

Loss Experience in GCC Private Credit: Is the Yield Real?

An evidence-led framework for testing whether the Gulf's double-digit coupons survive realised defaults and losses, and what loss experience the yield can actually withstand.

Chennakeshav Adya

Independent Researcher

June 2026

ABSTRACT

The companion papers established why Gulf real-asset coupons are high, how to underwrite and price the credit behind them, and how to engineer the recovery that limits loss. This paper asks the question on which the whole investment case ultimately turns: once realised defaults and losses are taken into account, is the yield real? A high gross coupon is worth nothing if losses consume it, and the central anxiety of any credit investor approaching a new market is that the yield is an illusion that a full credit cycle will dissolve. The paper builds a transparent, evidence-led framework for testing the yield against loss. It models default and loss experience across underwriting tiers and economic scenarios, decomposes the gross coupon into the net yield that survives losses, costs and fees, and examines how much loss the coupon can absorb before the net yield turns unattractive or negative. The central finding is that, for well-underwritten, well-secured Gulf credit, the gross coupon provides a substantial buffer against loss, so that the net yield remains attractive across a wide range of loss outcomes and is eroded to unattractive levels only under severe, sustained stress concentrated in the weakest transactions. The yield is real, but it is a net figure that must be defended through underwriting, structuring and diversification rather than assumed from the headline. The analysis is illustrated with modelled loss curves, default and loss tables, a net-yield bridge and scenario analysis, and is intended for credit funds and institutional allocators testing the durability of Gulf private-credit returns. All figures are modelled and illustrative, not forecasts, and are calibrated to general private-credit loss experience and to market structure rather than to proprietary data.

JEL Classification: G21, G24, G32, G17, C53

Keywords: loss given default, default rate, credit losses, private credit, net yield, loss curves, GCC, United Arab Emirates, stress testing, credit cycle

1. INTRODUCTION

There is a question that sits beneath every conversation about Gulf private credit, often unspoken because it sounds naive, yet more important than any other: is the yield real? A coupon of fifteen or eighteen per cent is easy to quote and easy to admire, but the return an investor keeps is the coupon less the losses, costs and fees that the full life of the book imposes. If losses are modest, the high coupon translates into a high net yield and the investment case is made; if losses are large, the coupon is merely the compensation for a loss that arrives, and the net yield collapses. The headline coupon, in other words, is a gross figure that says nothing definitive about the return until it is tested against loss.

This anxiety is rational and should not be dismissed. New, high-yielding markets have a long history of disappointing investors who confused the coupon with the return, deploying capital on the strength of a headline yield only to watch a credit cycle reveal that the yield was compensation for risk all along. A disciplined investor approaching Gulf private credit is right to ask whether the Gulf will prove the same, and to demand evidence rather than assurance. The purpose of this paper is to supply a framework for generating and interrogating that evidence, by modelling loss experience explicitly and testing the yield against it.

The framework rests on a distinction that the earlier papers developed and this one applies: the gross coupon is mostly compensation for illiquidity, complexity and the supply of scarce capital, with a smaller component compensating for expected credit loss. If that decomposition is correct, then the credit-loss component, the part that realised losses will consume, is a minority of the coupon, and the remainder is a buffer that losses must exhaust before the net yield is impaired. The question whether the yield is real therefore becomes a question about the size of that buffer relative to the losses a full cycle can impose: if the buffer is large relative to plausible losses, the yield is real and durable; if it is thin, the yield is fragile.

The paper approaches this through modelling rather than through a proprietary loss dataset, for the honest reason that Gulf private credit is too young to have generated a long, public record of realised losses through a full cycle. This is a genuine limitation, and the paper neither hides it nor pretends to data it does not have. Instead it models loss experience transparently, calibrated to the general loss experience of comparable private-credit markets and to the structure of the Gulf market, and it stresses the model hard, so that the conclusions rest not on optimistic assumptions but on how the yield behaves across a wide range of loss outcomes including severe ones. The aim is not to predict Gulf losses but to establish how much loss the yield can withstand.

The remainder of the paper is structured as follows. Section 2 reviews the literature on credit losses, loss distributions and the credit cycle, and develops the propositions the analysis examines. Section 3 sets out the data limitations and the stylised loss-modelling methodology. Section 4 presents the results: modelled default and loss experience across tiers and vintages, the net-yield bridge, the comparison with developed-market loss experience, and the size of the yield buffer. Section 5 stresses the framework through scenario analysis spanning benign conditions to crisis. Section 6 discusses implementation. Section 7 concludes.

2. LITERATURE REVIEW AND PROPOSITIONS

The framework draws on four bodies of work: the measurement of credit losses and their components; the modelling of loss distributions and portfolio credit risk; the literature on the credit cycle and the time-variation of losses; and the emerging evidence on private-credit loss experience. Each informs how loss is modelled and how the yield is tested against it.

2.1 Measuring Credit Losses

The measurement of credit loss rests on the decomposition, familiar from the underwriting paper, of expected loss into a probability of default and a loss given default. The literature, from Altman (1968) on default prediction to the large body of work on realised recoveries, establishes that both components vary across borrowers, over time and with the economic cycle, and that realised losses can differ substantially from the expected losses assumed at origination. For the present paper the key lesson is that loss is not a single number but a distribution, and that testing whether

the yield is real requires examining the yield not only against the expected loss but against the range of losses the distribution can produce, including its adverse tail.

2.2 Loss Distributions and Portfolio Credit Risk

The portfolio-credit-risk literature, including the work of Vasicek (2002) and the frameworks underlying modern credit-risk management, models the distribution of losses on a portfolio of credits and shows that this distribution is asymmetric: most of the time losses are modest, but occasionally, when a common factor causes many credits to default together, losses are severe. The shape of the distribution is governed by the correlation between credits: higher correlation produces a fatter tail, a greater chance of severe simultaneous losses. For Gulf real-asset credit, where many loans share exposure to the property cycle, the correlation is material, and the loss distribution accordingly has a meaningful tail that the modelling must capture.

2.3 The Credit Cycle and Time-Varying Loss

A third strand examines how losses vary over the credit cycle. The literature establishes that defaults and losses are strongly cyclical, low in expansions and high in downturns, and that the two components of loss reinforce each other in stress, as rising defaults coincide with falling recoveries. This procyclicality means that a book originated in a benign period and never tested by a downturn will show low losses that flatter its apparent quality, while the true test of the yield comes only when a downturn arrives. The paper reflects this by modelling loss across the cycle, from benign vintages through downturn vintages, rather than assuming the benign experience persists.

2.4 Evidence on Private-Credit Losses

The fourth strand is the growing body of evidence on realised losses in private credit and direct lending in developed markets, which, while not directly transferable to the Gulf, provides a reference point. This evidence suggests that well-underwritten senior direct lending has experienced relatively modest realised losses through recent cycles, lower than commonly feared, while more subordinated and aggressive strategies have experienced higher and more variable losses. The paper uses this developed-market experience as a calibration anchor and a comparison, asking how Gulf loss experience might compare given the differences in collateral, market structure and enforcement that the earlier papers examined.

2.5 Synthesis and Propositions

Taken together, these literatures imply that loss is a distribution rather than a number, that its tail is governed by correlation, that it varies strongly over the cycle, and that the relevant test of the yield is whether the gross coupon buffers the range of losses the distribution can produce. These implications motivate the following propositions.

Proposition 1. The gross coupon on well-underwritten Gulf real-asset credit substantially exceeds the expected loss, providing a buffer that realised losses must exhaust before the net yield is impaired.

Proposition 2. The net yield remains attractive across a wide range of loss outcomes, and is eroded to unattractive or negative levels only under severe, sustained stress concentrated in the weakest transactions.

Proposition 3. Losses are cyclical and correlated, so the benign experience of an untested book understates the loss a full cycle can impose; the yield must be tested against the downturn, not the expansion.

Proposition 4. The size of the yield buffer, and hence the durability of the net yield, depends on underwriting tier: it is largest for senior, well-secured transactions and smallest for subordinated, weakly secured ones.

Proposition 5. Diversification across uncorrelated dimensions widens the effective buffer at the portfolio level, but cannot eliminate the systematic property-cycle exposure that the buffer must ultimately absorb.

2.6 How the Propositions Connect

The propositions form a test of the yield. Proposition 1 establishes that a buffer exists; Proposition 2 asserts that it is large enough to keep the net yield attractive across most outcomes; Proposition 3 insists that the test be conducted against the downturn rather than the benign present; Proposition 4 locates the buffer's size in the underwriting tier; and Proposition 5 qualifies the role of diversification. If they hold, the yield is real, a net figure with a substantial buffer that survives all but severe stress, especially for well-underwritten senior credit, and the investor's task is to defend the buffer through underwriting, structuring and diversification rather than to assume it from the coupon. The sections that follow build and test this case.

2.7 Why Loss, Not Default, Determines the Net Yield

A point worth establishing before the modelling is that the quantity which erodes the net yield is loss, not default, and the distinction is more than semantic. A default is an event; a loss is its financial consequence after recovery. A book can experience a high default rate yet suffer modest losses if its recoveries are strong, because the collateral and enforcement examined in the companion paper return most of the exposure even when borrowers fail. Conversely, a book with a low default rate can suffer large losses if its recoveries are poor. It follows that the headline default rate, which alarms investors unfamiliar with a market, is an incomplete and often misleading indicator of the loss that actually reaches the net yield.

This reframes how the yield should be tested. The relevant test is not whether defaults occur, they will, in any credit book, but whether the losses net of recovery, accumulated over the cycle, exhaust the buffer the coupon provides. A market with visible defaults but strong recoveries can deliver a more durable net yield than a market with few defaults but weak recoveries, and the Gulf, with its strong real-asset collateral and the recovery engineering described earlier, is positioned to convert defaults into limited losses. The modelling that follows therefore tracks loss, not merely default, as the variable against which the yield is judged, and it is loss that the buffer must absorb.

2.8 The Evidentiary Standard for a Young Market

A question that pervades any analysis of a young market is what standard of evidence is fair to demand. To insist on a long, realised, through-the-cycle loss record before allocating would rule out every new market and asset class at precisely the moment their returns are highest, since by the time the record exists the opportunity has matured and the premium compressed. To allocate with no evidence at all, on the strength of the coupon alone, is the error the paper warns against. The reasonable standard lies between: to model loss transparently, calibrate it to the best

available comparators, stress it hard, and test whether the yield survives, while tracking realised experience and revising as it accumulates. This is the evidentiary standard the paper adopts and recommends.

Holding this standard distinguishes disciplined early allocation from speculation. The speculator allocates because the coupon is high and assumes the losses will be low; the disciplined early allocator allocates because the modelled and stress-tested net yield is attractive even under conservative loss assumptions, and monitors the realised experience to confirm or refute the thesis. The difference is not the absence of uncertainty, which is irreducible in a young market, but the rigour with which the uncertainty is confronted. An investor who applies this standard can participate in the Gulf opportunity while the premium is wide, without abandoning the evidentiary discipline that separates investment from gambling.

3. DATA AND METHODOLOGY

The study is a transparent loss-modelling exercise rather than an empirical study of realised Gulf losses, which do not yet exist in a long, public, through-the-cycle record. The methodology models loss experience explicitly and tests the yield against it, with every output traceable to a stated assumption and the assumptions calibrated to general private-credit loss experience and to the structure of the Gulf market. The paper is candid that this is a modelling exercise: it does not claim to measure Gulf losses but to establish how much loss the yield can withstand.

3.1 Data Limitations and the Modelling Response

The central data limitation is the youth of the market: Gulf private credit has not yet experienced a full credit cycle that would reveal its loss experience directly, and proprietary loss data, where they exist, are not public. The disciplined response is threefold: to model loss transparently rather than assert it; to calibrate the model to the loss experience of comparable private-credit markets, adjusted for the Gulf's differences in collateral and enforcement; and to stress the model across a wide range of scenarios so that the conclusions do not depend on a single, possibly optimistic, loss assumption. The limitation is real and is stated plainly; the modelling is designed to be robust to it.

3.2 The Loss-Modelling Framework

Loss is modelled at two levels. At the transaction level, expected loss is the probability of default multiplied by the loss given default, both varying by underwriting tier as in the underwriting paper. At the portfolio level, losses are aggregated across transactions allowing for correlation through the property cycle, so that the model produces not only an expected portfolio loss but a distribution of possible losses with an adverse tail. Loss is also modelled over time, through cumulative loss curves that show how losses accrue over the life of a vintage, and across the cycle, through scenarios from benign to crisis. The framework thus captures the level, distribution, timing and cyclicity of loss.

3.3 The Net-Yield Bridge

The test of the yield is the net-yield bridge, which moves from the gross coupon to the net yield by deducting realised losses, servicing costs, currency and hedging costs and fund fees, and adding back the recovery that limits the loss. The net yield is what the investor keeps, and the bridge makes explicit how much of the gross coupon survives a given loss experience. By

running the bridge across scenarios, the framework shows how the net yield behaves as losses rise, and so reveals how much loss the yield can absorb before it ceases to be attractive.

3.4 Defining the Yield Buffer

The yield buffer is the central concept of the paper: the excess of the gross coupon over the expected loss, which realised losses must exhaust before the net yield falls below an acceptable level. A large buffer means the net yield is robust to loss; a thin buffer means it is fragile. The buffer is measured both at the transaction level, where it depends on the tier, and at the portfolio level, where diversification across uncorrelated dimensions can widen it. The size of the buffer relative to plausible losses is the quantitative answer to the question whether the yield is real.

3.5 Base-Case Assumptions

Table 1 sets out the central assumptions, calibrated conservatively. Default and loss rates are set toward the higher end of plausible ranges, correlation is set high enough to produce a meaningful tail, and fees are set at standard levels, so that the net yields the framework produces are if anything understated. The full parameter set and the ranges used in the scenario analysis appear in Appendix A.

Table 1. Base-Case Assumptions for the Loss Framework

Parameter	Base-case assumption
Gross coupon (blended book)	15.0%
Annual default rate by tier	2%-12%
Loss given default by tier	20%-60%
Expected portfolio loss (base)	1.8% per year
Loss correlation (property cycle)	Moderate-to-high
Servicing & operating cost	0.8%
FX / hedging cost	0.3%
Fund fee load (return-equivalent)	2.4%
Hold period	3-5 years
Recovery lag	12-24 months

Indicative, illustrative assumptions used to drive the loss framework; calibrated to general private-credit experience, not Gulf-specific data, and not forecasts.

4. RESULTS AND DISCUSSION


This section presents the modelled loss experience and the test of the yield: cumulative loss curves, default and loss by tier, the net-yield bridge, the comparison with developed markets, and the size of the yield buffer.

4.1 Cumulative Loss Curves by Vintage

Figure 1 shows modelled cumulative loss curves for three vintages: a benign vintage originated and run in good conditions, an average vintage, and a downturn vintage that meets a property downturn during its life. The curves diverge sharply. The benign vintage accumulates modest losses; the average vintage more; the downturn vintage substantially more, with losses

accelerating in the years the downturn bites. The figure makes Proposition 3 concrete: the loss a book experiences depends heavily on whether it meets a downturn, and a book judged only by its benign early years will understate the loss a full cycle can impose.

Figure 1. Cumulative Loss Curves by Vintage




Loss experience depends heavily on whether a vintage meets a downturn; benign early years understate full-cycle loss.

The vintage analysis carries an investment lesson about timing and patience. A book originated near the top of a cycle and tested by the subsequent downturn will show the downturn curve; one originated after a correction, into a recovering market, will show the benign curve. The investor cannot control which vintage it is buying into with certainty, but it can recognise where in the cycle it is originating, lend more cautiously near the top and more confidently after a correction, and judge a manager's loss experience in the light of the vintages it has actually lived through rather than crediting benign early returns as proof of durable skill.

4.2 Default and Loss by Tier

Figure 2 shows modelled annual default and loss rates across the five underwriting tiers. Both rise with the tier, and the loss rate rises faster than the default rate, because the higher tiers combine higher default probability with higher loss given default, the two reinforcing each other. The figure quantifies Proposition 4: the loss a book experiences depends on its tier composition, and a book concentrated in the high tiers will experience materially higher losses than one tilted toward the low tiers, even at the same average coupon. The tier composition is therefore as important to loss experience as the average coupon is to gross income.

Figure 2. Default and Loss Rate by Underwriting Tier



Both rise with the tier, and the loss rate rises faster, because high tiers combine higher default and higher loss given default.

Table 2 sets out the tier-by-tier default and loss detail and the resulting expected loss, the input to the net-yield bridge. The table shows that even the higher tiers have expected losses well below their gross coupons, so that a buffer exists at every tier; but the buffer narrows as the tier rises, because the loss grows faster than the incremental coupon. This is the quantitative content of Proposition 4: the yield is real at every tier, but its buffer, and hence its durability under stress, is largest at the senior end and thinnest at the subordinated end.

Table 2. Modelled Default, Loss and Buffer by Tier (Illustrative)


Tier	Default rate	Loss given default	Expected loss	Gross coupon	Buffer
Tier 1	2.0%	20%	0.4%	9.5%	9.1%
Tier 2	3.5%	26%	0.9%	11.5%	10.6%
Tier 3	5.5%	33%	1.8%	14.0%	12.2%
Tier 4	8.0%	40%	3.2%	17.0%	13.8%
Tier 5	12.0%	46%	5.5%	21.0%	15.5%

Illustrative; buffer is gross coupon less expected loss, before costs and fees. Buffer in absolute terms is large at all tiers but proportionally thinner at high tiers.

4.3 The Net-Yield Bridge

Figure 3 presents the net-yield bridge for the blended book, moving from the gross coupon through the default loss, the recovery that offsets part of it, servicing and currency costs and fund fees, to the net yield. The bridge shows that, under base-case losses, the net yield remains firmly in double digits, well below the gross coupon but well above developed-market alternatives. The default loss, before recovery, is the second-largest deduction after fees, but the recovery, the subject of the security and enforcement paper, offsets a substantial part of it, illustrating how the recovery engineering of that paper feeds directly into the net yield of this one.

Figure 3. From Gross Coupon to Net Yield After Losses




Under base-case losses the net yield remains double-digit; recovery offsets much of the default loss, fees are the largest deduction.

The bridge is the operational answer to the paper's title question under base-case conditions: the yield is real, a net figure several points below the headline but substantial, and the deductions that separate gross from net are dominated not by losses but by fees, with losses, after recovery, a smaller drag. This reinforces a theme of the series: for well-underwritten, well-secured Gulf credit, the credit-loss component of the coupon is a minority of the total, so realised losses, while real, do not consume the yield under normal conditions.

4.4 Comparison with Developed-Market Loss Experience

Figure 4 compares modelled Gulf loss rates with the loss experience of developed-market private credit at matched risk tiers. The modelled Gulf loss rates are somewhat higher, reflecting the younger market and the property-cycle concentration, but not dramatically so, because the strong collateral and recovery engineering examined in the earlier papers limit loss given default. The comparison matters because the Gulf coupon is substantially higher than the developed-market coupon at the same tier, while the modelled loss is only modestly higher; the gap between the larger coupon premium and the smaller loss premium is the extra buffer that makes the Gulf net yield attractive.

Figure 4. Modelled GCC versus Developed-Market Loss Rates




Modelled Gulf losses are modestly higher than developed-market losses; the coupon premium is far larger than the loss premium.

This is the crux of the comparative case. If the Gulf coupon were higher only because Gulf losses were correspondingly higher, the extra yield would be an illusion, fully consumed by the extra loss. The modelling suggests the opposite: the coupon premium far exceeds the loss premium, so that even after the higher losses, the Gulf net yield exceeds the developed-market net yield. The extra coupon is mostly the illiquidity and complexity premium documented in the first paper, which is not consumed by loss, rather than a credit premium that is.

4.5 The Yield Buffer

Figure 9 presents the yield buffer over time: the cumulative gross coupon a book earns against the cumulative loss it suffers, under base and stressed conditions. The gross coupon accumulates steadily and substantially; even under stress, cumulative losses remain well below cumulative coupon, so the book remains in positive net territory throughout. The figure is the visual statement of Proposition 1: the coupon provides a large buffer, and losses, even stressed losses, exhaust only part of it. The yield is real because the buffer is large relative to the losses a plausible cycle imposes.

Figure 5. The Yield Buffer: Coupon Accumulation versus Loss




Cumulative coupon far exceeds cumulative loss under base and stress; the buffer absorbs losses and leaves positive net yield.

4.6 Loss by Vintage and Year

Figure 6 presents modelled loss as a surface over vintage and year, showing how each vintage's losses evolve and how a downturn year elevates losses across vintages simultaneously. The surface illustrates the correlation that the portfolio model captures: a bad year is bad for many vintages at once, because the common property-cycle factor strikes them together. This simultaneity is what produces the adverse tail of the loss distribution, and it is why diversification across vintages, while helpful, does not eliminate the systematic risk that all vintages share.

Figure 6. Modelled Loss Rate by Vintage and Year



A downturn year elevates losses across vintages at once; the common factor produces the tail of the loss distribution.

4.7 Interpreting the Results

The base-case results support the propositions. The gross coupon exceeds the expected loss by a wide margin, providing a substantial buffer (Proposition 1); the net yield remains attractive under base-case losses (Proposition 2, tested further under stress in Section 5); losses are cyclical, so the benign vintage understates full-cycle loss (Proposition 3); the buffer is largest for senior tiers (Proposition 4); and the correlation across vintages limits the benefit of diversification (Proposition 5). The yield, on the base-case evidence, is real; the question Section 5 addresses is how it behaves under stress.

4.8 The Asymmetry of the Buffer

An important feature of the yield buffer is its asymmetry across the loss distribution. In benign and base-case conditions, losses consume only a small fraction of the buffer, and the net yield sits close to the gross coupon less fees; the buffer is barely touched. As conditions deteriorate, losses rise and consume more of the buffer, but because the buffer is large the net yield falls gradually rather than collapsing. Only in the extreme tail, where losses are several multiples of the base case, does the buffer approach exhaustion and the net yield fall toward zero. This asymmetry, a yield that is stable across most of the distribution and falls steeply only in the extreme tail, is precisely the profile a credit investor wants: predictable returns in normal times and a defined, bounded downside in the worst.

The asymmetry also explains why the average loss is a poor summary of the risk. An investor who looks only at the expected loss sees a small number and may conclude the risk is trivial; an investor who looks only at the crisis scenario sees a large number and may conclude the risk is unacceptable. The truth is in the shape of the whole distribution: a yield that is robust across the bulk of outcomes and impaired only in the tail. Sizing the allocation requires understanding this shape, allocating to the point where even the tail outcome remains tolerable, rather than reacting to either the comforting average or the alarming extreme in isolation.

4.9 Loss Timing and the Recycling of Capital

The timing of loss within the life of a book matters as much as its level, because of how capital recycles. In a short-tenor book, capital returns and is redeployed frequently, so a loss on one transaction is set against the coupons earned on many redeployments; the buffer is effectively replenished as capital cycles. In a long-tenor book, capital is locked for longer, so losses are set against fewer coupon cycles and the buffer recovers more slowly. The cumulative loss curves of Figure 1 capture the level of loss over time, but the investor should also consider how quickly the coupon stream rebuilds the buffer after a loss, which depends on the tenor and the redeployment opportunity.

This interacts with the cycle. A book that suffers losses early, then redeloys into a recovering market, can rebuild its buffer and recover its net yield; a book that suffers losses late, near the end of its life, has fewer coupon cycles left to absorb them. The practical implication is that loss timing relative to the cycle and the tenor shapes the realised net yield beyond what the average loss rate conveys, and that a manager able to redeploy returning capital into a recovering market can repair a book that a static analysis would write down. The buffer is not a fixed reservoir but one continually replenished by the coupon stream, and its resilience depends on how fast it refills.

4.10 Loss and the Choice of Strategy

The loss analysis bears directly on the choice between credit strategies, because each strategy occupies a different position on the trade-off between gross coupon and loss volatility. A senior

strategy earns a lower gross coupon but enjoys a large buffer and stable losses, so its net yield is modest in benign conditions and remarkably resilient in stress. A special-situations strategy earns a much higher gross coupon but suffers higher and more volatile losses, so its net yield is spectacular in benign conditions and can turn negative in stress. Neither is superior in the abstract; the choice depends on the investor's tolerance for loss volatility and its confidence in its underwriting, and a thoughtful programme often blends the two to balance the steady net yield of the senior strategy against the higher expected return of the subordinated one.

What the loss framework adds to this familiar trade-off is precision about where each strategy sits and how its net yield behaves across the distribution. By modelling the buffer and the net yield under stress for each strategy, the investor can see not only the expected return of each but its resilience, and can construct a blend whose combined net yield meets its target under base conditions while remaining acceptable under stress. The strategy choice, in this light, is not a bet on which approach is best but a deliberate construction of a loss profile, and the framework gives the investor the tools to construct it consciously rather than to inherit whichever profile a single strategy happens to deliver.

4.11 Fees, Losses and the Order of Deductions

A subtle but practically important point concerns the relative size and order of the deductions that separate gross coupon from net yield. The net-yield bridge showed that, under base-case conditions, fund fees are a larger deduction than losses net of recovery, which has a counterintuitive implication: for a well-underwritten book in normal conditions, the investor loses more of the gross coupon to the manager than to defaulting borrowers. This does not diminish the importance of loss, which dominates under stress, but it reframes the fee discussion. An investor focused entirely on credit loss while accepting a high fee load may be guarding against the smaller deduction while conceding the larger one, and the access structure, direct, fund or co-investment, therefore deserves as much attention as the credit risk.

Under stress the order reverses: losses swell and become the dominant deduction, while fees, often charged on invested capital, remain relatively fixed. The two deductions thus matter in different regimes, fees in benign conditions, losses in stress, and a complete analysis of the net yield must weigh both. The investor who minimises fees through an efficient access structure widens the buffer in benign conditions and preserves more of the coupon to absorb losses in stress, so fee discipline and loss resilience are complementary rather than separate concerns. The net yield is defended on both fronts at once: by underwriting and recovery against loss, and by access structure against fees.

5. ROBUSTNESS AND SCENARIO ANALYSIS


The base-case net yield is reassuring, but the test of whether the yield is real is how it behaves under stress. This section runs the net yield across scenarios from benign to crisis, stresses the loss drivers, and states the conditions under which the yield ceases to be real.

5.1 Loss Across Scenarios

Figure 5 shows the modelled portfolio loss rate across five scenarios, from benign through base, mild stress and severe stress to crisis. The loss rate rises sharply across the scenarios, roughly quadrupling from benign to severe stress and rising further in crisis, reflecting the cyclicity and correlation that the literature emphasises. The scenarios span a wide enough range to include

outcomes more severe than the Gulf has experienced in its short history, deliberately, so that the test of the yield is demanding rather than complacent.

Figure 7. Loss Rate Across Scenarios




Modelled loss rises sharply from benign to crisis; the scenarios deliberately span outcomes more severe than experienced to date.

5.2 Net Yield Across Scenarios

Figure 7 runs the net-yield bridge across the same scenarios, showing the net yield that survives each loss level. The result is the central finding of the paper. Under benign, base and mild-stress conditions, the net yield remains comfortably in double digits; under severe stress it falls but remains positive and respectable; only under the crisis scenario does it fall to a low single-digit level, and even then it remains positive for the blended book. The yield is real across a wide range of loss outcomes, and is impaired but not destroyed even under severe stress, which is the quantitative content of Proposition 2.

Figure 8. Net Yield After Loss Across Scenarios




Net yield remains double-digit through mild stress, positive and respectable under severe stress, low but positive in crisis.

The scenario analysis distinguishes the durable from the fragile parts of the yield. For the blended, well-underwritten book the yield survives even severe stress, because the buffer is large; but a book concentrated in the high tiers, with a thinner buffer, would see its net yield turn negative under the same stress, because there the coupon was in larger part genuine credit compensation. The durability of the yield, in other words, is not uniform: it is robust for senior, diversified books and fragile for subordinated, concentrated ones, exactly as Proposition 4 implies.

5.3 Sensitivity of the Net Yield

Figure 8 reports the sensitivity of the net yield to its loss drivers: the default rate, the recovery rate, loss clustering (correlation), the gross coupon, fund fees and servicing cost. The default rate, recovery rate and clustering dominate, because together they determine the realised loss and its tail; the coupon and fees matter but less. The tornado instructs the investor to focus on the drivers of loss, default, recovery and correlation, in assessing whether a given book's yield is durable, and to treat the coupon and fee levels as secondary, since they are known with more certainty.

Figure 9. Net-Yield Sensitivity to Loss Drivers



Default rate, recovery and loss clustering dominate the net-yield outcome; the coupon and fees matter but less.

5.4 When the Yield Is Not Real

Honesty requires stating the conditions under which the yield is not real, that is, under which realised losses consume it. The yield is not real if the book is concentrated in the high tiers, where the buffer is thin and severe stress turns the net yield negative. It is not real if underwriting is poor, so that realised default and loss exceed the modelled rates and the buffer is smaller than assumed. It is not real if recovery fails, so that loss given default is higher than the security and enforcement paper's engineering would deliver. And it is not real if losses cluster more severely than modelled, so that the tail is fatter than assumed. Each of these is a real risk, and each is addressable through the disciplines the companion papers describe: tier discipline, sound underwriting, recovery engineering and diversification.

5.5 Robustness of the Central Conclusion

Across the scenarios and sensitivities, the central conclusion, that the yield on well-underwritten, well-secured, diversified Gulf credit is real and survives all but severe stress, holds robustly, while the yield on poorly-underwritten, weakly-secured, concentrated books is fragile and can be consumed by loss. This is the appropriately nuanced answer to the paper's title question: the yield is real for those who earn it through discipline, and an illusion for those who chase it without. The modelling does not promise that Gulf losses will be low; it establishes that the yield can withstand losses materially higher than the base case before it ceases to be attractive, provided the book is well-constructed.

5.6 Putting the Stress Scenarios in Perspective

It is worth situating the stress scenarios against the kinds of events that have tested credit markets elsewhere, so that the severity of the test is clear. The severe-stress and crisis scenarios modelled here correspond, in their loss intensity, to major property and credit downturns, episodes in which collateral values fell substantially, defaults rose sharply and enforcement lengthened across the market at once. By calibrating the stress to events of this magnitude, the test deliberately exceeds anything the young Gulf private-credit market has experienced, ensuring that the conclusion, that the yield survives all but the most severe stress, is not an artefact of a mild test. The yield is being asked to withstand a serious downturn, not a minor wobble.

This perspective matters for interpreting the result. That the net yield remains positive even under a crisis-magnitude scenario, for a well-constructed book, is a strong statement, because it means the yield can absorb a downturn comparable to the worst that comparable markets have suffered. That it falls to a low single-digit level in that scenario is a reminder that the yield is not invulnerable: a sufficiently severe and sustained crisis, concentrated in a weakly underwritten book, can consume most of it. The honest reading is neither complacent nor alarmist: the yield is durable against serious stress for the disciplined investor, and fragile against it for the undisciplined one, which is the consistent finding of the entire series.

5.7 The Limits of the Model

Every modelling exercise has limits, and stating them plainly is part of using the model responsibly. The loss model assumes a particular structure for the loss distribution and its correlation, calibrated to comparable markets, and the true Gulf distribution may differ in ways that only realised experience will reveal; the tail in particular is the hardest part of any loss distribution to estimate, and the model's tail is an informed assumption rather than a measurement. The model also treats the scenarios as exogenous shocks of given severity, whereas in reality the depth and duration of a downturn, and the policy and market responses to it, are uncertain and interact with the losses in ways a stylised model cannot fully capture. These are real limits, and they counsel humility about any single number the model produces.

The appropriate response is to use the model for what it does well, illuminating how the net yield behaves across a range of loss outcomes, rather than for what it cannot do, predicting the precise loss the Gulf will suffer. The value of the exercise is not the point estimate of loss in any scenario but the shape it reveals: a net yield that is stable across most of the distribution and falls steeply only in the extreme tail, with the size of the buffer determining how far into the tail the yield remains acceptable. That qualitative shape is robust to the modelling assumptions in a way that any single loss number is not, and it is on the shape, not the point estimate, that the paper's conclusion rests.

6. IMPLEMENTATION CONSIDERATIONS

Turning the evidence into practice requires attention to loss tracking, tier discipline, diversification, manager assessment and the honest treatment of an untested record.

6.1 Tracking and Building Loss Data

Because the market lacks a long public loss record, the investor's own loss data are uniquely valuable, and the discipline of tracking them rigorously is a priority. Every default, recovery and loss should be recorded against the original underwriting assessment, so that, over time, the investor accumulates a proprietary dataset that reveals its actual loss experience and calibrates its future modelling. A lender or allocator that tracks losses rigorously will, after a few years, hold evidence that no newcomer can replicate, and that evidence is both a source of edge and the only true answer to the question whether the yield is real for its particular book.

6.2 Tier Discipline

Because the yield buffer narrows as the tier rises, tier discipline, controlling the proportion of the book in the high tiers, is the principal lever over the durability of the net yield. A book tilted toward the senior tiers sacrifices some gross coupon for a larger buffer and a net yield that survives stress; a book that maximises coupon by concentrating in the high tiers earns more in benign conditions and risks a negative net yield in stress. The investor should set and hold a tier allocation that delivers an acceptable net yield under stress, not merely an attractive one under base-case conditions.

6.3 Diversification and Its Limits

Diversification across sponsor, asset type, sub-region and vintage widens the effective buffer by reducing idiosyncratic loss, and the investor should pursue it deliberately. But the scenario and vintage analysis showed that diversification cannot eliminate the systematic property-cycle exposure that all the loans share, so the investor must size the overall allocation on the assumption that the systematic factor will eventually bind. Diversification reduces the loss in normal conditions and the idiosyncratic part of the tail; it does not protect against the correlated downturn, which only the buffer and conservative sizing can absorb.

6.4 Assessing a Manager's Loss Experience

For the allocator, assessing a manager's loss experience is delicate precisely because the market is young and most managers' records are benign and untested. The allocator should ask not only what losses the manager has experienced but through what vintages, distinguishing a record built entirely in benign conditions from one tested by a downturn. It should examine the manager's tier composition, its underwriting and recovery discipline, and its loss-tracking rigour, treating these as predictors of future loss more reliable than a short, benign track record. A manager with a disciplined process and a modest but tested record is preferable to one with a spectacular but untested one.

6.5 The Honest Treatment of an Untested Record

Both managers and allocators should treat an untested record honestly, neither dismissing it as worthless nor crediting it as proof of durable skill. A few years of low losses in benign conditions is genuine evidence of competent origination and underwriting, but it is not evidence that the yield survives a downturn, because the downturn has not occurred. The honest position is to

present the record for what it is, to model the loss a downturn would impose, and to test the yield against that modelled stress, exactly as this paper does. An investor who confuses a benign record with a stress-tested one is making the very error the paper warns against.

6.6 Pricing the Buffer into the Allocation Decision

The yield buffer should be priced into the allocation decision as a margin of safety. An allocator should size its Gulf credit allocation, and choose its tier composition, so that the net yield remains acceptable under a stress scenario it regards as plausible, treating the base-case net yield as the expected outcome and the stressed net yield as the floor it is willing to tolerate. This converts the abstract question whether the yield is real into a concrete allocation discipline: allocate to the point where the stressed net yield is still acceptable, and no further. The buffer, properly priced, becomes the governor of the allocation size.

6.7 Communicating Loss Risk to the Committee

Finally, the loss analysis should be communicated to the investment committee in the same decomposed, stress-tested form as the coupon and the recovery. A committee should see the expected loss, the buffer, the net yield under base and stressed conditions, and the conditions under which the yield would cease to be real. Presenting the analysis this way disciplines the committee against both the optimism of crediting the benign record and the pessimism of dismissing the market, and it grounds the allocation decision in evidence about how the yield behaves under loss rather than in a reaction to the headline coupon.

6.8 Vintage Diversification as a Practical Defence

One practical defence against the cyclicity of loss is vintage diversification: deploying capital steadily over several years rather than concentrating it in a single period, so that the book is not entirely composed of loans originated at one point in the cycle. A book built over multiple vintages spreads its exposure across different entry points, so that some loans are originated into favourable conditions and others into less favourable ones, smoothing the loss experience relative to a single-vintage book that bets everything on the conditions of one year. Vintage diversification is the temporal analogue of the cross-sectional diversification across sponsor and asset type, and it addresses the timing risk that the vintage analysis highlighted.

Like other forms of diversification, it has limits. Vintage diversification smooths the idiosyncratic timing of entry but cannot escape a prolonged downturn that affects several vintages at once, as the loss surface of Figure 6 showed. Its value is in avoiding the concentration of an entire book at a cyclical peak, not in immunising against the cycle altogether. For the investor building a Gulf credit allocation, the practical guidance is to deploy steadily rather than rush capital in at a single moment, accepting a slower build in exchange for a smoother and more resilient loss experience across the cycle.

6.9 What Would Change the Conclusion

Intellectual honesty requires stating what evidence would change the paper's conclusion that the yield is real for the disciplined investor. If realised losses, once the market experiences a full cycle, prove systematically higher than the conservative modelling assumes, because collateral is weaker, recoveries poorer or correlation higher than calibrated, the buffer would be thinner than estimated and the conclusion would weaken. If enforcement proves less effective in practice than the security paper anticipates, loss given default would be higher and the net yield lower. And if

the structural premium compresses faster than expected, as the first paper warned it eventually must, the gross coupon and hence the buffer would shrink. Each of these is a genuine possibility that the accumulation of realised data will test.

The appropriate response to this uncertainty is not to avoid the market but to monitor the evidence and adjust. An investor should track its realised losses against the modelled expectations, watch for the compression of the coupon premium, and stand ready to reduce its allocation or raise its tier discipline if the evidence turns against the thesis. The conclusion that the yield is real is a conditional one, conditional on underwriting, recovery and the persistence of the buffer, and the disciplined investor holds it as a hypothesis to be tested by its own experience rather than as a faith to be maintained regardless of the data. That stance, evidence-led and ready to revise, is itself the central discipline the paper recommends.

6.10 Reporting Net Yield Honestly

A practical discipline that follows from the analysis is the honest reporting of net yield to investors. A manager that reports the gross coupon, or the net yield in benign conditions, without disclosing the loss experience and the behaviour of the yield under stress, presents an incomplete and potentially misleading picture, however accurate the figures themselves. The honest report presents the gross coupon, the realised and expected losses, the costs and fees, the resulting net yield, and crucially the behaviour of that net yield under stress, so that the investor understands not only what the book has earned in benign conditions but what it might earn through a downturn. This is the reporting analogue of the decomposition discipline that runs through the series.

Honest net-yield reporting protects both the investor and the manager. The investor is protected from the disappointment of discovering, in a downturn, that the benign net yield was not durable; the manager is protected from the reputational damage of having presented an incomplete picture, and from the pressure to chase yield in ways that thin the buffer. A manager confident in its underwriting and its buffer has no reason to hide the loss analysis, and its willingness to present the net yield honestly, stress and all, is itself a signal of the discipline that makes the yield real. Investors should prefer the manager who shows them the stressed net yield to the one who shows them only the benign headline.

6.11 Building the Allocation Around the Buffer

The culmination of the analysis for an allocator is to build the Gulf credit allocation around the buffer rather than the coupon. This means choosing a tier composition and a diversification that deliver an acceptable net yield under a chosen stress scenario, and then sizing the allocation so that even the stressed loss leaves the overall portfolio within its risk tolerance. The coupon sets the expected return, but the buffer and its behaviour under stress set the risk, and a disciplined allocation is sized to the risk, not the return. An allocator that sizes to the benign net yield will be over-allocated when stress arrives; one that sizes to the stressed net yield holds an allocation it can live with through the cycle.

Building the allocation around the buffer also clarifies how Gulf credit fits within a broader portfolio. Its attraction is a high net yield with a defined, bounded downside for a well-constructed book, which makes it a return-enhancing sleeve rather than a core holding, sized to contribute meaningfully to portfolio return without exposing the portfolio to an intolerable loss in the tail. The companion paper on portfolio construction develops this sizing in detail; the

contribution of the present paper is to supply the loss and net-yield evidence on which that sizing rests. The allocation is built, in the end, not on the promise of the coupon but on the demonstrated resilience of the net yield, which is the proper foundation for committing capital to a young but compelling market.

6.12 A Checklist for Testing the Yield

The analysis reduces, for the practitioner, to a checklist for testing whether a given Gulf credit book's yield is real. Establish the gross coupon and decompose it, confirming that the credit-loss component is a minority of the total. Estimate the expected loss from independent default and recovery assumptions, and compute the buffer. Examine the tier composition, since the buffer narrows with the tier. Model the net yield under base and stressed conditions, and confirm it remains acceptable under the stress the investor regards as plausible. Assess the diversification across sponsor, asset, sub-region and vintage. And review the manager's loss-tracking and the vintages its record actually spans. A book that passes this checklist has a yield that is real; one that fails it has a yield that is fragile, whatever the headline.

The checklist operationalises the paper's central message and connects it to the rest of the series. Each item draws on a companion paper, the decomposition on the coupon paper, the loss estimate on the underwriting paper, the recovery on the security paper, and the diversification and sizing on the portfolio paper, so that testing the yield is, in effect, applying the whole series to a single book. The investor who works through the checklist is not relying on assurance but generating evidence, and it is evidence, not the coupon, that answers the question the paper poses. The yield is real when the checklist says so, and the discipline of working through it is the price of knowing.

6.13 The Patient-Capital Advantage in Loss

A final implementation point connects loss resilience to the nature of the investor's capital. Much of the buffer's value is realised only by an investor able to hold through a downturn rather than being forced to sell or withdraw at the worst moment. An investor with stable, long-horizon capital can wait out a liquidity-driven dip in collateral values, allow recoveries to be realised in an orderly way, and let the coupon stream rebuild the buffer, capturing the full net yield the analysis describes. An investor with unstable capital, by contrast, may be forced to crystallise losses at the trough, converting a temporary impairment into a permanent one and forfeiting the recovery that patience would have delivered. The same book can thus deliver a durable net yield to patient capital and a poor one to impatient capital.

This reinforces a theme of the whole series: the Gulf opportunity rewards patient, capable capital disproportionately. The illiquidity premium that makes the coupon high is compensation for bearing the lock-up, and the loss resilience that makes the yield real depends on the ability to hold through stress. An investor whose funding matches the illiquidity of the asset captures both the premium and the resilience; one whose funding does not is exposed to forced selling that the framework cannot protect against. Matching the stability of the capital to the illiquidity of the asset is therefore not a detail but a precondition for the yield being real in practice as well as in the model.

6.14 Loss Experience and Manager Incentives

A final consideration links loss experience to the incentives under which a credit manager operates, because the two are more connected than is often recognised. A manager paid

principally on deployed volume or on gross income has an incentive to grow the book and to favour higher-coupon, higher-tier transactions, both of which raise loss experience and thin the buffer. A manager paid on net, risk-adjusted returns realised over time has the opposite incentive, to underwrite carefully, structure for recovery and avoid the concentrations that swell losses in stress. The loss experience an investor ultimately suffers is therefore shaped not only by the manager's skill but by the incentives that direct it, and an investor assessing a manager should examine the incentive structure as carefully as the track record.

This connection matters most precisely because the market is young and the record benign. In benign conditions, a volume-driven manager and a risk-driven manager may show similar returns, because losses have not yet tested the difference; it is only in the downturn that the volume-driven book reveals the higher losses its incentives encouraged. An investor who waits for the downturn to discover the difference has waited too long. The remedy is to infer the future loss experience from the incentive structure now: a manager whose compensation rewards durable, net, risk-adjusted performance is likely to deliver a loss experience consistent with the durable net yield this paper describes, while one whose compensation rewards deployment is likely to deliver the higher losses that thin the buffer. Incentives are a leading indicator of loss, and reading them is part of testing whether the yield is real.

7. CONCLUDING COMMENTS

Before drawing the threads together, it is worth restating the question in the plainest terms, because it is the question every prospective investor in this market silently asks. Strip away the modelling and the scenarios and the issue is simple: if I lend into the Gulf at these coupons, and a normal share of my borrowers fail, will I still make money? The whole apparatus of this paper exists to answer that plain question with evidence rather than assertion, and the evidence says yes for the disciplined lender and no for the careless one. Everything else is detail in service of that answer.

This paper has addressed the question on which the Gulf private-credit case ultimately turns: once realised defaults and losses are taken into account, is the yield real? The answer, developed through transparent loss modelling and demanding stress testing, is that for well-underwritten, well-secured, diversified Gulf credit the yield is real, a net figure with a substantial buffer that survives all but severe stress, while for poorly-underwritten, weakly-secured, concentrated books the yield is fragile and can be consumed by loss. The yield is real for those who earn it through discipline and an illusion for those who chase it without.

The practical implications follow. The gross coupon provides a large buffer over expected loss, but the buffer narrows as the tier rises, so tier discipline is the principal lever over the durability of the net yield. Losses are cyclical and correlated, so the benign experience of a young, untested book understates the loss a full cycle can impose, and the yield must be tested against the downturn rather than the expansion. Diversification widens the buffer but cannot eliminate the systematic property-cycle exposure, which only the buffer and conservative sizing can absorb. And the investor's own loss data, tracked rigorously, are the only true answer to whether the yield is real for its particular book.

The argument is properly and importantly qualified by the data limitation at its heart: Gulf private credit has not yet experienced a full credit cycle, so the loss experience modelled here is calibrated to comparable markets and to market structure rather than measured directly. The

paper does not hide this; it responds to it by modelling transparently, calibrating conservatively and stressing hard, so that the conclusion rests not on optimistic loss assumptions but on how much loss the yield can withstand. The honest claim is not that Gulf losses will be low but that the yield can absorb losses materially higher than the base case before it ceases to be attractive, for a well-constructed book.

The study points to clear extensions. As the market ages and experiences its first full cycle, the modelled loss curves can be replaced by realised ones, and the framework recalibrated to actual Gulf experience. The correlation structure across exposures could be estimated rather than assumed as it becomes observable. And the loss experience of different strategies, senior, mezzanine, special situations, could be tracked separately to refine the tier-specific buffers. Each extension would replace modelled evidence with realised evidence, strengthening the answer to the paper's question.

The paper completes a logical arc across the credit series. The first paper explained why the coupon is high and that it is mostly an illiquidity and complexity premium; the second showed how to underwrite and price the credit risk; the third showed how to engineer the recovery that limits loss; and this fourth paper has tested whether, after the losses that recovery does not prevent, the yield survives. The answer ties the arc together: the yield survives because the premium is mostly non-credit compensation that loss does not consume, because disciplined underwriting keeps the credit-loss component small, and because recovery engineering limits the loss that does occur. The durability of the net yield is the cumulative result of the disciplines the earlier papers described, which is why it accrues to the disciplined investor and eludes the undisciplined one.

For the allocator, the final message is one of conditional confidence. The Gulf yield is real and durable for a well-constructed book, but the construction is not optional: it requires tier discipline, sound underwriting, recovery engineering, diversification and honest, evidence-led monitoring. An allocator willing to do this work, or to back a manager who does, can capture a net yield that survives a serious downturn while the structural premium remains wide. An allocator unwilling to do it should not be reassured by the coupon, because for the undisciplined the yield is exactly the illusion they fear. The work is the price of the yield, and the yield, for those who pay that price, is real.

In sum, the paper's contribution is to subject the Gulf yield to the test it must pass to be worth pursuing: survival of realised loss through a cycle. On the modelled evidence, the yield passes the test for the disciplined investor and fails it for the undisciplined one, which is precisely the result the rest of the series would predict. The yield is real, but it is a net figure to be defended through underwriting, structuring and diversification, not a headline to be assumed.

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 general private-credit loss experience and to market structure rather than to Gulf-specific or any proprietary data, and are not forecasts of realised losses or returns. 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] Bharath, S. T., and Shumway, T. (2008). Forecasting Default with the Merton Distance to Default Model. *Review of Financial Studies*, 21(3), 1339-1369.
- [4] Duffie, D., and Singleton, K. J. (2003). *Credit Risk: Pricing, Measurement, and Management*. Princeton University Press.
- [5] Giesecke, K., Longstaff, F. A., Schaefer, S., and Strebulaev, I. (2011). Corporate Bond Default Risk: A 150-Year Perspective. *Journal of Financial Economics*, 102(2), 233-250.
- [6] Gupton, G. M., Finger, C. C., and Bhatia, M. (1997). *CreditMetrics Technical Document*. J.P. Morgan.
- [7] Moody's / rating-agency methodologies on through-the-cycle default and recovery (general reference).
- [8] Schuermann, T. (2004). What Do We Know About Loss Given Default? Wharton Financial Institutions Center Working Paper.
- [9] Vasicek, O. (2002). The Distribution of Loan Portfolio Value. *Risk*, 15(12), 160-162.
- [10] International Monetary Fund (2024). *Global Financial Stability Report: The Rise of Private Credit*. IMF, Washington, D.C.
- [11] Markowitz, H. (1952). Portfolio Selection. *Journal of Finance*, 7(1), 77-91.

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 3. Full Parameter Set for the Loss Framework

Parameter	Base case	Low	High
Gross coupon (blended)	15.0%	13.0%	17.0%
Default rate Tier 1	2.0%	1.0%	3.5%
Default rate Tier 3	5.5%	3.5%	9.0%
Default rate Tier 5	12.0%	8.0%	20.0%
LGD Tier 1	20%	12%	30%
LGD Tier 3	33%	22%	48%
LGD Tier 5	46%	35%	62%
Expected portfolio loss	1.8%	0.8%	3.5%
Stressed portfolio loss	5.5%	3.5%	8.5%

Loss correlation	Moderate	Low	High
Servicing cost	0.8%	0.5%	1.2%
FX / hedging cost	0.3%	0.2%	0.5%
Fund fee load	2.4%	1.5%	3.0%
Hold period	4 years	3 years	5 years

Indicative, illustrative parameter set; low and high columns define the ranges used in the scenario 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.

Buffer (yield buffer) The excess of the gross coupon over the expected loss, which realised losses must exhaust before the net yield is impaired.

Correlation The tendency of losses on different loans to move together, which fattens the tail of the loss distribution.

Credit cycle The cyclical pattern of low losses in expansions and high losses in downturns.

Cumulative loss curve A plot of total losses accrued over the life of a vintage, year by year.

Default rate The proportion of borrowers that default in a given period.

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

Gross coupon The interest rate a borrower pays, before any deductions to the lender.

Loss distribution The range of possible losses on a portfolio and their probabilities, typically asymmetric with an adverse tail.

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

Loss rate Losses as a proportion of the portfolio, equal to the default rate multiplied by the loss given default.

Net yield The return an investor keeps after deducting losses, costs and fees from the gross coupon.

Procyclicality The tendency of losses to rise in downturns and fall in expansions, reinforcing the cycle.

Recovery The value realised from collateral after default, which offsets part of the loss.

Scenario analysis The examination of outcomes under a range of defined economic conditions, from benign to crisis.

Stress testing Assessing how a portfolio performs under severe but plausible adverse conditions.

Systematic risk Risk common to many exposures, here the property cycle, that diversification cannot remove.

Tail (of a distribution) The region of rare but severe outcomes; a fat tail means severe outcomes are more likely.

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

Tier (underwriting tier) A category grouping transactions of similar assessed risk; higher tiers carry higher default and loss.

Vintage The set of loans originated in a given period, whose loss experience depends on the conditions it meets.

Idiosyncratic risk Risk specific to an individual loan, which diversification across many loans can reduce.

Realised loss The loss actually incurred, as opposed to the expected loss assumed at origination.

Servicing cost The expense of administering and monitoring a loan book.

Fund fee load The management fee and carry, expressed as a return-equivalent drag on investor returns.

Downturn vintage A set of loans whose life includes a property downturn, experiencing elevated losses.