loan-to-value rises. At low loan-to-value the collateral cushion absorbs the haircuts and discounts and recovery is near complete; as loan-to-value rises, the cushion thins and recovery falls, slowly at first and then sharply once the cushion is exhausted. The relationship is non-

linear, which is why marginal increases in leverage at high loan-to-value are so dangerous: each additional point of leverage removes a point of cushion exactly where the cushion is most needed. The lender's loan-to-value discipline is, in effect, its loss-given-default discipline.

Figure 3. Net Recovery as a Function of Loan-to-Value




Recovery is near-complete at low leverage and falls sharply once the collateral cushion is exhausted.

4.4 The Risk-Based Required-Coupon Curve

Figure 8 presents the central output of the framework: the coupon a rational lender should require as a function of the probability of default, decomposed into the credit-spread component and the rest. The curve rises with default probability, as it must, because higher default risk demands higher compensation; and the credit-spread component, the part that compensates specifically for loss, rises faster than the total, because at higher default probabilities more of the coupon must go to covering and capitalising loss rather than to margin. The curve is the lender's pricing discipline made visible: for any estimated default probability, it gives the coupon required, against which the offered coupon can be judged.


Figure 4. The Risk-Based Pricing Curve



Required coupon rises with default probability; the credit-spread component rises faster than the total.

Figure 3 translates the curve into discrete underwriting tiers, showing the required coupon for each of five risk tiers from low to high. The tiering is how the framework is used in practice: a transaction is assessed, assigned to a tier, and priced to the required coupon for that tier. A transaction offered at or above its tier's required coupon creates value; one offered below it, however high the absolute coupon, does not. This is the operational form of Proposition 3: the same offered coupon can clear the bar for a low tier and fall short for a high one, so the offered coupon is meaningless without the tier.

Figure 5. Risk-Based Required Coupon by Underwriting Tier




Each tier carries a required coupon; an offered coupon creates value only if it clears its tier's bar.

4.5 Default Probability by Sponsor and Asset

Figure 5 shows how estimated default probability varies with the two factors the underwriter assesses most carefully: sponsor quality and asset quality. The strongest sponsors with prime assets carry the lowest default probabilities; the weakest sponsors with secondary assets carry the

highest, several times greater. The figure quantifies the intuition that underwriting is largely about these two factors, and it shows their interaction: a strong sponsor mitigates a weaker asset, and a prime asset mitigates a weaker sponsor, but the combination of weak sponsor and secondary asset compounds into a default probability that few coupons can adequately price.

Figure 6. Default Probability by Sponsor and Asset Quality



Default risk is lowest for strong sponsors with prime assets and compounds for weak sponsors with secondary assets.

4.6 The Underwriting Scorecard

Table 2 sets out an underwriting scorecard that quantifies the assessment of a transaction across the factors that drive default and recovery. Each factor is scored and weighted, and the total score maps to a risk tier and hence to a required coupon. The scorecard is the structured, quantified form of the five Cs adapted to Gulf real-asset lending. Its purpose is not to replace judgement but to discipline it: by forcing every transaction through the same factors with the same weights, it makes underwriting consistent, comparable and auditable, and it produces a risk assessment independent of the coupon on offer.

Table 2. Illustrative Underwriting Scorecard

Factor	Weight	What it captures	Score range
Sponsor quality	25%	Track record, financial strength, alignment	1-5
Asset quality	20%	Location, saleability, completion status	1-5
Security & structure	20%	Rank, charges, controls, enforceability	1-5
Leverage (LTV)	15%	Collateral cushion to net recovery	1-5
Cash flow / exit	10%	Repayment source and refinancing risk	1-5
Conditions / market	10%	Property-cycle and	1-5


		liquidity backdrop	
--	--	--------------------	--

Illustrative scorecard; weighted score maps to a risk tier and required coupon. Weights and ranges are stylised.

4.7 Underwriting Quality and Realised Loss

Figure 6 plots underwriting score against realised loss across a set of illustrative transactions. The relationship is strongly negative: higher-scoring transactions experience lower losses, and the relationship is tighter than any relationship between the coupon and realised loss would be. This is the evidence for Proposition 4. The coupon is a price, set partly by market conditions and negotiation; the underwriting score is an assessment of risk, and it is the assessment, not the price, that predicts loss. A lender that selects on score will experience lower losses than one that selects on coupon, even if both books carry the same average coupon.

Figure 7. Underwriting Score versus Realised Loss




Higher underwriting scores predict lower realised losses; the score predicts loss better than the coupon does.

4.8 Expected and Unexpected Loss by Tier

Figure 9 shows, for each tier, the expected loss and the buffer of unexpected loss against which the lender must hold capital. The expected loss is priced into the coupon; the unexpected loss is capitalised. Both rise with the tier, but the unexpected-loss buffer rises faster, because higher-risk transactions are not only more likely to lose but more likely to lose much more than average in a bad year. This is why pricing to a return on capital matters most for the highest tiers: it is there that the capital charge, the cost of holding the buffer, is largest, and there that a lender who prices only for expected loss most badly under-prices the risk.

Figure 8. Expected and Unexpected Loss by Tier



Both rise with the tier, but the unexpected-loss buffer rises faster, raising the capital cost of high-tier lending.

4.9 Reading the Results Together

The results support the propositions. The required coupon is computable from default, recovery, capital and target return (Proposition 1); loss given default is built from controllable components (Proposition 2); required coupons vary across tiers so the offered coupon is uninformative alone (Proposition 3); the underwriting score predicts loss better than the coupon (Proposition 4); and the unexpected-loss buffer, largest for high tiers, is what the capital charge must cover for the book to survive a cycle (Proposition 5). Together they describe a lending process driven by underwriting, with the coupon as a check.

4.10 A Worked Pricing Example

It is worth working a single transaction through the framework to show how the pieces combine. Consider a mezzanine loan to a mid-quality sponsor on a prime asset at a moderate loan-to-value. The structured assessment places the sponsor in the middle band and the asset in the top band, yielding an estimated default probability of around six per cent. The collateral analysis, after a valuation haircut, a forced-sale discount, enforcement costs and a time discount, yields a net recovery that implies a loss given default of around forty per cent. The expected loss is therefore roughly two and a half per cent. On top of the five per cent funding cost, the framework adds this expected loss, the cost of capital held against an unexpected loss several times larger, the operating cost and the target margin, arriving at a required coupon in the mid-teens.

Suppose the transaction is offered at sixteen and a half per cent. The comparison is now concrete: the offered coupon exceeds the required coupon by a margin, and the transaction creates value. Had the same transaction been offered at thirteen per cent, it would have fallen short of the required coupon despite being, in absolute terms, a high-yield loan, and the framework would have counselled declining it or repricing it. The worked example shows the framework doing its job: converting an assessment of risk into a required price, and converting the accept-or-decline decision from a reaction to the headline into a comparison of two numbers.

4.11 The Margin of Safety

A disciplined lender does not fund transactions that merely clear the required coupon by a hair; it seeks a margin of safety, a cushion between the offered and required coupons that protects against errors in the estimates. The margin of safety matters because the inputs, especially default probability, are estimated with uncertainty, and a transaction priced exactly at its required coupon under central estimates may be under-priced if the estimates prove optimistic. By requiring a margin, the lender builds in protection against its own estimation error, and it concentrates its capital in the transactions where the price most comfortably exceeds the risk. The size of the margin should reflect the confidence in the estimates: wider for novel, hard-to-assess transactions, narrower for familiar ones with strong data.

The margin of safety also disciplines portfolio growth. In a competitive market the temptation is to thin the margin to win more transactions and deploy capital faster; the framework counsels the opposite, holding the margin constant and accepting slower deployment when prices tighten. A lender that maintains its margin through a competitive phase will deploy less but lose less; one that thins its margin to keep deploying will grow its book and its losses together. The margin of safety is the mechanism by which underwriting discipline survives commercial pressure.

4.12 Tier Migration and Repricing

A transaction's risk is not static, and the framework must accommodate migration between tiers over the life of a loan. A construction project that falls behind schedule, a sponsor whose other commitments deteriorate, or a property market that softens can move a transaction from a lower tier to a higher one, raising its expected loss and the capital it consumes. Where the loan structure permits, the lender should reprice or restructure to reflect the higher risk, through margin ratchets, additional security, or amortisation that reduces the exposure. Where it does not, the lender bears the migration as a loss of margin, which is itself an argument for building repricing and protective mechanisms into the structure at origination.

Migration also runs in the favourable direction. A project that completes on time, a sponsor that strengthens, or a market that firms can move a transaction to a lower tier, reducing its risk and freeing capital. A lender that monitors actively can recognise favourable migration and redeploy the freed capital, improving the efficiency of its book. The discipline is symmetric: the same monitoring that detects deterioration and triggers protective action also detects improvement and frees capital, and a lender that tracks tier migration in both directions manages its capital more efficiently than one that sets the tier at origination and forgets it.

4.13 The Limits of Quantification

For all its apparatus of probabilities, recoveries and required coupons, the framework rests on judgements that resist full quantification, and honesty requires acknowledging where the numbers run out. The estimate of default probability for a special-purpose borrower with no trading history is, ultimately, a structured opinion about a sponsor, an asset and a plan, dressed in the language of probability. The estimate of recovery depends on how a hypothetical future enforcement plays out in a jurisdiction where each case has its own complications. The framework's value is not that it removes judgement but that it organises it, forces it to be explicit, makes it consistent across transactions, and exposes it to challenge and to subsequent comparison with outcomes.

This is a strength rather than a weakness, provided it is understood. A spurious precision, treating a judgemental six per cent default estimate as if it were measured to the basis point, would be dangerous, because it would invite false confidence and discourage the margin of safety that


estimation uncertainty demands. The disciplined user of the framework holds two ideas at once: the numbers are the best available organisation of the available judgement, and they are uncertain enough to require conservative calibration and a margin of safety. Quantification serves judgement; it does not supersede it, and a framework that forgot this would underwrite worse, not better, than one that kept judgement in its proper place.

5. ROBUSTNESS AND SCENARIO ANALYSIS

The framework's outputs depend on its inputs, and its credibility rests on how those outputs respond to plausible variation. This section stresses the pricing inputs, examines a downturn, and states the conditions under which the discipline fails.

5.1 Sensitivity of the Required Coupon

Figure 4 reports the sensitivity of the required coupon to its inputs. Loss given default and probability of default dominate, as the pricing equation implies, because together they determine the expected loss and the capital charge that the coupon must cover. The base rate matters because it sets the funding cost; the illiquidity premium, origination cost and target return matter less. The tornado instructs the underwriter to estimate default and recovery with the greatest care, because errors there translate most directly into mispricing, and to treat the other inputs as second-order.



Loss given default and default probability dominate; errors there translate most directly into mispricing.

Because recovery is both high-impact and controllable, the sensitivity analysis reinforces the structuring imperative. A lender uncertain about its recovery estimate should resolve the uncertainty by improving the security, faster enforcement, better rank, tighter controls, rather than by raising the coupon, because improving recovery reduces both the expected loss and the capital charge, while raising the coupon only compensates for a risk left in place. Structuring dominates pricing as a response to recovery risk.

5.2 The Downturn Scenario

The scenario that tests the framework is a property downturn in which default probabilities rise, recoveries fall, and enforcement lengthens, all together and across the book at once. Table 3 shows the effect on a tiered book. The low tiers, underwritten to strong sponsors, prime assets and conservative leverage, see expected losses rise but remain well within the capital held against them; the high tiers see expected losses rise sharply and can exhaust the capital allocated to them. The scenario validates the tiering: the capital charge that looked expensive for the high tiers in good times is exactly what absorbs their losses in bad ones, and a lender that skimmed on capital for the high tiers to flatter its returns would find itself under-capitalised when it mattered.

Table 3. Loss Under Base and Downturn Scenarios by Tier (Illustrative)

Tier	Base EL	Downturn EL	Capital held	Covered?
Tier 1	0.5%	1.2%	1.5%	Yes
Tier 2	1.0%	2.4%	2.8%	Yes
Tier 3	2.0%	4.6%	4.8%	Marginal
Tier 4	3.5%	7.8%	7.5%	Stretched
Tier 5	5.5%	12.5%	11.5%	Breached

Illustrative downturn: default rates roughly double and recoveries fall. High tiers can breach the capital held against them.

The scenario carries a portfolio lesson beyond the individual transaction. Because losses cluster in a downturn, a book concentrated in the high tiers is exposed to a correlated shock that the average-loss pricing does not capture; the unexpected-loss buffer must be sized for the correlated bad year, not the independent average. A disciplined high-yield lender therefore limits its concentration in the highest tiers, holds capital against the correlated scenario, and prices the high tiers to a return on that larger capital, accepting that this makes some high-coupon transactions uneconomic. That is the framework working as intended: declining the transactions whose coupon does not compensate for the capital their risk consumes.

5.3 When the Discipline Fails

The discipline fails under three conditions. It fails if default and recovery are estimated with systematic optimism, so that the required coupons are too low and transactions clear the bar that should not; the remedy is conservative calibration and honest post-mortems on realised losses. It fails if the scorecard is applied mechanically, as a box-ticking exercise rather than a structured judgement, so that it produces a score without insight; the remedy is to treat the scorecard as a discipline on judgement, not a substitute for it. And it fails if the lender prices to expected loss alone, ignoring the capital charge, so that the book looks profitable until the bad year arrives; the remedy is to price every transaction to a return on the capital its risk requires.

5.4 Robustness of the Central Conclusion

Across the variations considered, the central conclusion, that the lending decision should be driven by an underwriting process that estimates loss, prices to a return on capital and structures for recovery, with the coupon as a check rather than a driver, holds robustly. The framework is conservative by construction, and a transaction that clears it under conservative inputs clears it with margin. The conclusion is most valuable precisely for the highest-coupon transactions, where the temptation to let the yield drive the decision is greatest and the cost of doing so is highest.

5.5 Correlated Recovery Risk

The downturn scenario assumes that defaults rise and recoveries fall together, and it is worth dwelling on why this correlation is so damaging. In benign conditions, a default is usually idiosyncratic, a particular sponsor's difficulty, and the collateral can be sold into a functioning market at close to its appraised value, so recovery is high. In a downturn, defaults are driven by a common factor, a falling property market, and the same factor depresses the value of the collateral and lengthens the time to sell it. Default probability and loss given default, which are largely independent in good times, become correlated in bad ones, and the expected loss computed on the assumption of independence understates the loss that actually occurs when the two move together.

The implication is that a high-yield book must be capitalised and priced for the correlated scenario, not the independent average. This is why the unexpected-loss buffer is larger than a naive reading of historical average losses would suggest, and why the capital charge weighs most heavily on the high tiers, where the correlation bites hardest. A lender that prices each transaction in isolation, ignoring the correlation that links its recoveries in a downturn, will hold too little capital and price too little spread, and will discover the shortfall precisely when it can least afford to. Pricing and capitalising for correlated recovery risk is the discipline that separates a high-yield book that survives a cycle from one that does not.

5.6 Benchmarking Required Coupons Against the Market

A useful external check on the framework is to compare the required coupons it produces with the coupons the market actually quotes. Where the framework's required coupon sits below the market coupon for a given tier, the market is offering more than the risk requires, evidence of the structural premium documented in the companion paper, and the lender should be an enthusiastic participant. Where the framework's required coupon sits above the market coupon, the market is under-pricing the risk, and the lender should decline, however tempting the participation. The comparison turns the framework into a market-discipline tool: it tells the lender not only what to charge but when the market as a whole is mispricing risk, in either direction.

Over time, tracking the gap between required and market coupons reveals how the market is evolving. A widening gap in the lender's favour signals that the structural premium is intact or growing; a narrowing or reversing gap signals that competition is compressing the premium and that the lender should expect lower returns or migrate to less contested tiers. The framework thus doubles as an early-warning system for the maturation of the market, the process that the companion paper on the coupon predicted would eventually compress the Gulf premium. A lender that watches the gap will see the compression coming before it shows up in realised returns.

5.7 Stress-Testing the Whole Book

Beyond stressing individual transactions, the lender should stress the book as a whole, because the risks that matter most are those that strike many transactions at once. A book-level stress test applies a common shock, a property downturn of a given severity, to every transaction simultaneously, allowing defaults to rise and recoveries to fall together, and aggregates the resulting losses against the capital held. The exercise reveals whether the book, not merely each loan, is adequately capitalised for a correlated event, and it exposes concentrations, in a sponsor,

an asset type, a sub-region, that individual transaction analysis misses. A book that looks comfortable loan by loan can be dangerously exposed in aggregate if its risks are correlated.

The book-level stress test should drive concentration limits. If the aggregate loss under a plausible downturn exceeds the capital the lender is willing to put at risk, the remedy is not to raise coupons but to reduce concentration, capping exposure to any single sponsor, asset type or sub-market so that no single shock can overwhelm the book. The discipline mirrors at the portfolio level the discipline applied at the transaction level: estimate the loss under stress, hold capital against it, and decline or trim the exposures whose contribution to aggregate risk the available capital cannot support. A high-yield book survives a cycle not because each loan is safe but because the book as a whole is sized and diversified to absorb the correlated bad year.

6. IMPLEMENTATION CONSIDERATIONS

Turning the framework into a live underwriting capability requires decisions about process, people, data and governance.

6.1 Building the Underwriting Process

The process should run every transaction through the same sequence: structured assessment of sponsor, asset and structure; independent estimation of default and recovery; computation of the required coupon; and comparison with the offered coupon. The sequence must be followed in order, with the risk assessment completed before the coupon is considered, so that the assessment is not contaminated by the price. A process that lets the underwriter see the coupon first will, consciously or not, reverse-engineer an assessment that justifies it, which defeats the purpose.

6.2 Structuring for Recovery

Because recovery is the highest-leverage variable, the structuring function deserves the most senior attention. The lender should seek first-ranking security, clear and enforceable title, completion and cost-overrun protections, cash-flow control and step-in rights, and should map the enforcement path before funding, not after default. Each structural feature that raises net recovery lowers both the expected loss and the capital charge, and so improves the economics more than an equivalent increase in the coupon. Structuring is examined in depth in the companion paper on security and enforcement.

6.3 Data and Calibration

A young market provides little historical loss data, which is a genuine constraint on calibration. The disciplined response is to calibrate conservatively, to record every transaction's assessment and outcome systematically so that a proprietary dataset accumulates, and to update the calibration as realised losses provide evidence. A lender that captures its own loss experience rigorously will, over a few years, hold a dataset that no newcomer can replicate, and that dataset is itself a source of edge. The absence of market data is a barrier to entry that rewards the lender who builds its own.

6.4 People and Judgement

The framework disciplines judgement; it does not replace it. The estimation of default probability, in particular, rests on experienced assessment of sponsors, assets and structures that no model can fully capture in a data-scarce market. The lender's underwriting talent, its ability to read a sponsor, value an asset and anticipate how a structure behaves in stress, is the ultimate

determinant of loss experience. The framework ensures that talent is applied consistently and checked against the coupon; it does not manufacture the talent, which must be hired, developed and retained.

6.5 Governance and the Credit Committee

The credit committee is where the framework becomes governance. A committee presented with a transaction should see the underwriting score, the estimated default and recovery, the required coupon, the offered coupon and the margin between them, and should make its decision on that basis rather than on the headline yield. Presenting transactions in this decomposed form disciplines the committee as the scorecard disciplines the underwriter, and it creates an auditable record of why each transaction was approved at the price it was. Over time the record of decisions against outcomes is the institution's most valuable underwriting asset.

6.6 Pricing to Win Without Mispricing

A practical tension in any competitive lending market is between pricing to win the transaction and pricing to compensate for the risk. The framework resolves the tension by making the required coupon explicit: the lender may choose to price at or slightly above the required coupon to win a transaction whose risk it likes, but it should never price below the required coupon, because to do so is to lend at a loss in expectation. A lender that loses transactions because competitors price below the required coupon is not losing business it should regret; it is declining to fund losses that its competitors will eventually realise.

6.7 For Allocators: Assessing a Manager's Underwriting

For the allocator backing a Gulf credit manager rather than lending directly, the framework is a diligence tool. The allocator should ask the manager to demonstrate its underwriting process: how it estimates default and recovery, how it computes a required coupon, how it structures for recovery, how it capitalises unexpected loss, and how its realised losses compare with its underwriting scores. A manager that can answer these questions in decomposed, evidenced terms has the discipline the high-yield book requires; one that points to its coupons as evidence of its skill has confused price with risk, and is the manager most likely to deliver the loss the coupon was meant to compensate.

6.8 Common Underwriting Errors

Several errors recur and the framework is built to prevent them. The first is underwriting to the appraised value rather than the net recovery, which understates loss given default and underprices the credit. The second is letting the coupon drive the assessment, which reverses the logic of underwriting. The third is pricing to expected loss while ignoring the capital charge, which flatters returns until the bad year. The fourth is treating the scorecard mechanically, which produces a number without judgement. The fifth is concentrating in the high tiers because their coupons are attractive, which builds a book exposed to the correlated downturn. Each error is a failure of discipline, and each is corrected by following the process the framework lays out.

6.9 From Underwriting to Monitoring

Underwriting does not end at funding; the assessment made at origination must be maintained over the life of the loan, because the factors that drive default and recovery evolve. A sponsor's circumstances change, an asset's market shifts, a construction project slips, and each change alters the risk the lender bears. A disciplined lender therefore treats monitoring as continuous

underwriting: it tracks the covenants, the project milestones, the sponsor's position and the market, and it updates its default and recovery estimates as conditions change. The earlier a deterioration is detected, the more options the lender has to protect its position, restructure, enforce, or exit, and the higher its eventual recovery.

Monitoring is also where the security and controls negotiated at origination earn their keep. Cash-flow control, reporting obligations, milestone conditions and step-in rights are valuable not as paperwork but as the means by which the lender detects and responds to deterioration. A lender that structures these controls but does not use them has paid for protection it does not collect; one that monitors actively converts the structural features of the loan into early warning and effective intervention. The companion paper on security and enforcement examines how these controls translate into recovery in practice.

6.10 Speed Versus Rigour

A recurring tension in this market is between the speed sponsors demand and the rigour underwriting requires. Much of the lender's premium, as the companion paper on the coupon argued, is compensation for providing capital quickly when banks cannot, and a lender that takes months to underwrite forfeits the speed that is part of its value. Yet rigour cannot be sacrificed to speed without inviting the losses that discipline is meant to prevent. The resolution is not to choose between them but to build a process that is both fast and rigorous: a structured, well-resourced underwriting capability that can assess a transaction thoroughly in days rather than weeks.

Achieving this requires investment in the underwriting function, experienced people, clear processes, ready access to valuation and legal expertise, and a credit committee that can convene quickly. It also requires selectivity at the top of the funnel, so that scarce underwriting capacity is spent on transactions likely to clear the bar rather than dissipated across long-shot inquiries. A lender that can underwrite well and quickly captures the speed premium without bearing the losses that haste would otherwise bring; this combination of speed and rigour is itself a source of competitive edge in a market where few possess both.

6.11 Documenting and Learning from the Decision

The final implementation point concerns institutional learning. Every transaction the lender underwrites is an opportunity to improve the framework, but only if the assessment and its outcome are documented and compared. The lender should record, for each transaction, the estimated default and recovery, the required and offered coupons, the tier and the decision, and should later record the actual outcome. Comparing assessments with outcomes reveals where the estimates were systematically optimistic or pessimistic, and allows the framework to be recalibrated against the lender's own experience, the proprietary dataset that a young market otherwise denies.

This discipline of documentation and review is what converts a lending business into a learning institution. A lender that funds transactions and forgets the reasoning behind them repeats its mistakes; one that records its reasoning and confronts it with outcomes improves its underwriting year by year and accumulates an edge that capital alone cannot buy. In a market without deep public data, the lender's own documented experience is the most valuable dataset available, and the discipline of building it is as important to long-run returns as the underwriting of any individual transaction.

6.12 Aligning Incentives Inside the Lender

A final implementation consideration is internal: the incentives of the people who originate and underwrite must be aligned with the long-run performance of the book, not with short-run deployment. An originator rewarded for volume will push transactions through; an underwriter rewarded for caution alone will block good ones. The lender must structure incentives so that the organisation as a whole is rewarded for funding transactions where the offered coupon comfortably exceeds the required coupon and for declining those where it does not, with reward tied to realised outcomes over time rather than to deployment in the moment. Deferred and outcome-linked compensation, shared accountability between origination and credit, and a culture that celebrates a well-judged decline as much as a completed deal all serve this alignment.

Misaligned incentives are how disciplined frameworks fail in practice. A lender can build an impeccable required-coupon model and still accumulate losses if its people are paid to deploy regardless of price, because they will find ways to make weak transactions appear to clear the bar. The framework's discipline survives only in an organisation whose incentives reward the discipline, which is why governance and compensation are not peripheral to underwriting but central to it. The most important control on a high-yield book is not the model but the alignment of the people who feed it, and a lender that gets the incentives right will out-perform one that has a better model but rewards the wrong behaviour.

7. CONCLUDING COMMENTS

This paper has set out a disciplined framework for underwriting and pricing high-yield Gulf real-asset credit, built around the two quantities that determine whether a coupon is sufficient: the probability of default and the loss given default. From these, together with the cost of capital and a target return, it derives the coupon a rational lender should require, and it argues that the lending decision should be driven by this required coupon, with the offered coupon serving as a check rather than a driver.

The central message is that underwriting, not yield, should drive lending. A high coupon is a question, whether it adequately prices the risk, and the answer differs across transactions that quote the same headline. A lender that estimates loss independently, structures to control recovery, capitalises unexpected loss and prices to a return on that capital will build a book whose realised losses are low and whose returns survive a cycle. A lender that lets the coupon drive the decision will, sooner or later, fund the transaction whose coupon was genuine compensation for a loss that duly arrives.

The framework is properly qualified. It is a structured, assumption-driven engine rather than a fitted statistical model, because the market lacks the data such a model would require; its figures are illustrative, not forecasts; and the estimation of default probability rests on judgement that no framework can fully systematise. These are limitations, but they are also the conditions that make disciplined underwriting valuable: in a data-rich, efficient market the edge would be competed away, while in a data-scarce, segmented one it accrues to the lender who underwrites best.

The study suggests several extensions. As realised loss data accumulate, the judgemental default and recovery estimates can be calibrated against outcomes, tightening the framework. The correlation structure of losses across the book, treated here through scenario analysis, could be modelled explicitly to size the unexpected-loss buffer more precisely. And the scorecard could be

validated against realised losses to refine its weights. Each extension would sharpen the framework without changing its philosophy.

The framework also reframes what competitive advantage means in this market. In a mature, data-rich market, advantage comes from scale, cost of funding and access to flow; in a young, data-scarce one, it comes from underwriting judgement, the discipline to price for capital and not just loss, and the patience to decline transactions that competitors fund. These are capabilities a lender builds rather than buys, and they compound: each year of disciplined underwriting adds to the proprietary loss dataset, sharpens the calibration, and widens the gap between the disciplined lender and the yield-chaser. The advantage is durable precisely because it cannot be replicated with capital alone.

Finally, the framework should be read alongside its companions. The paper on the coupon explains why the Gulf premium exists and why it is mostly compensation for illiquidity and complexity; this paper explains how to price the credit risk that remains; the paper on security and enforcement examines how recovery is actually realised; the paper on loss experience asks whether the yield survives realised defaults; and the paper on portfolio construction shows how the individual transactions combine into a book. Together they describe a complete discipline for investing in Gulf private credit, of which sound underwriting, the subject of this paper, is the indispensable core.

In sum, the paper's contribution is to insist that a double-digit coupon be earned through underwriting rather than assumed through yield. The coupon is where the analysis begins, not where it ends; the work is to price the risk behind it, and to fund only those transactions where the price on offer comfortably exceeds the price the risk requires.

DISCLOSURES

The author is an independent researcher and transaction adviser who advises capital providers and capital seekers on real-asset and private-credit transactions in the Gulf and may have a commercial interest in the matters discussed. This paper is provided for information and education only. It is not investment, legal, tax or accounting advice, and it is not an offer or solicitation in respect of any security, loan or transaction. All figures are modelled and illustrative, calibrated to market structure rather than to any individual transaction, and are not forecasts. Private credit involves the risk of loss, including loss of principal, and illustrative performance is not indicative of future results. Readers should obtain independent professional advice before acting.

References

- [1] Altman, E. I. (1968). Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy. *Journal of Finance*, 23(4), 589-609.
- [2] Altman, E. I., Brady, B., Resti, A., and Sironi, A. (2005). The Link between Default and Recovery Rates. *Journal of Business*, 78(6), 2203-2228.
- [3] Basel Committee on Banking Supervision (2017). Basel III: Finalising Post-Crisis Reforms. Bank for International Settlements.
- [4] Crouhy, M., Galai, D., and Mark, R. (2000). A Comparative Analysis of Current Credit Risk Models. *Journal of Banking and Finance*, 24(1-2), 59-117.
- [5] Duffie, D., and Singleton, K. J. (2003). *Credit Risk: Pricing, Measurement, and Management*. Princeton University Press.
- [6] Gupton, G. M., Finger, C. C., and Bhatia, M. (1997). *CreditMetrics Technical Document*. J.P. Morgan.

[7] Merton, R. C. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. Journal of Finance, 29(2), 449-470.

[8] Schuermann, T. (2004). What Do We Know About Loss Given Default? Wharton Financial Institutions Center Working Paper.

[9] Saunders, A., and Allen, L. (2010). Credit Risk Measurement In and Out of the Financial Crisis. Wiley Finance.

[10] Vasicek, O. (2002). The Distribution of Loan Portfolio Value. Risk, 15(12), 160-162.

[11] International Monetary Fund (2024). Global Financial Stability Report: The Rise of Private Credit. IMF, Washington, D.C.

About the Author

Chennakeshav Adya is an independent researcher with more than 25 years of industry experience as both an operator and an investment banker. His work focuses on capital formation, private credit and real-asset finance in the Gulf and on connecting international institutional and private capital with opportunities in the region. He holds a Bachelor of Engineering from Visvesvaraya Technological University (VTU), an MBA from London Business School, and is pursuing a Master of Laws (LLM) at University College London (UCL). He writes on capital structure, private markets and cross-border investment, with an emphasis on making complex financing decisions intelligible to both capital providers and capital seekers.

APPENDIX A. BASE-CASE ASSUMPTIONS

Table 4. Full Parameter Set for the Underwriting Framework

Parameter	Base case	Low	High
Funding cost	5.0%	4.0%	6.0%
PD Tier 1	2.0%	1.0%	3.5%
PD Tier 3	6.0%	3.5%	9.0%
PD Tier 5	12.0%	8.0%	20.0%
LGD Tier 1	15%	10%	25%
LGD Tier 3	35%	25%	50%
LGD Tier 5	60%	45%	75%
Valuation haircut	12%	8%	18%
Forced-sale discount	15%	10%	25%
Enforcement cost	7%	4%	12%
Enforcement time	12 months	6 months	30 months
Capital multiple (vs EL)	2.5x	2.0x	3.5x
Target return on capital	15%	12%	20%
Operating cost	0.8%	0.5%	1.2%

Indicative, illustrative parameter set; low and high columns define the ranges used in the robustness analysis. Not forecasts.

APPENDIX B. GLOSSARY

The following terms are used in this paper. Definitions are provided for the non-specialist reader and are deliberately plain rather than technical.

Capital charge The cost of holding capital against the possibility that losses exceed the expected; part of the required coupon.

Collateral cushion The excess of collateral value over the loan, which absorbs haircuts and discounts before the lender takes a loss.

Covenant A condition in a loan agreement that the borrower must meet, breach of which can trigger default or enforcement.

Credit committee The body within a lender that approves or declines transactions and sets their terms.

Economic capital The capital a lender holds against unexpected loss, sized to survive a stressed outcome.

Enforcement The legal process by which a lender realises its security after a default.

Expected loss The average loss anticipated over time, equal to the probability of default multiplied by the loss given default.

Five Cs The traditional underwriting framework: character, capacity, capital, collateral and conditions.

Forced-sale discount The reduction in realised price when collateral is sold quickly in a distressed rather than orderly sale.

Funding cost The interest rate a lender pays on the capital it lends out; the floor beneath the required coupon.

Haircut A reduction applied to an appraised or market value to reflect uncertainty or conservatism.

Loan-to-value (LTV) The ratio of a loan to the value of the collateral securing it.

Loss given default (LGD) The fraction of an exposure a lender loses when a borrower defaults, after recoveries.

Net recovery The value a lender actually realises from collateral after haircuts, discounts, costs and delay.

Probability of default (PD) The likelihood, over a given period, that a borrower will default.

Recovery rate The fraction of an exposure recovered after default; one minus the loss given default.

Required coupon The coupon a rational lender must charge to cover funding, expected loss, capital cost and a target margin.

Return on capital The profit a lender earns relative to the capital it holds against a transaction's risk.

Risk-based pricing Setting the price of credit to reflect the specific risk of the transaction.

Risk tier A category, here one to five, grouping transactions of similar assessed risk and required coupon.

Scorecard A structured tool that scores and weights underwriting factors to produce a risk assessment.

Security package The set of charges, guarantees and controls a lender takes to protect and recover its loan.

Seniority The rank of a claim in the capital structure; senior claims are repaid first.

Sponsor The owner or developer behind a transaction, whose quality strongly influences default risk.

Step-in rights A lender's contractual right to take over a project or asset on default to protect its position.

Through-the-cycle An average over a full economic cycle, including good and bad years.

Unexpected loss The amount by which losses could exceed the expected in a stressed year; capitalised, not priced.

Valuation haircut A reduction to an appraised value reflecting appraisal uncertainty.