PARA BELLUM ADVISORS

PRACTITIONER PAPER

The Credit Hedge Illusion

A Practitioner's Guide to What Actually Works

By Mike Duncan, Para Bellum Advisors

Version 3, December 2025

www.parabellumadvisors.com

Executive Summary

Many institutional credit hedges are expensive box-ticking exercises. They look disciplined in risk reports and satisfy governance committees, but in real credit stress they often fail to protect against the losses that matter most.

> *The problem is not that credit is unhedged; it is that most hedges do not hedge the risk that actually matters?*

Credit risk is not one thing. It includes default, recovery, spread, and liquidity risk, and most hedges mainly address spread volatility rather than economic loss.

That is why CDS often disappoints. It can reduce reported volatility, but it brings roll costs, basis risk, counterparty exposure, and poor monetisation when markets are under pressure.

The core problem is structural mismatch. Real credit losses are usually driven by defaults, restructurings, and impaired recoveries over multi-year horizons, while most hedges are built for mark-to-market movements and quarterly hedge tests.

For many portfolios, especially private credit, infrastructure, insurance general accounts, and long-duration investors, the better answer is not a more sophisticated hedge. It is stronger portfolio construction, better subordination, more diversification, larger loss reserves, and, in some cases, accepting the risk unhedged.

This paper argues that credit hedging should be judged on whether it transfers real economic loss, not whether it looks effective in governance or accounting terms.

1. The Credit Hedge Illusion

Credit hedging is usually treated as a solved problem. Buy protection. Size it to exposure. Monitor hedge effectiveness. Report reduced risk. For many institutions, that process is considered prudent by default. In practice, it rarely protects against the losses that actually matter.

Most credit losses are not driven by gradual spread volatility. They arise from structural events: defaults, restructurings, prolonged impairments, and recovery shortfalls that play out over years, not quarters. They are binary, slow-moving, and path dependent.

Yet most credit hedges are designed to manage mark-to-market spread movements, index behaviour, and accounting optics. They are optimised for quarterly reporting, regulatory comfort, and committee visibility – not for absorbing or transferring the economic loss that occurs when a credit genuinely fails.

> *This mismatch is not accidental. It is structural.*

Credit default swaps dominate institutional credit hedging not because they are effective at transferring loss risk, but because they are liquid in normal markets, familiar to consultants, and easy to explain to boards. They fit governance frameworks cleanly. They look like risk management. What they do not reliably do is pay when the portfolio suffers real credit damage – or pay in a way that can be monetised, timed, and applied to offset that damage.

As a result, many portfolios pay ongoing premiums for protection that covers the wrong risk, fails to monetise under stress, introduces liquidity, basis, and counterparty risk, and often worsens outcomes during actual credit events.

Credit exposure is not the same as credit outcome. A hedge can reduce spread volatility and still fail to protect against default loss. It can satisfy accounting tests and still create liquidity stress at the worst possible time.

2. What Credit Risk Actually Is

Before assessing whether a credit hedge works, you need clarity on what you are hedging. Credit risk is not one exposure. It is a bundle of distinct risks with different economic consequences and different hedge requirements.

2.1 Default Risk

The binary event – the borrower stops paying, enters restructuring, triggers cross-defaults, or becomes impaired. Default is where losses become final, and where most spread hedges stop being relevant.

2.2 Recovery Risk

Even when default probability is properly assessed, most credit models systematically underweight recovery uncertainty. Recovery assumptions that look stable in spreadsheets rarely hold in real restructurings, particularly for private credit with bespoke structures and thin legal precedent, emerging market corporates where processes are opaque, infrastructure assets where post-default viability is uncertain, and crossover credits where subordination gets tested.

The gap between assumed and realised recovery is often larger than spread moves. Most hedges do not address this risk – they implicitly assume recoveries are known and stable.

2.3 Downgrade and Spread Risk

Mark-to-market volatility from credit migration. This matters or does not matter depending entirely on how the portfolio is run. For insurance general accounts or hold-to-maturity portfolios, spread widening is often tolerable if default loss is contained. For leveraged vehicles, daily margining, or structures with liquidity constraints, spread volatility can become existential.

2.4 Liquidity Risk

The inability to exit, finance, or restructure positions when needed. Not a stress scenario – a predictable feature of credit markets once conditions tighten. Most credit hedges only address one of these risks – usually downgrade and spread risk.

Yet when credit portfolios actually lose serious money, it is typically because defaults clustered (correlation), recoveries were worse than modelled, positions could not be exited (liquidity), and hedges could not be monetised (liquidity again). The spread volatility that a CDS hedge captures is often immaterial compared to these structural failure modes.

Case Study 1: Asian Life Insurer – Hedging the Wrong Risk

Large Asian life insurer held $2.8bn investment-grade corporate credit portfolio, primarily in Asia ex-Japan, with long-term liability matching. Regulatory capital framework treated spread widening as capital-intensive, incentivising a spread hedge to reduce capital charges.

Implementation: rolling 5-year iTraxx Asia IG index hedge, 40% notional coverage. Spread risk dampened. Accounting tests passed. Risk committee satisfied.

December 2018: single A-rated Chinese property developer in portfolio defaulted and entered restructuring. Recovery ultimately 60 cents on the dollar. Economic loss: $108m over 24-month workout.

The hedge paid out nothing. Not in the index. Wrong credit. Wrong structure. Meanwhile, ongoing hedge costs over 36 months: $22m in premium decay plus $8m in basis losses when rolling hedges.

Net outcome: $108m real economic loss, $30m in hedge costs, zero hedge offset. The portfolio would have been better off unhedged, holding the $30m in premium as additional loss absorption.

3. Why CDS Dominates (and Why That's a Problem)

Credit default swaps are the default choice for institutional credit hedging. Not because they are particularly effective at transferring credit loss, but because they are familiar, liquid in normal markets, and sit comfortably within governance and accounting frameworks.

3.1 Liquidity Optics

CDS trades in size. Quotes are available. Market makers exist. This creates the appearance of hedgeability – if we can trade it, we can hedge it. But tradability in normal markets is not the same as liquidity when it matters.

March 2020 demonstrated this clearly. Investment-grade CDS indices continued to trade, but single-name CDS – particularly in financials and energy – saw bid-offer spreads widen dramatically. Monetising protection at anything close to fair value became difficult precisely when portfolios needed it most. The liquidity is real in conditions where the hedge is least needed. It evaporates in conditions where it is.

3.2 Governance and Accounting Comfort

CDS fits neatly into hedge accounting frameworks (IAS 39, IFRS 9, FAS 133). Hedge effectiveness can be demonstrated. Quarterly tests can be passed. Auditors understand the instrument. CFOs can explain it to boards.

This is seductive. It reframes the hard question – are we protected against credit loss? – into an easier one: does the hedge pass accounting and governance tests? Consultants, rating agencies, and risk committees are comfortable with CDS. They have models, benchmarks, and precedents. Proposing CDS hedges carries little career risk. Proposing not to hedge, or proposing non-standard structures, requires justification that few organisations reward.

> *The result is institutional inertia. CDS becomes the default not because it works, but because it is defensible.*

What CDS manages effectively	What CDS manages poorly
Quarterly P&L volatility	Realised default losses
Regulatory capital calculations	Recovery shortfalls
Board and consultant expectations	Liquidity crises under stress
Audit and governance requirements	Correlation events across credit

For portfolios that mark to market daily and face redemptions or margin calls, managing reporting risk may be rational. But for insurance general accounts, infrastructure equity, family office evergreen vehicles, and other long-duration holders, the trade-off is different. These portfolios often pay ongoing premiums to hedge accounting volatility while remaining exposed to the economic risks that ultimately drive outcomes.

4. The Hidden Costs of CDS Hedging

CDS premium is not the cost. It is the visible cost – the one that appears in P&L and gets debated in investment committee. The real costs are structural, often invisible, and frequently larger than the premium itself.

4.1 Roll Decay

CDS contracts have fixed maturities. Most credit exposures do not. That mismatch creates a persistent drag. Each time a maturing five-year CDS is rolled into a new contract, time decay is crystallised and protection is reset at prevailing market levels. If spreads have tightened, the hedge was expensive relative to what is now available. If spreads have widened, protection is rolled after the damage has already occurred.

Over time, this creates a systematic cost. A continuously rolled hedge behaves like short gamma on credit spreads – value is lost across cycles regardless of direction. As a rule of thumb, a 100bp five-year CDS rolled annually costs roughly 20-25bp per year in premium plus 5-8bp in roll friction. Over a decade, cumulative drag of 250-330bp is common for protection that may never meaningfully pay out.

4.2 Basis Risk

Credit exposure is specific. CDS contracts are standardised. The two rarely align. The hedge references a legal entity; the exposure sits in specific bonds or loans with defined seniority, maturity, and restructuring terms. Index hedges add another layer of mismatch by averaging credit quality and correlation that rarely reflects the underlying portfolio.

Restructuring definitions compound this risk. Modified restructuring clauses can result in situations where bonds are impaired or restructured while CDS protection fails to trigger, or triggers late and at recovery levels that bear little resemblance to realised outcomes. This basis risk is not an edge case. It is embedded in the structure of CDS hedging.

4.3 Jump-to-Default Illusion

CDS pricing implicitly assumes continuous deterioration. Credit reality is discrete. Defaults often occur suddenly – failed refinancings, fraud revelations, regulatory intervention – leaving little opportunity for hedges to adjust. When default occurs, CDS payouts are determined through post-event recovery auctions that occur weeks or months after default, rely on settlement mechanisms detached from actual bond recoveries, and often produce outcomes misaligned with portfolio-level losses.

> *The hedge "worked" in a technical sense. Economically and operationally, it did not.*

4.4 Liquidity Disappearance When Needed Most

CDS markets are liquid when credit risk is benign. They are not liquid when protection is needed. March 2020 illustrated this clearly. Investment-grade CDS indices continued to trade, while single-name CDS in financials and energy saw bid-offer spreads widen by an order of magnitude. This pattern repeats across cycles – 2008, 2015, 2018 – because credit stress produces the same conditions every time: dealer balance sheets contract, protection demand spikes, single-name risk becomes toxic, and liquidity evaporates.

> *The hedge exists on paper. In practice, it cannot be monetised without severe slippage or cannot be unwound when conditions normalise faster than expected. This is not a tail risk. It is the dominant failure mode of CDS hedging.*

Case Study 2: Pension Fund – The Roll-Decay Death Spiral

A$4.5bn pension fund held a diversified investment-grade corporate credit portfolio with a six-year average duration. Risk committee mandated 30% notional hedge coverage using five-year iTraxx Main, rolled semi-annually to maintain constant maturity.

Between 2015 and 2019: eight rolls executed, average roll cost ~12bp per roll, cumulative roll cost reached 96bp. Underlying credit portfolio delivered ~35bp of net positive performance. Hedge P&L over the period: -88bp. Combined outcome: -53bp versus remaining unhedged.

The hedge functioned as designed. Its cost exceeded the credit risk it was intended to protect against. The fund discontinued the programme in 2019 after accumulating ~$22m in negative carry.

The question that mattered was never asked upfront: whether the option value of protection justified the certainty of ongoing cost.

5. Index vs Single-Name CDS – Different Failures, Same Outcome

The choice between index and single-name CDS hedging feels consequential. Portfolios debate which approach better matches their exposures. In practice, both tend to fail – just in different ways.

5.1 Index CDS: Correlation Lies

Index CDS (iTraxx, CDX) offers clean execution, tight bid-offer spreads, and easy governance. The problem is correlation. Portfolios are rarely aligned with index composition – they are overweight certain sectors, underweight others, and often hold off-index credits entirely.

In normal markets, correlation might sit around 0.65-0.75. Adequate for optics, inadequate for economics – the portfolio is structurally underhedged with permanent basis leakage. In stress, correlation spikes as spreads widen together. The hedge briefly appears to work – until dispersion takes over. The weakest credits deteriorate far more than the index average, while stronger names recover faster. Post-crisis, correlation collapses. The index tightens as healthy constituents recover while problem credits remain wide or default. The hedge loses money while the portfolio is still impaired.

> *Index CDS provides average protection against average credit behaviour. Credit losses are not average. They are idiosyncratic, clustered, and binary.*

5.2 Single-Name CDS: Operational and Liquidity Toxicity

Single-name CDS promises precision. Specific exposures hedged with specific protection. Different problems emerge. Operational complexity rises sharply. Managing dozens of CDS positions against a changing bond portfolio requires continuous rebalancing, documentation, collateral management, and monitoring.

Liquidity is the larger issue. Outside the most actively traded names, single-name CDS liquidity deteriorates rapidly. In lower-rated credits, Asia ex-Japan, and cyclical sectors, bid-offer spreads widen materially, and two-way markets disappear under stress. Counterparty concentration compounds the risk. Wrong-way risk is common: protection is often bought from counterparties whose credit is correlated with the exposure being hedged. Credit risk is replaced with counterparty risk – often with worse correlation properties.

5.3 Why Both Fail Lean Teams

Effective CDS hedging requires resources most teams do not have. Index hedging demands ongoing correlation analysis, dynamic hedge adjustment, and tolerance for persistent basis risk. Single-name hedging requires operational infrastructure, legal and collateral frameworks, counterparty credit management, and realistic exit planning. Most teams have neither.

> *The relevant question is not "index or single-name?" – it is "do we have the infrastructure and discipline to make either work, and if not, should we be hedging credit at all?"*

6. Credit Hedging as a Liquidity Problem

Credit risk eventually manifests as a liquidity problem. Either you cannot sell the deteriorating credit, or you cannot monetise the hedge that is supposed to protect you, or both. Understanding this changes how credit hedging should be evaluated.

6.1 Credit Stress Does Not Mean Orderly Markets

Most credit hedging frameworks assume positions can be exited or monetised when needed. Credit reality is that when exits matter, markets are one-way. Bond markets seize. Dealers step back. Bid-offer spreads blow out. In theory, the CDS hedge should offset this. In practice: index CDS carries 30-40% basis to the specific deteriorating credits; single-name CDS suffers the same liquidity constraints as the underlying bonds; and attempting to exit both simultaneously hits the bid side of two stressed markets.

6.2 When Hedges Cannot Be Exited or Monetised

March 2020 illustrated this clearly. Investment-grade corporate bonds widened 150-200bp in two weeks. Single-name CDS on those same credits should have produced equivalent gains. Index CDS traded with 10-15bp slippage (versus 2-3bp in normal conditions). Single-name financials CDS saw 40-80bp slippage. High-yield and distressed single-name CDS often had no executable bid.

Hedges showed gains on paper. Monetising them required giving up 30-50% of those gains to slippage or waiting weeks for markets to stabilise – by which point spreads had already tightened and the opportunity had passed. Any hedge strategy that relies on monetising protection during stress assumes liquidity that is unlikely to exist.

6.3 Why Funding Matters More Than Spread Moves

What damages portfolios during credit stress is rarely the spread move itself. It is the funding and liquidity consequences that follow. For a leveraged, mark-to-market portfolio: the CDS hedge may show gains but sits with a different counterparty (no netting benefit), monetisation requires slippage, and settlement is T+2 while margin is due today. The portfolio is forced to liquidate bonds at distressed levels despite being "hedged."

For an insurance company under regulatory capital pressure: the CDS hedge helps regulatory calculations if structured correctly, but premium drag reduces return on capital, basis risk limits offset, and forced deleveraging still requires exiting both bonds and hedges at poor levels. For unlevered, patient capital: adding CDS introduces margining, collateral, and monetisation decisions that did not previously exist. A hold-to-maturity exposure becomes a derivatives position requiring active liquidity management. The hedge creates a liquidity problem where none existed.

Case Study 3: Multi-Asset Fund – When Hedges Create Liquidity Crises

$1.2bn multi-asset fund with 30% in IG corporate credit, hedged 40% notional using iTraxx Main. Governance metrics looked robust.

February 2020: spreads widened ~80bp. CDS hedge showed ~$15m of gains. Simultaneously, equities fell ~12%, redemptions reached ~8% of AUM, and immediate liquidity was required.

Monetising CDS gains appeared logical. Execution reality: iTraxx bid-offer widened to ~12bp, ~$1.8m lost to slippage, T+2 settlement versus T+0 redemptions. The fund sold equities and government bonds instead. Two weeks later, spreads tightened ~40bp, halving unrealised CDS gains.

Outcome: equities sold near the lows, CDS hedge eventually realised ~$7.5m, value destruction from forced asset sales exceeded hedge benefit. A simpler approach – holding more cash and accepting spread volatility – would likely have produced better outcomes.

The hedge created the illusion of liquidity that evaporated under stress, leading to worse decisions than if the team had accepted the illiquidity upfront.

7. Alternatives That Actually Transfer Risk

If CDS hedging is expensive box-ticking, what actually works? The answer depends entirely on what specific outcome you are trying to achieve. Several structural approaches transfer genuine economic risk rather than managing accounting optics.

7.1 Structural Subordination

The most explicit form of credit protection: accept you are senior to someone else who absorbs losses before you do. Instead of buying CDS on BBB corporate exposure, own senior secured debt while someone else – equity, mezzanine, junior unsecured – sits below you. Default happens; you are first in line for recoveries. The subordination is structural, not contractual.

Why it actually works: no counterparty risk unlike CDS; no liquidity risk – you are not trying to monetise a derivative; no roll costs – subordination is a permanent feature of the capital structure; no basis risk – you are protected by actual capital, not correlation assumptions. The cost is lower yield: senior secured typically yields 100-200bp less than unsecured exposure on the same credit. But that yield give-up is transparent, one-time, and permanent versus ongoing CDS premium that bleeds forever.

> *This is not hedging – it is portfolio construction. But portfolio construction that avoids risk is better than hedging that does not work.*

7.2 Asset-Level Hedging vs Portfolio-Level Optics

Most credit hedging is portfolio-level: buy index CDS, reduce aggregate risk metrics, satisfy governance. This optimises for reporting, not economic reality. Asset-level thinking asks a different question: for each specific credit exposure, what is the actual loss scenario and how would we prevent it?

For EM corporate exposure, the real risk is often not default but a currency crisis preventing dollar debt service – hedging with BRL/USD put spreads at 80bp per year addresses the actual failure mode and is materially cheaper than CDS. For project finance exposure, the real risks are construction delay or technical failure, not sponsor default – structural protections like construction bonding, technology guarantees, or subordinated equity from the sponsor are structurally superior. For leveraged loan exposure, the real risk is covenant-lite structure allowing deterioration without default – the protection is demanding maintenance covenants, requiring amortisation schedules, and walking away from deals that do not meet standards.

7.3 Where Insurance-Style Thinking Beats Trading Logic

CDS is a trading instrument designed for mark-to-market books, active positioning, and tactical hedging. Most institutional credit portfolios are insurance-style: buy, hold, collect income, manage binary default risk over multi-year horizons. Insurance-style protection looks different.

It delivers binary payouts on defined events rather than MTM protection, operates as a long-dated non-recourse arrangement with no collateral or margin calls, and prices diversification into the premium where law of large numbers applies. For long-dated, hold-to-maturity credit portfolios, the ideal protection looks more like trade credit insurance or political risk insurance than CDS: upfront or level premium with no mark-to-market

volatility, defined trigger events, par payout or recovery shortfall payment, no collateral requirements, and multi-year commitment with no roll risk.

> *The irony: insurance companies managing credit portfolios use CDS (trading logic) when they should use insurance (their own business model logic).*

7.4 The Truthful Assessment: When Is Imperfect Better Than Fragile?

Perfect hedges do not exist. The choice is between fragile precision – CDS that perfectly matches your exposure on paper but fails via basis risk, liquidity risk, counterparty risk, or roll costs – and robust imperfection: structural subordination, insurance-based protection, or simply accepting unhedged exposure with appropriate portfolio construction.

Robust imperfection often wins: no operational complexity, no ongoing costs if risk does not materialise, no liquidity risk when you need protection, no basis risk, no counterparty risk. The cost is imperfect correlation and potentially wider distributions of outcomes. But if your portfolio can absorb those wider distributions, robust imperfection is economically superior to fragile precision.

Framework: When structural alternatives beat CDS

If you answer YES to 3 or more of these, structural alternatives beat CDS:

- Hold period >3 years for majority of positions
- No forced MTM (insurance, family office, sovereign wealth)
- Ability to influence capital structure at origination
- Limited operational resources for derivative management
- Primary concern is default loss, not spread volatility
- Portfolio is diversified (>30 positions, <5% in largest)

If you answer NO to 3 or more, CDS might be appropriate:

- Daily MTM and redemptions (mutual fund, hedge fund)
- Leveraged portfolio with margin requirements
- Concentrated exposures you cannot diversify
- Dedicated derivatives team and infrastructure

8. Private Credit and Infrastructure: The Hardest Case

If CDS hedging is problematic for liquid IG corporates, it is nearly impossible for private credit and infrastructure debt. The mismatch between illiquid assets and liquid hedges creates structural problems that no amount of sophistication can solve.

8.1 Illiquid Assets, Liquid Hedges

Private credit has fundamental characteristics: bespoke structures negotiated bilaterally, limited or no secondary market, hold-to-maturity accounting, quarterly or annual valuations not daily marks, covenants and monitoring as primary risk control, recovery via workout not liquidation. Now try to hedge this with CDS: standardised contracts on public reference entities, daily mark-to-market with collateral requirements, correlation assumptions to proxy hedges since direct hedges do not exist, payout based on bond recovery auctions not actual loan recoveries.

Every dimension is mismatched. Valuation mismatch creates accounting volatility you were trying to avoid. Recovery mismatch means you might recover 70 cents, your hedge assumes 40 cents, and you have paid for protection against losses that do not match reality. Trigger mismatch means your loan might go non-performing but not technically default. Correlation mismatch means you are hedging private sponsor-backed LBO loans with a public high-yield index at correlations that range from 0.20 to 0.70 depending on whether stress is sector-wide or idiosyncratic.

8.2 Why Mismatch Is Structural

This is not a "find a better hedge" problem. It is a fundamental impossibility. Private credit works because it is illiquid and bespoke. The illiquidity premium is compensation for no exit option during stress, concentration risk from large positions, operational complexity of monitoring and workouts, and relationship-based restructurings when things go wrong. Trying to hedge this with liquid instruments means either paying away your illiquidity premium to gain liquidity you explicitly chose not to have or accepting massive basis risk that renders the hedge performative.

> *There is no magic third option.*

European Private Debt Fund: The Unhedgeable Illiquidity Premium

European private debt fund, €2bn AUM, 35 sponsor-backed LBO loans. Investors demanded "credit protection" before committing additional capital.

Fund analysis: no liquid CDS on any portfolio company (all private). Proxy hedging via HY index would cost 180bp per year. Illiquidity premium earned on portfolio: 220bp over liquid HY. Correlation between portfolio and HY index: 0.35-0.55. Expected hedge effectiveness: ~40%.

Economics: pay 180bp for 40% effective protection = 450bp cost per unit of effective protection, versus 220bp of excess return. Hedging would guarantee negative alpha.

Fund response: declined to hedge. Instead, increased diversification to 40 loans, enhanced monitoring and covenant packages, built larger loss reserves, and accepted concentration limits. Cost them some investors who wanted the comfort of "hedged" even if the hedge did not work.

8.3 Infrastructure Debt: Even Worse

Infrastructure debt has all the private credit problems plus several more. Ultra-long tenor: 15-25 year debt maturities against CDS markets that barely function beyond 10 years. Project-level risk from construction completion, regulatory approvals, and offtake agreements – none captured in sponsor-level CDS. Ring-fencing: infrastructure SPVs are legally isolated from sponsor; sponsor default does not trigger project default, so sponsor CDS pays out while your infrastructure loan keeps performing. Recovery is via ongoing cashflows not asset sales, making CDS settlement auctions completely misaligned.

Asian Toll Road: Addressing the Right Risks

Asian toll road project finance, $400m senior debt, 20-year maturity, availability-based payments from government. Risks in order of materiality: construction delay (bonded separately), traffic volume below forecast (subordinated equity absorbs first), government payment default (political risk insurance), FX volatility (currency hedged separately), sponsor financial distress (distant fifth – SPV is ring-fenced).

CDS on sponsor addresses risk #5. The other four account for 85% of potential loss scenarios.

Structuring team proposed: no CDS, use construction bonding, demand larger equity cushion, buy political risk insurance for government exposure. Total cost: 60bp over life of loan.

Alternative CDS approach would have cost 140bp annually (2,800bp cumulative over 20 years) to hedge a risk already mitigated structurally.

8.4 When Not Hedging Credit Is Rational

For most private credit and infrastructure portfolios, credit hedging is negative expected value once you properly account for ongoing premium costs, roll friction, basis risk, operational complexity, and opportunity cost of the hedging budget.

The rational approach is to accept unhedged credit risk and manage it via diversification (40+ positions, sector limits), structural seniority (senior secured, meaningful equity subordination, strong covenants), active monitoring (quarterly reviews, early warning systems, workout capability), conservative underwriting (lower leverage multiples, debt service coverage margins), and appropriate pricing – earning enough spread to absorb expected losses over the cycle.

> *Culturally this is hard. "We're unhedged" sounds reckless. "We're hedged with CDS" sounds prudent, even when the CDS does not work. Breaking this requires realism about what hedges actually achieve versus what they cost.*

Case Study 4: Infrastructure Equity – The Hedge That Made Everything Worse

Pension fund: $800m infrastructure equity fund, European renewables, regulated utilities, and transport assets. Year 3: portfolio down 8% on mark-to-market due to rising discount rates, though underlying cashflows performing in line with plan. Trustees demanded "downside protection."

GP implemented reluctantly: €200m notional CDS on utility and transport sector indices at 95bp annually = €1.9m per year, 3-year commitment.

Year 4-5: interest rates stabilised, infrastructure valuations recovered, NAV +12%. CDS hedge: year 4 lost €2.1m (spreads tightened), year 5 lost €1.7m (spreads tightened further). Cumulative hedge P&L: -€5.7m over 3 years. Underlying portfolio: +18% NAV appreciation. Every project operational and cashflow-positive.

The hedge protected against credit deterioration that never happened, cost €5.7m in cumulative losses, created quarterly P&L volatility from hedge marks, required GP time managing derivatives instead of assets, and reduced net IRR by 0.8% over fund life.

The hedge addressed trustee political risk (fear of being blamed for losses), not economic risk. The portfolio was never in credit distress. The right answer was educating trustees that infrastructure equity has mark-to-market volatility but cashflow stability. The actual answer was implementing an expensive hedge to make trustees comfortable, then destroying value over a multi-year period proving it was unnecessary.

9. Designing Credit Protection That Survives Stress

If credit hedging is necessary – due to concentration, leverage, or genuine risk transfer objectives – the challenge is designing protection that functions when markets break, not just when conditions are benign.

9.1 Align Hedge Horizon with Loss Realisation, Not Reporting

A common design failure is matching hedge tenor to reporting and governance cycles rather than to the timeline over which credit losses are realised. Quarterly effectiveness testing drives 3-5 year maturities. Annual budget cycles drive rolling 12-month hedges. But public corporate default to final recovery takes 18-36 months, private credit impairment to resolution takes 24-48 months, and infrastructure distress to recovery takes 36-60 months. The timelines rarely align.

A more robust approach matches hedge maturity to expected loss resolution: longer-dated protection (7-10 years) for public IG corporates where available; accepting that CDS is structurally mismatched for private credit; hedging through construction and ramp-up for project finance then letting exposure run; short-dated protection for distressed or HY where losses emerge quickly. This approach sacrifices tidy hedge effectiveness testing. It improves alignment with economic loss.

9.2 Pre-Define Failure States

Credit hedges underperform because success is defined after the fact. Effective design starts with a specific question: what outcome are we trying to prevent, and how does it actually occur?

For a leveraged fund with redemption risk, the failure state is spread widening triggering NAV decline leading to forced selling – the design objective is liquidity, not default protection, pointing toward short-dated index CDS sized to offset spread shocks that trigger redemptions. For an insurance general account, the failure state is rating migration driving regulatory capital pressure – the approach is CDS structured to meet capital offset rules. For a family office with concentration risk, the failure state is a single large position defaulting – the approach is single-name CDS on the concentrated position only, exited if exposure is reduced.

9.3 Accept Imperfect Hedges Over Fragile Ones

A hedge that functions imperfectly under stress is preferable to one that fails completely. Fragile hedges typically share common features: high correlation in normal markets, structural complexity, reliance on tight liquidity, leverage or collateral sensitivity, and concentrated counterparty exposure. These attributes improve back-tests and presentations. They degrade under stress.

More robust hedges tend to look inefficient in benign markets: looser correlation, simpler structures, demonstrated stress-period liquidity, minimal funding requirements, diversified counterparties. They cost more in quiet periods. They function when conditions deteriorate.

9.4 Stress-Test the Hedge, Not Just the Portfolio

Most portfolios are stress-tested. Few hedges are. Questions that matter: What were bid-offer spreads in prior stress periods? Can gains be monetised, not just marked? What is settlement timing versus liquidity needs? How exposed are you if the dealer deteriorates? How does the hedge behave if correlation collapses? Can the hedge be managed under reduced staffing and time pressure? How are large hedge losses or gains explained internally?

> *If a hedge fails these tests, it should be redesigned or not implemented.*

9.5 When "Protected" Is Worse Than "Unhedged"

In some cases, carrying a hedge creates worse outcomes than remaining unhedged. Common consequences include false confidence leading to higher risk-taking, opportunity cost from premium spend versus holding liquidity, increased political risk if a hedge fails visibly, and strategic rigidity created by hedge roll commitments.

For some portfolios, the most robust decision is to forgo hedging and instead tighten concentration limits, increase diversification, hold additional liquidity, and improve monitoring. This requires explanation and control. It avoids the perception of protection while preserving flexibility when conditions change.

10. Structural Diagnostic – Is Your Credit Hedge Real?

Five uncomfortable questions. If you cannot answer at least four clearly and defensibly, your credit hedge is unlikely to survive stress.

Question 1: Can you explain, specifically, what loss event this hedge prevents?

Not "credit risk" or "spread widening" – the specific scenario with names, amounts, timeline, and consequences, and how hedge payoff matches that scenario.

Good answer: "If our top 5 BBB-rated positions migrate to BB simultaneously, we face $45m in additional regulatory capital charges. This hedge offsets $35m of that charge, keeping us inside our capital buffer."

Bad answer: "It provides protection against credit deterioration in our portfolio."

Question 2: Have you stress-tested this hedge in actual historical crises?

Not back-tested returns – stress-tested operational functionality. Pull actual transaction data from March 2020, September 2008, and August 2011 for your specific hedge instrument. Look at bid-offer spreads, trade sizes, and settlement times.

Good answer: "In all three periods, this specific index traded with bid-offer spreads of 8-12bp. We could have monetised in 2-3 days with documented dealer capacity."

Bad answer: "Our risk models show 0.82 correlation in stressed scenarios."

Question 3: If you monetised this hedge tomorrow, what percentage would you lose to slippage?

Good answer: "Index CDS: 5-8bp in current market, potentially 15-20bp in stress based on March 2020 data. Acceptable for the protection provided."

Bad answer: "It's liquid, we can exit anytime."

Test: get actual dealer quotes to unwind the hedge today. Not indicative levels – actual executable bids. If you are shocked by the slippage, imagine 3x worse in a crisis.

Question 4: What happens if your hedge counterparty's stock drops 50%?

Good answer: "We are cleared through LCH, no counterparty exposure." Or: "We have CSA agreements with daily margining and segregated collateral with a third-party custodian."

Bad answer: "Our dealer is a major bank; they're not going to fail."

> *Lehman Brothers was a major bank. Calculate your mark-to-market exposure to each hedge counterparty under stressed scenarios. If you do not know, you have unhedged counterparty risk inside your hedge.*

Question 5: Would you personally pay this ongoing cost if it was your own money?

Good answer: "Yes. I have concentrated exposure I cannot diversify, leverage I cannot reduce, and regulatory requirements I must meet. The cost is justified."

Bad answer: "It's institutional money, governance requires it."

Calculate cumulative hedge costs over 10 years. Compare to expected loss from unhedged portfolio over the same period. If hedge costs exceed expected losses by 2x or more, and you would not pay this personally, why are you recommending it institutionally?

Red Flags – When Your Hedge Is Actively Harmful

Red Flag 1: The hedge has never been monetised. If you've held a hedge for 3+ years and never exited it, you don't know if it works.

Red Flag 2: Nobody can explain the correlation assumption. "It's in the model" means nobody understands your hedge.

Red Flag 3: Hedge sizing changed to pass effectiveness tests. You've optimised for accounting, not economics.

Red Flag 4: The hedge costs more than the illiquidity premium you're earning. Paying away 83% of excess return to hedge an inherent risk of the asset class is incoherent.

Red Flag 5: You've never calculated total cost of ownership. Premium + roll costs + bid-offer + collateral + operational overhead + opportunity cost.

Red Flag 6: The hedge survived a portfolio review but not independently. If the hedge cannot justify itself outside the context of "we need to do something," it's probably not the right something.

Case Study 5: The Honest Diagnostic That Killed a Hedge Programme

$8bn insurance general account, BBB average credit quality, 6-year duration. Five-year-old hedge programme: $2.5bn notional 5-year CDS on iTraxx and CDX, rolled annually. New CIO commissioned independent review.

Question 1 (Specific loss event): FAILED. Team could not articulate the specific scenario the hedge prevented. Question 2 (Historical stress test): FAILED. No analysis of March 2020, 2011, or 2008 execution feasibility. Dealers confirmed they would have charged 40-60bp bid-offer on this size in March 2020.

Question 3 (Monetisation slippage): FAILED. Current exit cost: $18-25m in slippage. Had never been calculated. Monetisation had never been attempted in 5-year programme history.

Question 5 (Personal money test): FAILED. Investment team admitted they would not pay these costs personally. Programme existed because "previous CIO implemented it."

Cumulative hedge costs over 5 years: $142m. Realised benefits: $8m. Net cost: $134m. Expected credit losses over same period: $95m. The programme had cost 41% more than unhedged expected losses.

CIO decision: terminated entire hedge programme. Redeployed capital to enhanced credit surveillance ($20m), dedicated workout team ($15m), remainder to investment capital for loss absorption. Five years later: cumulative savings vs continuing hedge: $168m. Net position: $133m better off with no hedge than with it.

11. The Real Conversation: When Credit Risk Should Stay Unhedged

The assumption is always that credit risk exists, therefore we must hedge it. But hedging is a cost-benefit decision, not a moral imperative. Sometimes the right hedge ratio is zero.

11.1 Accepting Economic Reality vs Regulatory Fiction

Regulations often require institutions to measure, report, and appear to manage credit risk. This creates pressure to hedge even when hedging is economically irrational. The regulatory fiction: credit risk can be precisely measured; credit hedges reduce risk; hedged portfolios are safer than unhedged. The economic reality: credit models miss tail risk, correlation shifts, and recovery uncertainty; credit hedges often transfer one risk while creating others; "safer" on regulatory metrics often means more fragile in actual stress.

The real conversation with regulators, rating agencies, or boards: "We have analysed available hedging approaches and concluded they would destroy value. Available hedges cost 140bp annually for protection with 60% expected effectiveness. Our portfolio earns 200bp above risk-free specifically for taking unhedgeable credit risk. Hedging would pay away 100% or more of excess return to achieve partial protection. We have adequate capital to absorb expected credit losses stress-tested to 2008 levels. We are choosing to stay unhedged."

11.2 Cost of Honesty vs Cost of Pretence

Pretending to manage credit risk via ineffective hedges has costs. Direct costs include ongoing hedge premium, roll costs, bid-offer and slippage, operational overhead, and opportunity cost of capital. Indirect costs include false confidence leading to higher risk-taking, complexity creating operational risk, governance attention diverted from real risks, and career risk when the hedge inevitably fails. Cultural costs include normalising compliance over effectiveness, rewarding appearance over substance, and punishing honesty.

> *The cost of honesty is immediate and visible: "We are unhedged, we might lose money from credit events." The cost of pretence is delayed and diffuse: ongoing value destruction, false confidence, eventual larger losses when the hedge fails. Most organisations choose pretence because the pain is distributed and deniable.*

11.3 Why Doing Nothing Is Often the Right Hedge

"Doing nothing" sounds reckless. It is often the most sophisticated choice. When you are adequately diversified (50+ positions, <3% in largest single-name), expected credit losses over a decade are 80-120bp while available hedges cost 120bp or more annually at 60-70% effectiveness – doing nothing is economically superior, and the diversification is your hedge. When you have structural seniority – senior secured positions with meaningful subordination, infrastructure debt with 1.4x or greater debt service coverage – adding CDS on top is redundant protection. When you can absorb the losses with long time horizon,

no forced selling, and adequate capital reserves, you do not need to transfer risk you can absorb.

Accept the risk if you are being paid for unhedgeable risk – illiquidity premium, complexity premium, concentration premium. These premiums exist because the risk cannot be efficiently hedged. Attempting to hedge pays away the premium you are earning.

11.4 Reframing Unhedged as a Positive Choice

Language matters. "We manage credit risk through portfolio construction rather than derivatives" means diversification requirements, conservative underwriting, active monitoring, structural protection, and adequate capital reserves. "We have analysed credit hedging and concluded the costs exceed the benefits" means we have done the work, understand the trade-offs, and are making an informed choice. "Our returns include compensation for unhedgeable credit risk, and hedging would destroy value" means we understand our return sources and are optimising for stakeholder outcomes, not optics.

11.5 When Unhedged Becomes Indefensible

Staying unhedged is not always right. It becomes indefensible when concentration exceeds capital (single position >10% of capital, or top 5 positions >40% of capital), when leverage amplifies credit risk and creates forced selling, when you are not being compensated for unhedgeable exposure (credit spread over benchmarks <100bp), when regulatory and rating agency requirements are binding constraints, or when stakeholders genuinely cannot tolerate the outcomes.

In these cases, staying unhedged is choosing institutional friction over fiduciary duty. Hedge, or change the constraints – deleverage, diversify, exit the position, educate stakeholders.

12. Implications for Portfolio Construction

Credit hedging fails when it is treated as a separate problem from portfolio construction. The two are inseparable.

12.1 Credit Hedging as Capital Allocation Decision

Every dollar spent on credit hedging is a dollar not available for holding more liquid instruments (actual liquidity, not derivative liquidity), investing in higher-quality credits (structural protection vs derivative protection), building loss reserves (real capital vs contingent protection), or enhancing monitoring and systems (prevent losses vs transfer losses).

Option A: Hedge with CDS	Option B: Build quality and liquidity
$500m BBB portfolio, 30% notional CDS hedge ($150m)	$475m portfolio (smaller, reallocating hedge budget)
Annual cost: 95bp = $1.425m per year	Upgrade $75m from BBB to A rated: 25bp yield give-up
Expected effectiveness: 65%	Hold $25m in cash/T-bills: actual 5% liquidity buffer
True cost per unit of protection: 146bp	Invest $200k per year in credit surveillance
10-year cumulative cost: $14.25m	$538k per year to additional loss reserves
Counterparty risk, basis risk, operational complexity	No counterparty risk, no basis risk, structurally more resilient

12.2 Hedging Credit Exposure vs Hedging Credit Outcomes

Credit exposure is the notional size of credit risk in the portfolio. Credit outcome is the specific bad event you are trying to prevent. Most hedges target exposure. Better hedges target outcomes.

For a pension fund with liability-driven investing: exposure-based hedging costs $12m annually and protects against spread widening and portfolio MTM volatility. Outcome-based thinking recognises the real risk is inability to meet benefit payments in 15-20 years – the relevant approach is upgrading credit quality to A-rated, extending duration to match liabilities, accepting spread volatility (not economic risk for hold-to-maturity), and building a dedicated reserve for expected defaults over 20 years. Cost: $4m annually, 66% cheaper, better aligned to actual risk, no derivative complexity.

> *The outcome lens asks: "What specific bad future are we preventing?" Not "reduce credit VAR" – real outcomes like "unable to meet pension obligations in 2043" or "forced to sell assets at 40% loss during redemption spike."*

12.3 When to Hedge, When to Diversify, When to Accept

The decision sequence:

1) **can you diversify?** If yes, diversify first – cheaper than hedging, no basis risk, no counterparty risk. If no,
2) **can you upgrade quality or gain structural seniority?** If yes, move up the capital structure. If no, can you absorb the loss? If yes, accept unhedged with monitoring, reserves, and surveillance. If no,
3) **can you get effective hedging at reasonable cost?** If yes, hedge selectively using the simplest structure. If no, you have a fundamental problem: reduce position size, exit entirely, change constraints, or accept you are hedging for governance not economics.

> *Most portfolios jump straight to "hedge" without working through the first three steps. This is why so many hedges are ineffective.*

12.4 Reframing Hedge Effectiveness Around Outcomes, Not Accounting

Current definition of hedge effectiveness: accounting correlation above 80%, retrospective testing, quarterly verification. This optimises for IFRS 9 compliance, not economic protection. Better definition: a hedge is effective if it meaningfully improves the probability of achieving portfolio objectives when credit stress occurs.

This requires defining the specific portfolio objective, the credit stress scenario that threatens it, and how the hedge changes the probability of achieving it. The result is an answerable cost-benefit question: is 1.2% of returns worth 5 percentage points of lower drawdown probability? Compare to current effectiveness testing: "The hedge has a correlation of 0.83 over the last quarter, exceeding our 0.80 threshold." That tells you nothing about whether the hedge is worth having.

Conclusion

Credit hedging persists because it is defensible, not because it reliably works. It provides governance comfort, career cover, and the appearance of control, while often leaving the underlying economic risk largely intact.

The issue is not that hedging is always wrong. The issue is that most institutional credit hedges are designed for spread volatility, not default loss, and they often fail precisely when liquidity disappears and losses become real.

For many portfolios, the more robust solution is simpler: diversify more, move up the capital structure, improve underwriting, hold more liquidity, and build reserves instead of paying away value through ongoing hedge costs.

Where hedging is necessary, it should be used only when it clearly transfers the specific economic loss the portfolio is trying to avoid. If it does not, the honest answer is not to pretend otherwise — it is to accept the risk, change the portfolio, or change the constraint.

The right question is not whether a hedge looks prudent. It is whether, under stress, it actually improves outcomes .

> *If the hedge does not improve outcomes under stress, it is not a hedge — it is an expense.*

Further Reading

Para Bellum Advisors publishes practitioner papers and CIO Briefs focused on real-world portfolio construction, risk and capital efficiency:

www.parabellumadvisors.com/insights/

About Para Bellum Advisors

Para Bellum Advisors is an independent advisory firm specialising in derivatives structuring, structured finance, balance sheet efficiency, and capital optimisation for institutional investors and corporate treasury teams.

The firm works with lean investment and treasury teams managing complex, multi-asset exposures – long-dated assets, illiquid portfolios, and non-standard risk profiles – where structural precision makes a material difference to outcomes.

Its engagements typically involve designing and re-engineering hedges across FX, rates, credit, equity, and volatility; identifying and releasing trapped capital; and providing embedded structuring capability where permanent headcount is neither practical nor warranted.

Para Bellum does not distribute products or earn transaction volume. Its value is in structure: how exposures are designed, how capital is consumed, and how portfolios behave when conditions deteriorate.

The firm is practitioner-led, drawing on three decades of experience across trading, structuring, and portfolio management in banks, asset managers, and institutional balance sheets in Asia-Pacific and global markets.

For more information: www.parabellumadvisors.com

For discussion or enquiries: mike.duncan@parabellumadvisors.com

Connect on LinkedIn: www.linkedin.com/in/mikeduncan-structuring/