Practitioner Paper

The Credit Hedge Illusion - A Practitioner's Guide to What Actually Works

When Protection Costs More Than the Risk It's Meant to Transfer

Version 1, December 2025

By Mike Duncan, Para Bellum Advisors

www.parabellumadvisors.com

Executive Summary

Many institutional credit hedges are expensive box-ticking exercises. They look professional in risk reports, satisfy governance committees, and create the appearance of prudent risk management. But when credit markets actually blow up – when defaults cluster, liquidity vanishes, and correlations spike to one – these hedges typically fail to deliver meaningful protection.

This isn't a theoretical concern. Across three decades structuring credit protection for insurers, pension funds, asset managers, and family offices across Asia-Pacific, I've watched the same pattern repeat: sophisticated investors pay ongoing premiums for CDS hedges that protect against the wrong risks, fail to pay out when needed, or generate basis losses that dwarf the protection they were meant to provide.

The core problem is simple: **credit hedging as commonly practised addresses spread volatility and accounting optics, not actual economic loss from default or severe credit deterioration.**

Real credit losses are structural events. A private credit borrower defaults and goes through restructuring. An infrastructure loan gets impaired and sits non-performing for eighteen months. A high-yield issuer blows out and recovery takes years. These are slow-motion train wrecks with binary outcomes – you either get your money back or you don't.

Yet most credit hedges are designed around mark-to-market spread movements, index tracking, and quarterly hedge effectiveness tests. They're optimised for governance comfort and regulatory reporting, not for transferring the actual economic risk of credit loss.

This paper examines:

- What credit risk actually consists of (and why most hedges only address one component)
- Why CDS dominates despite systematic failures
- The hidden costs that make credit hedging negative sum for many portfolios
- Why index and single-name CDS fail in different ways but produce the same outcome
- How credit hedging becomes a liquidity problem precisely when you need it most
- What actually works when the goal is transferring economic risk, not managing optics

Key thesis: For most institutional portfolios, credit hedging increases capital drag, creates false confidence, and often worsens outcomes during actual credit stress. The honest conversation isn't "how do we hedge credit risk?" – it's "what specific credit outcome are we trying to prevent, and does this hedge actually achieve that?"

If you're running private credit, infrastructure debt, insurance general accounts, or multi-asset portfolios with meaningful credit exposure, this paper is designed to help you distinguish between hedges that transfer risk and hedges that transfer accountability.

Who This Paper Is (and Isn't) For

This paper is written for institutions that hold credit risk over meaningful horizons and are accountable for outcomes, not optics.

This paper is for:

- CIOs, portfolio managers, and heads of risk
- Insurance general accounts and balance-sheet investors
- Private credit, infrastructure, and long-dated credit portfolios
- Multi-asset portfolios where credit loss is absorbed, not traded
- Teams willing to question whether hedging improves outcomes at all

This paper is not for:

- Trading desks managing short-term P&L or daily VAR
- Portfolios where credit risk is actively turned over
- Teams seeking a product recommendation or hedge ratio
- Situations where the decision to hedge has already been made
- Readers looking for validation that "doing something" is always prudent

The arguments that follow may lead to uncomfortable conclusions, including that some credit risks are better owned than hedged.

That is intentional.

Legal Notice – Attribution & Reuse Notice

Table of Contents

1. Introduction: 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
- Often worsens outcomes during actual credit events

This paper starts from a simple distinction: **credit exposure is not the same as credit outcome**.

A hedge that dampens spread volatility can still fail to protect against default loss. A hedge that passes accounting tests can still create liquidity stress at the worst possible time. A portfolio can be "hedged" on paper and fully exposed in reality.

The purpose of this paper is not to argue that credit risk should never be hedged. It is to examine **when credit hedging transfers real economic risk, when it merely transfers accountability, and when not hedging is the most rational choice available**.

The analysis that follows focuses on behaviour under stress, not elegance in normal markets. It treats credit loss as a capital and liquidity problem, not a pricing exercise.

It evaluates hedges by what they deliver when credit actually breaks – not by how they look in a risk report.

The aim is not complexity.

It is clarity.

2. What Credit Risk Actually Is

Before assessing whether a credit hedge works, you need clarity on what you're hedging. Credit risk isn't 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 or debtor-friendly
- infrastructure assets where post-default viability is uncertain
- 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 / Spread Risk

Mark-to-market volatility from credit migration. This matters or doesn't 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.

This is most visible in lower-rated credit, single-name CDS under stress, bespoke private credit, and parts of emerging market credit during local currency or funding crises.

Most credit hedges only address one of these risks – usually downgrade/spread risk. They typically do not address recovery risk, liquidity risk, or correlation risk.

Yet when credit portfolios actually lose serious money, it's typically because:

- defaults clustered (correlation)
- recoveries were worse than modelled (recovery)
- positions couldn't be exited (liquidity)
- hedges couldn't 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. The incentive was to hedge spread risk to reduce capital charges.

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

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 – because the loss mechanism was single-name default with poor recovery in an illiquid restructuring.

Portfolio would have been better 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.

Understanding why CDS dominates helps explain why most credit hedging fails.

3.1 Liquidity Optics

CDS trades in size. Quotes are available. Market makers exist. Bloomberg has pages of data. 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 the conditions where the hedge is least needed. It evaporates in the 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.

The Hard Truth

CDS hedging primarily manages reporting risk, not economic risk.

It is effective at managing:

- quarterly P&L volatility
- regulatory capital calculations
- board and consultant expectations
- audit and governance requirements

It is far less effective at managing:

- realised default losses
- recovery shortfalls

- liquidity crises
- correlation events

For portfolios that mark to market daily and face redemptions or margin calls, managing reporting risk may be rational – mark-to-market losses quickly become economic losses.

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.

The relevant question is not *"should we use CDS?"*
It is *"what risk are we actually trying to manage, and does CDS manage it?"*

In many institutions, that question has never been asked truthfully.

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, depending on conditions. 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 and tends to surface precisely when credit stress becomes idiosyncratic rather than systemic.

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
- often produce outcomes misaligned with portfolio-level losses

By the time the hedge pays, bonds may already be marked down 40–70%, liquidity constrained, and capital tied up in workouts lasting years.

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. Monetising gains or exiting positions became prohibitively expensive.

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. The risk committee mandated 30% notional hedge coverage using five-year iTraxx Main.

The hedge was rolled semi-annually to maintain constant maturity.

Between 2015 and 2019:

- Eight rolls were executed
- Average roll cost was ~12bp per roll
- Cumulative roll cost reached 96bp
- Underlying credit portfolio delivered ~35bp of net positive performance

Hedge P&L over the period was -88bp. The combined outcome was -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 implicit justification was protection against a future crisis that never materialised.

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. Protection is bought on a basket of names, equally weighted and highly liquid. Simple.

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. The hedge relies on an assumed correlation between portfolio behaviour and index behaviour.

That correlation is neither stable nor sufficient.

In normal markets, correlation might sit around 0.65–0.75. Adequate for optics, inadequate for economics. You are 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. 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. Correlation solved.

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. This is not trivial infrastructure.

For a mid-sized portfolio, this typically means dedicated operational resources and external legal and system costs before any protection is purchased.

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.

Protection can usually be put on. Getting it off – or monetising gains – is far less certain.

Counterparty concentration compounds the risk. As spreads widen, dealers owe increasing amounts under CDS contracts. In systemic stress, collateral and close-out assumptions become unreliable.

Wrong-way risk is common. Protection is often bought from counterparties whose credit is correlated with the exposure being hedged. Credit stress then weakens both the hedge and the counterparty simultaneously.

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.

As a result, CDS hedges are implemented that satisfy governance – *we are hedged* – while introducing cost, complexity, and questionable protection.

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

For many organisations, the answer is no. Hedges are implemented anyway because not hedging feels irresponsible.

The result is expensive optics-driven activity.

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 ≠ 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. A £50m position can suddenly require £3m of market impact to exit – 600bp of slippage on an asset that may already be down 500–800bp in spread terms.

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
- Attempting to exit both simultaneously hits the bid side of two stressed markets

The protection exists on paper. The ability to capture it does not.

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.

What actually happened:

- 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.

This is not a crisis-specific phenomenon. Credit stress reliably produces the same conditions:

- Liquidity tightens simultaneously across cash and derivatives
- Demand for protection spikes at the same time for all participants
- Dealers ration balance sheet to the highest-priority counterparties
- Everyone else faces punitive pricing or no market at all

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.

Scenario 1: Leveraged, mark-to-market portfolio
The portfolio is down 8% from spread widening at 2.5× leverage. Margin calls arrive immediately. The CDS hedge shows gains, but:

- It sits with a different counterparty (no netting benefit)
- Monetisation requires slippage
- Settlement is T+2 while margin is due today

The portfolio is forced to liquidate bonds at distressed levels despite being "hedged." The hedge did not fail – it simply arrived too late to help.

Scenario 2: Insurance company under regulatory capital pressure
Spread widening increases capital charges. Exposure must be reduced or capital raised.

The CDS hedge helps regulatory calculations if structured correctly, but:

- Premium drag reduces return on capital
- Basis risk limits offset
- Forced deleveraging still requires exiting both bonds and hedges at poor levels

The hedge improves reported ratios. It does not resolve the liquidity constraint driving the decision.

Scenario 3: Unlevered, patient capital
A family office holds private credit with no mark-to-market pressure and no forced selling.

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. The CDS hedge showed ~$15m of gains.

Simultaneously:

- Equities fell ~12%
- Redemptions reached ~8% of AUM
- 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

The hedge made money. It failed to solve the liquidity problem when decisions mattered.

A simpler approach – holding more cash and accepting spread volatility – would likely have produced better outcomes.

The hedge created the illusion of liquidity. When stress arrived, that illusion mattered more than the hedge ever did.

The credit hedge created the illusion of liquidity that evaporated under stress, leading to worse decisions than if they'd 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're trying to achieve. But 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're senior to someone else who absorbs losses before you do.

How it works: Instead of buying CDS on your BBB corporate exposure, you own senior secured debt while someone else (equity, mezzanine, junior unsecured) sits below you. Default happens; you're first in line for recoveries. The subordination is structural, not contractual.

Why it actually works:

- No counterparty risk (unlike CDS)
- No liquidity risk (you're not trying to monetise a derivative)
- No roll costs (subordination is permanent feature of capital structure)
- No basis risk (you're protected by actual capital, not correlation assumptions)

The cost: Lower yield. Senior secured typically yields 100-200bp less than unsecured exposure on same credit. But that yield give-up is transparent, one-time, and permanent – versus ongoing CDS premium that bleeds forever.

When it makes sense:

- Private credit and direct lending (you control structure)
- Infrastructure debt (asset-level subordination via project finance)
- Insurance general accounts (long hold periods, value certainty over yield)
- Any portfolio where you can move up the capital structure instead of hedging

When it doesn't:

- Public markets where you can't dictate structure
- When yield requirements are absolute
- When you're stuck with existing unsecured positions

This isn't hedging – it's portfolio construction. But portfolio construction that avoids risk is better than hedging that doesn't 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 is different: for each specific credit exposure, what's the actual loss scenario and how would we prevent it?

Example 1: EM Corporate Exposure

Portfolio approach: Buy EM CDX index, hedge 30-40% of notional, call it protected.

Asset-level approach:

- Exposure is $100m to Brazilian utility, dollar-denominated bonds
- Real risk isn't default – it's currency crisis preventing dollar debt service
- Hedge the specific risk: long BRL/USD put spreads, costs 80bp/year
- Much cheaper than CDS, actually addresses the failure mode

Example 2: Project Finance Exposure

Portfolio approach: Infrastructure debt portfolio, buy CDS on parent sponsor company.

Asset-level approach:

- Exposure is to SPV with ring-fenced cashflows from availability payments
- Real risk is construction delay or technical failure, not sponsor default
- Credit protection via construction bonding, technology guarantees, or subordinated equity from sponsor
- Structural, upfront, survives sponsor bankruptcy

Example 3: Leveraged Loan Exposure

Portfolio approach: Buy loan CDS index.

Asset-level approach:

- Real risk is covenant-lite structure allowing deterioration without default
- Protection via:
    - Demanding actual maintenance covenants
    - Requiring amortisation schedules that deleverage over time
    - Accepting lower leverage multiples at origination
    - Walking away from deals that don't meet standards

Again, this is underwriting discipline, not hedging. But it prevents losses rather than trying to transfer them after the fact.

7.3 Where Insurance-Style Thinking Beats Trading Logic

CDS is a trading instrument. It's designed for mark-to-market books, active positioning, and tactical hedging. Most institutional credit portfolios aren't trading books – they're insurance-style portfolios: buy, hold, collect income, manage binary default risk over multi-year horizons.

Insurance-style protection looks different:

Characteristic 1: Binary payouts, not MTM protection

Insurance pays out on defined events (default, restructuring below X recovery, payment failure). You don't care about spread movements between purchase and payout.

CDS is priced and traded based on spread movements. You're paying for liquidity and optionality you don't need.

Characteristic 2: Long-dated, non-recourse

Insurance contract is set-and-forget. You pay premium upfront or over known schedule. No collateral, no margin calls, no mark-to-market disputes.

CDS requires ongoing collateral management, dealer relationships, and active monitoring. Operational overhead for protection that might never pay out.

Characteristic 3: Diversification assumed in pricing

Insurance companies' price assuming they'll have 1000 exposures and law of large numbers applies. They can offer cheap protection per unit because losses are uncorrelated.

CDS is priced assuming correlation risk, jump risk, and concentrated exposures. You're paying for worst-case scenarios even when your portfolio is diversified.

What this means practically:

For long-dated, hold-to-maturity credit portfolios, the ideal protection looks more like credit insurance than CDS:

- Upfront or level premium (no mark-to-market volatility)
- Defined trigger events (default, restructuring, payment failure)
- Par payout or recovery shortfall payment (not spread-based)
- No collateral or margin requirements
- Multi-year commitment (no roll risk)

This exists in specialty insurance markets (political risk insurance, trade credit insurance, surety bonds) but rarely gets used for institutional portfolio hedging because:

- Investment teams think in CDS terms
- Insurance solutions don't fit hedge accounting frameworks
- Consultants don't know how to analyse insurance-based protection
- Governance committees want "market standard" (i.e., CDS)

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 don't 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.

Robust imperfection: Structural subordination, insurance-based protection, or simply accepting unhedged exposure with appropriate portfolio construction.

Robust imperfection often wins because:

- No operational complexity
- No ongoing costs if risk doesn't materialise
- No liquidity risk when you need protection
- No basis risk between hedge and exposure

- No counterparty risk

The cost is imperfect correlation and potentially wider distributions of outcomes. But if your portfolio can absorb those wider distributions (long time horizon, no forced selling, adequate capital), robust imperfection is economically superior to fragile precision.

Framework for choosing:

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

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

If you answer NO to 3+, CDS *might* be appropriate:

1. Daily MTM and redemptions (mutual fund, hedge fund)
2. Leveraged portfolio with margin requirements
3. Concentrated exposures you cannot diversify
4. Tactical positioning around credit views
5. Regulatory capital optimisation is primary goal
6. Dedicated derivatives team and infrastructure

Most portfolios fall into the first category but implement hedges designed for the second. This is the central dysfunction.

8. Private Credit and Infrastructure: The Hardest Case

If CDS hedging is problematic for liquid IG corporates, it's 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 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 and restructuring, not liquidation

Now try to hedge this with CDS:

- Standardised contracts on public reference entities
- Daily mark-to-market with collateral requirements

- Liquid trading (in theory) requiring active management
- Correlation assumptions to proxy hedges since direct hedges don't exist
- Payout based on bond recovery auctions, not actual loan recoveries

Every single dimension is mismatched.

Specific failures:

Valuation mismatch: Your private loan is marked quarterly at par (assuming no credit event). Your CDS hedge marks daily and shows P&L volatility from spread movements unrelated to your loan's actual credit quality. This creates accounting volatility you were trying to avoid.

Recovery mismatch: Private credit recoveries are typically higher than public bond recoveries (stronger covenants, closer monitoring, ability to renegotiate). CDS pays out based on public bond recovery auctions. You might recover 70 cents, your hedge assumes 40 cents, you've paid for protection against losses that don't match reality.

Trigger mismatch: Your loan might go non-performing but not technically default (interest reserve draws, payment holidays, covenant waivers). CDS doesn't trigger. Or your loan might cure after a brief default while CDS settles at depressed recovery. Outcomes diverge completely.

Correlation mismatch: You're hedging private sponsor-backed LBO loan with CDS on public high-yield index. Correlation in normal times: maybe 0.40. Correlation when you need hedge: might spike to 0.70, might collapse to 0.20 depending on whether stress is sector-wide or idiosyncratic. You have no idea which.

8.2 Why Mismatch Is Structural

This isn't a "find better hedge" problem. It's a fundamental impossibility:

Private credit works *because* it's 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
- Relationship-based restructurings when things go wrong

Trying to hedge this with liquid instruments means either:

1. **Paying away your illiquidity premium** to gain liquidity you explicitly chose not to have, or
2. **Accepting massive basis risk** that renders the hedge performative risk management (i.e. useless)

There's no magic third option.

Real example: 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/year
- Illiquidity premium earned on portfolio: 220bp over liquid HY
- Correlation between portfolio and HY index: 0.35-0.55 depending on period
- 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 portfolio diversification (40 loans vs 35)
- Enhanced monitoring and covenant packages
- Built larger loss reserves
- Accepted concentration limits
- Educated investors on why hedging was value-destructive

The approach to follow. Cost them some investors who wanted the comfort of "hedged" even if the hedge didn't work.

8.3 Infrastructure Debt: Even Worse

Infrastructure debt has all the private credit problems plus:

Ultra-long tenor: 15-25 year debt maturities. CDS markets barely function beyond 10 years. You'd need to roll 5-year CDS three times to cover the exposure, compounding all the roll costs and basis issues.

Project-level risk: Cashflows depend on construction completion, regulatory approvals, offtake agreements, availability payments. None of these are captured in sponsor-level CDS.

Ring-fencing: Infrastructure SPVs are legally isolated from sponsor. Sponsor default doesn't trigger project default. Sponsor CDS pays out, your infrastructure loan keeps performing. You've hedged the wrong thing.

Recovery through asset operation: Infrastructure assets don't liquidate – they get operated by a new sponsor or taken over by lenders. Recovery is via ongoing cashflows, not asset sales. CDS settlement auctions assume bond recoveries from liquidation. Completely different recovery process.

Currency and regulatory risk: Cross-border infrastructure often has local currency revenue, dollar debt service, and regulatory risk around tariffs/tolls. The credit risk is inseparable from currency and political risk. CDS only addresses credit.

Case study: Asian toll road project finance, $400m senior debt, 20-year maturity, availability-based payments from government.

Risks in order of materiality:

1. Construction delay (bonded separately)
2. Traffic volume below forecast (subordinated equity absorbs first)
3. Government payment default (political risk insurance)
4. FX volatility affecting debt service (currency hedged separately)
5. 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 that was already mitigated structurally.

The right answer was portfolio construction and insurance-based solutions, not CDS.

8.4 When Not Hedging Credit Is Rational

Here's the uncomfortable truth: 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
- Opportunity cost of hedging budget

The rational approach is often:

Accept unhedged credit risk and manage it via:

1. **Diversification** – 40+ positions, <3% in any single-name, sector limits
2. **Structural seniority** – senior secured, meaningful equity subordination, strong covenants
3. **Active monitoring** – quarterly reviews, early warning systems, workout capability
4. **Conservative underwriting** – lower leverage multiples, debt service coverage margins, stress testing
5. **Appropriate pricing** – earn enough spread to absorb expected losses over cycle

This requires acknowledging to investors, boards, and risk committees that the portfolio is *unhedged* – and explaining why that's the right economic choice.

Culturally, this is hard. "We're unhedged" sounds reckless. "We're hedged with CDS" sounds prudent, even when the CDS doesn't work.

Breaking this requires realism about what hedges actually achieve versus what they cost. For illiquid credit portfolios, the answer is usually: hedging destroys value, accept the risk and manage it properly instead.

CASE STUDY 4: Infrastructure Equity – The Hedge That Made Everything Worse

Pension fund commitment: $800m infrastructure equity fund investing in European renewables, regulated utilities, and transport assets. 12-year fund life, J-curve expected.

Year 3: Portfolio down 8% on mark-to-market due to rising discount rates, though underlying asset cashflows performing in line with plan. Pension trustees panicked by NAV decline, demanded "downside protection."

Fund GP analysis showed:

- NAV decline was valuation, not credit deterioration
- All projects performing to forecast
- Cashflow-based returns still projected at 9.5% IRR
- Declines would reverse as rates stabilised

Trustees insisted on credit protection anyway. GP reluctantly implemented:

- €200m notional CDS on utility and transport sector indices
- Cost: 95bp annually = €1.9m/year
- 3-year commitment before reassessment

What happened:

Year 4-5: Interest rates stabilised, infrastructure asset valuations recovered. NAV back to par, then +12% as energy transition tailwinds strengthened.

CDS hedge:

- Year 4: Lost €2.1m (spreads tightened as credit markets recovered)
- Year 5: Lost €1.7m (spreads tightened further)
- Cumulative hedge P&L: -€5.7m over 3 years

Underlying portfolio: +18% NAV appreciation over same period, all projects 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 derivative positions instead of assets
- Reduced net IRR by 0.8% over fund life

Deeper problem: The hedge addressed trustee *political* risk (fear of being blamed for losses), not economic risk. The portfolio was never in credit distress. The NAV decline was interest rate-driven valuation, not credit impairment.

Right answer: educate trustees that infrastructure equity has mark-to-market volatility but cashflow stability, so MTM declines don't indicate credit problems. Accept the volatility.

Actual answer: implement expensive hedge to make trustees comfortable, destroy value over multi-year period, eventually exit hedge after wasting €6m proving it was unnecessary.

This is common. Credit hedges implemented for political/governance reasons, not economic reasons, in asset classes where hedging is structurally impossible to do effectively.

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 choice is matching hedge tenor to reporting and governance cycles rather than to the timeline over which credit losses are realised.

Hedges are typically shaped by:

- quarterly hedge effectiveness testing, driving 3–5 year maturities
- annual budget cycles, driving rolling 12-month hedges
- rating agency frameworks, driving index-based coverage
- committee reporting, driving visible “current” protection

Credit losses unfold differently:

- public corporate default to final recovery: ~18–36 months
- private credit impairment to resolution: ~24–48 months
- infrastructure distress to recovery: ~36–60 months

The timelines rarely align.

Five-year protection is often rolled before outcomes are known, crystallising roll costs and basis risk at precisely the wrong time.

A more robust approach is matching hedge maturity to expected loss resolution:

- public IG corporates: longer-dated protection (7–10 years) where available
- private credit: accept that CDS is structurally mismatched
- project finance: hedge through construction and ramp-up, then let exposure run

- distressed or HY: short-dated protection 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?

Not "hedge credit risk," but a concrete failure state.

Example 1: Leveraged fund with redemption risk
Failure state: spread widening triggers NAV decline → redemptions → forced selling → permanent capital loss.
Design objective: liquidity, not default protection.
Approach: short-dated index CDS sized to offset spread shocks that trigger redemptions.

Example 2: Insurance general account
Failure state: rating migration drives regulatory capital pressure → forced deleveraging.
Design objective: regulatory capital stability.
Approach: CDS structured to meet capital offset rules, sized to formulas rather than economics.

Example 3: Family office with concentration risk
Failure state: single large position defaults → material capital and reputational loss.
Design objective: protect against idiosyncratic loss.
Approach: single-name CDS on the concentrated position only, exited if exposure is reduced.

In each case, the hedge is designed around the failure mechanism, not around a generic notion of "credit risk."

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
- concentrated counterparty exposure

These attributes improve back-tests and presentations. They degrade under stress as correlation breaks, liquidity disappears, and operational constraints bind.

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 include:

Liquidity

- what were bid-offer spreads in prior stress periods?
- can gains be monetised, not just marked?
- what is settlement timing versus liquidity needs?

Counterparty

- how exposed are we if the dealer deteriorates?
- what happens if a counterparty fails?

Correlation

- how does the hedge behave if correlation collapses?
- can objectives still be met with significant basis slippage?

Operational

- can the hedge be managed under reduced staffing and time pressure?
- what decisions require committee approval versus delegation?

Governance

- how are large hedge losses or gains explained internally?
- what tolerance exists for outcomes that look imperfect but are intentional?

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 that leads to higher risk-taking
- opportunity cost from premium spend versus holding liquidity
- increased political risk if a hedge fails visibly
- strategic rigidity created by hedge roll commitments

An alternative approach is explicit risk ownership.

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.

This takes courage. It requires educating stakeholders on why not hedging is the right choice. But it's honest, and it prevents worse mistakes

10. Structural Diagnostic – Is Your Credit Hedge Real?

This section is intentionally practical. It is designed to be used in an investment or risk committee, not read passively

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.

Good answer: "If our top 5 BBB-rated positions migrate to BB simultaneously (2008-style rating agency catch-up), we face $45m in additional regulatory capital charges. This hedge offsets $35m of that charge, keeping us inside our capital buffer with room for one more downgrade."

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

Test: Write down the specific failure scenario (names, amounts, timeline, consequences). Show how hedge payoff matches that scenario. If you can't write this down with numbers, you don't have a good hedge.

Question 2: Have You Stress-Tested This Hedge in Actual Historical Crises?

Not back tested returns – stress-tested operational functionality.

Good answer: "We analysed March 2020, September 2008, and August 2011. In all three, this specific index traded with bid-offer spreads of 8-12bp (versus 2-3bp normal). We could have monetized in 2-3 days. We've documented dealer capacity and alternative execution venues."

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

Test: Pull actual transaction data from past crises for your specific hedge instrument. Look at bid-offer spreads, trade sizes, settlement times. If this data doesn't exist or you haven't reviewed it, your hedge might not work when needed.

Question 3: If You Monetized This Hedge Tomorrow, What Percentage Would You Lose to Slippage and Friction?

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." (Translation: you haven't actually analysed execution costs.)

Test: Get actual dealer quotes to unwind the hedge today. Not indicative levels – actual executable bids. If you're shocked by the slippage, imagine 3x worse in a crisis. That's your real cost. Many people don't do this.

Question 4: What Happens If Your Hedge Counterparty's Stock Drops 50%?

Good answer: "We're cleared through LCH, no counterparty exposure." Or: "We have CSA agreements with daily margining and segregated collateral with third-party custodian. Tested during dealer stress in 2020."

Bad answer: "Our dealer is a major bank; they're not going to fail." (Lehman Brothers was a major bank.)

Test: Calculate your mark-to-market exposure to each hedge counterparty under stressed scenarios. Map your collateral agreements. Identify what happens if a dealer fails mid-crisis. If you don't know, you have unhedged counterparty risk inside your "hedge." This is still not treated seriously by many buy-side investors. One of the major trading systems used by 80% of the market was surprised when I asked about how I check counterparty exposure for derivatives using their system. Not something their clients were looking at.

Question 5: Would You Personally Pay This Ongoing Cost If It Was Your Own Money?

Brutal reality check.

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

Bad answer: "It's institutional money, governance requires it." (Translation: you wouldn't do this with your own capital.)

Test: Calculate cumulative hedge costs over 10 years. Compare to expected loss from unhedged portfolio over same period. If hedge costs exceed expected losses by 2x+, and you wouldn't 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 monetized

If you've held a hedge for 3+ years and never taken profits or losses by exiting, you don't actually know if it works. You're running a permanent cost with unproven functionality.

Red Flag 2: Nobody can explain the correlation assumption

"We assume 0.75 correlation between our portfolio and the index." Why 0.75? Where did this come from? If the answer is "it's in the model," nobody understands your hedge. In most market crises, correlations rise sharply and converge.

Red Flag 3: Hedge sizing changed to pass effectiveness tests

You started with 30% notional coverage. Effectiveness test failed at 0.78 (need 0.80+). You increased to 35% notional and retested. You've optimized for accounting, not economics.

Red Flag 4: The hedge costs more than the illiquidity premium you're earning

You're earning 180bp over liquid benchmarks for illiquidity. Your hedge costs 150bp annually. You're paying away 83% of your excess return to hedge a risk that's inherent to the asset class you chose. This is incoherent.

Red Flag 5: You've never calculated total cost of ownership

Premium + roll costs + bid-offer + collateral costs + operational overhead + opportunity cost. If you haven't added all of these and compared to expected portfolio losses, you don't know if hedging is positive NPV.

Red Flag 6: The hedge survived a portfolio review but not independently

"We reviewed our credit portfolio and decided hedging was appropriate." Did you also review: not hedging and redeploying that budget to other risk controls? If the hedge can't justify itself outside the context of "we need to do something about credit risk," it's probably not the right something.

CASE STUDY 5: The Honest Diagnostic That Killed a Hedge Program

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

New CIO commissioned independent review of hedge effectiveness. Hired external consultant (former trader, not risk consultant) to stress-test the program.

Findings:

Question 1 (Specific loss event): FAILED

- Team couldn't articulate specific scenario hedge prevented
- Answer was "regulatory capital volatility" but capital model changes made hedge impact unclear
- No documented analysis of hedge benefit under actual capital rules

Question 2 (Historical stress test): FAILED

- No analysis of March 2020, 2011, or 2008 execution feasibility
- Assumptions based on normal-market liquidity
- 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 (60-80bp on $2.5bn notional)
- Never calculated this before
- Monetisation had never been attempted in 5-year program history

Question 4 (Counterparty risk): PARTIAL PASS

- Cleared through LCH, but legacy bilateral positions still outstanding ($800m)
- Uncollateralized exposure to single dealer: $12m mark-to-market
- No backup dealers documented

Question 5 (Personal money test): FAILED

- Investment team admitted they wouldn't pay these costs personally
- Program existed because "previous CIO implemented it" and stopping felt like admitting mistake

Red flags identified:

- Cumulative hedge costs over 5 years: $142m
- Realized benefits (two credit events where hedge paid out): $8m
- Net cost: $134m
- Expected credit losses over same period (modelled): $95m
- Program had cost 41% more than unhedged expected losses
- Operational overhead (systems, people, dealer management): additional $3m/year

CIO decision:

Terminated entire hedge program. Redeployed capital:

- $20m to enhanced credit surveillance systems
- $15m to dedicated workout team for distressed situations
- Remainder returned to investment capital for loss absorption

First-year results: no hedge, portfolio performed in line with expectations, one credit event (automotive supplier) with 52% recovery. Total credit losses: $45m. Would have been $42m with hedge (after slippage and costs). Saved $28m in annual hedge premium.

Five years later: Cumulative savings vs continuing hedge: $168m. Cumulative credit losses: $203m. Net cost of being unhedged vs hedged: $35m worse outcomes, but $168m saved in premium = $133m better off.

The hedge program had destroyed $134m over five years protecting against losses that would have been smaller than the protection costs.

Key insight: Nobody had ever done this analysis because:

- Original implementation was driven by rating agency pressure
- Continuing it became institutional inertia
- Stopping it required admitting the previous 5 years were mistaken
- New CIO had no legacy attachment and asked the uncomfortable questions

The honest diagnostic killed a program that should never have existed.

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

This is the section most practitioners skip. The assumption is always "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 models, PD/LGD, rating transitions)
- Credit hedges reduce risk (capital offset, effectiveness testing)
- Hedged portfolios are safer than unhedged (regulatory capital relief)

The economic reality:

- Credit models miss tail risk, correlation shifts, and recovery uncertainty
- Credit hedges often transfer one risk (default) while creating others (liquidity, counterparty, basis)
- "Safer" on regulatory metrics often means more fragile in actual stress

The real conversation with regulators, rating agencies, or boards:

"We understand the regulatory framework encourages credit hedging. We've analysed available hedging approaches and concluded they would destroy value for our stakeholders:

- Available hedges cost 140bp annually for protection with 60% expected effectiveness = 230bp cost per unit of actual protection
- Our portfolio earns 200bp above risk-free specifically for taking unhedgeable credit risk (illiquidity premium, complexity premium)

- Hedging would pay away 100%+ of excess return to achieve partial protection
- We have adequate capital to absorb expected credit losses (stress-tested to 2008 levels)
- We believe accepting unhedged credit risk with strong underwriting and monitoring provides better outcomes than expensive ineffective hedging

We are choosing to stay unhedged."

This is harder than implementing a bad hedge. It requires:

- Confidence in your analysis
- Ability to educate stakeholders
- Willingness to accept responsibility for losses
- Integrity to admit when hedging doesn't work

But it's honest and doing so prevents worse mistakes.

11.2 Cost of Honesty vs Cost of Pretence

Pretending to manage credit risk via ineffective hedges has costs:

Direct costs:

- Ongoing hedge premium (obvious)
- Roll costs, bid-offer, slippage (less obvious)
- Operational overhead (systems, people, dealers)
- Opportunity cost of capital (hedge budget could be used elsewhere)

Indirect costs:

- False confidence leading to higher risk-taking
- Complexity creating operational risk
- Governance attention diverted from real risks
- Career/political risk when hedge inevitably fails

Cultural costs:

- Normalizing "compliance over effectiveness"
- Rewarding appearance over substance
- Punishing honesty ("why aren't we hedged like everyone else?")
- Creating organizational debt that compounds

The cost of honesty is immediate and visible: "We're 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 hedge fails.

Most organizations choose pretence because the pain is distributed and deniable. The honest organisations choose transparency and accept the accountability.

11.3 Why Doing Nothing Is Often the Right Hedge

"Doing nothing" sounds reckless. It's often the most sophisticated choice.

When doing nothing is right:

1. You're adequately diversified

50+ positions, <3% in largest single-name, sector limits enforced, industry diversification. Expected credit losses over a decade: 80-120bp. Available hedges cost 120bp+ annually and provide 60-70% effectiveness.

Doing nothing is economically superior. The diversification *is* your hedge.

2. You have structural seniority

Senior secured positions with meaningful subordination. Infrastructure debt with 1.4x+ debt service coverage. Private credit with conservative leverage (3.5x or below).

Your "hedge" is being at the top of the capital structure. Adding CDS on top of this is redundant protection.

3. You can absorb the losses

Long-time horizon, no forced selling, adequate capital reserves, stress-tested to 2008-level scenarios and survived with capital intact.

You don't need to transfer risk you can absorb. The cost of transfer exceeds the benefit.

4. You're 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're earning.

Accept the risk. That's why you're being paid.

5. The hedge creates more problems than it solves

Mark-to-market volatility in otherwise hold-to-maturity portfolio. Counterparty risk. Collateral requirements. Operational complexity. Governance burden.

Sometimes the cure is worse than the disease. Doing nothing is the least-bad option.

Framework for "do nothing" decision:

If you answer YES to 4+ of these, staying unhedged is likely optimal:

1. Portfolio is diversified (>40 positions, <4% largest)
2. You have structural protection (senior secured, meaningful subordination)
3. Hold-to-maturity accounting (no MTM, no forced selling)
4. Adequate capital to absorb 2008-level credit losses
5. Being paid illiquidity/complexity premium >150bp

6. Available hedges cost >100bp annually with <70% effectiveness
7. No regulatory capital requirements driving hedge need
8. Long time horizon (>5 years average hold period)

This is most insurance general accounts, most infrastructure equity, most private credit, and most family office portfolios.

Yet most of these portfolios still hedge. Why? Because saying "we're unhedged" requires explaining why, and "we're hedged" requires no explanation.

11.4 Reframing Unhedged as a Positive Choice

Language matters. "We're unhedged" sounds like neglect. "We've chosen not to hedge" sounds like incompetence. Better framing:

"We manage credit risk through portfolio construction rather than derivatives."

What this means:

- Diversification requirements (minimum 50 positions)
- Conservative underwriting (maximum leverage, minimum coverage ratios)
- Active monitoring (quarterly reviews, early warning systems)
- Structural protection (senior secured, meaningful equity cushions)
- Adequate capital reserves (stress-tested to historical worst-case)

This is active risk management. It's not hedging, but it's not passive acceptance either.

"We've analysed credit hedging and concluded the costs exceed the benefits for our portfolio."

What this means:

- We've done the work
- We understand the trade-offs
- We're making an informed choice
- We're accountable for outcomes

This is professional judgment, not negligence.

"Our returns include compensation for unhedgeable credit risk, and hedging would destroy value."

What this means:

- We understand our return sources
- We're honest about what can and cannot be hedged
- We're optimizing for stakeholder outcomes, not optics
- We're not paying away returns for ineffective protection

This is fiduciary responsibility.

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. You don't have diversification. A single blow-up threatens solvency. Hedge the concentration or exit the position.

Leverage amplifies credit risk: Portfolio is funded with short-term debt or faces margin calls. Credit deterioration triggers forced selling. You've created liquidity risk on top of credit risk. Either deleverage or hedge.

You're being paid for market beta, not credit alpha: Your credit spread over benchmarks is <100bp. You're taking credit risk but not being compensated for unhedgeable exposure. You're mis-priced. Either demand higher spreads or hedge.

Regulatory/rating agency requirements are binding constraints: You need regulatory capital relief or rating agency capital credit. These requirements might be economically questionable, but they're real constraints. Sometimes you hedge to satisfy external requirements even when economics don't justify it.

Stakeholders cannot tolerate the outcomes: Insurance policyholders, pension beneficiaries, or investors who would withdraw capital after credit losses. Political tolerance for losses is below economic capacity. You might need to hedge for stakeholder management even when economics don't require it.

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's treated as a separate problem from portfolio construction. The two are inseparable. This section connects the diagnostic back to how you build portfolios in the first place.

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)
- Enhancing monitoring and systems (prevent losses vs transfer losses)

Example comparison:

Option A: Hedge with CDS

- $500m BBB portfolio

- 30% notional CDS hedge ($150m)
- Annual cost: 95bp = $1.425m/year
- Expected effectiveness: 65%
- True cost per unit of protection: 146bp
- 10-year cumulative cost: $14.25m

Option B: Build quality and liquidity

- $475m portfolio (smaller due to reallocating hedge budget)
- Upgrade $75m from BBB to A rated (25bp yield give-up = $187k/year)
- Hold $25m in cash/T-bills (200bp yield give-up = $500k/year)
- Invest $200k/year in credit surveillance and monitoring
- Remaining budget: $538k/year to loss reserves

10-year result:

- Higher average credit quality (less default probability)
- 5% liquidity buffer (actual, not theoretical)
- Better monitoring (earlier warning, better recoveries)
- $5.4m in additional loss reserves
- No counterparty risk, no basis risk, no operational complexity

Which portfolio is actually safer?

Option A looks "hedged." Option B is structurally more resilient. Most governance frameworks would choose Option A because it's easier to explain.

The capital allocation lens forces clarity:

"We have $1.5m annually to allocate toward credit risk management. Should we: a) Buy CDS with questionable effectiveness, or b) Improve portfolio quality, add liquidity, enhance monitoring, and build reserves?"

When framed as resource allocation rather than "should we hedge?", the answer often changes.

12.2 Hedging Credit Exposure vs Hedging Credit Outcomes

Critical distinction that most frameworks miss:

Credit exposure = notional size of credit risk in portfolio **Credit outcome** = specific bad event you're trying to prevent

Most hedges target exposure. Better hedges target outcomes.

Example: Pension fund with liability-driven investing

Exposure-based hedging:

- $3bn credit portfolio supporting long-dated liabilities
- Hedge 40% of credit exposure with CDS

- Cost: $12m annually
- Protects against: spread widening, portfolio MTM volatility

Outcome-based hedging:

- Real risk: inability to meet benefit payments in 15-20 years due to credit defaults
- Relevant horizon: match to liability duration (20+ years)
- Hedge design: Don't hedge MTM volatility (irrelevant for 20-year horizon). Instead:
    - Upgrade credit quality to A-rated (lower default probability)
    - Extend duration to match liabilities (reduce reinvestment risk)
    - Accept spread volatility (not economic risk for hold-to-maturity)
    - Build dedicated reserve for expected defaults over 20 years

Cost: $4m annually (upgrade yield give-up + reserve contributions) Outcome: Better aligned to actual risk (meeting liabilities), 66% cheaper, no derivative complexity

The outcome lens asks: "What specific bad future are we preventing?"

Not "reduce credit VAR" or "limit spread duration" – those are exposure metrics.

Real outcomes:

- "Unable to meet pension obligations in 2043"
- "Forced to sell assets at 40% loss during redemption spike"
- "Regulatory capital falls below minimum threshold"
- "Lose largest institutional client after 15% drawdown"

Design protection around outcomes, not exposures. Outcomes are specific, measurable, and reveal whether hedging is actually necessary.

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

Decision tree for credit risk management:

STEP 1: Can you diversify?

YES → Diversify first. Cheaper than hedging, no basis risk, no counterparty risk, no operational complexity. Build to minimum 40-50 positions, <3% largest name.

NO (concentrated exposure, illiquid market, specific mandate) → Go to Step 2.

STEP 2: Can you upgrade quality or gain structural seniority?

YES → Move up the capital structure. Senior secured instead of unsecured. A-rated instead of BBB. Lower leverage multiples. Structural protection beats derivative protection.

NO (stuck with current positions, quality already maximized) → Go to Step 3.

STEP 3: Can you absorb the loss?

YES (adequate capital, long horizon, no forced selling) → Accept unhedged. Monitor closely. Build reserves. Enhance surveillance. Save the hedge premium.

NO (insufficient capital, leverage, redemption risk, regulatory constraints) → Go to Step 4.

STEP 4: Can you get effective hedging at reasonable cost?

YES (liquid markets, good correlation, acceptable costs, proven stress liquidity) → Hedge selectively. Use simplest structure. Stress-test extensively.

NO (illiquid, poor correlation, excessive costs, questionable effectiveness) → Go to Step 5.

STEP 5: You have a problem

You can't diversify, can't upgrade, can't absorb losses, and can't hedge effectively. Your options:

- **Reduce position size** (accept lower returns, reduce risk)
- **Exit entirely** (admit this risk doesn't fit your portfolio)
- **Change constraints** (raise more capital, reduce leverage, extend horizon)
- **Accept that you're gambling** (hedge box-ticking for governance, pray it works)

Most portfolios jump straight to Step 4 ("let's hedge") without working through Steps 1-3. 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 is optimizing for IFRS 9 compliance, not economic protection.

Better definition:

"A hedge is effective if it meaningfully improves our probability of achieving portfolio objectives when credit stress occurs."

This requires defining:

1. Portfolio objective Not "maximize risk-adjusted returns" – too vague.

Specific: "Generate 6.5% net returns over 10 years with maximum 15% drawdown in any rolling 3-year period."

2. Credit stress scenario that threatens objective

"If more than 3 defaults occur within 24 months with average recovery <50%, we will breach the 15% drawdown threshold and potentially trigger institutional client redemptions."

3. How hedge changes probability of achieving objective

Without hedge:

- Probability of >3 defaults in 24 months: 8% (modelled)
- Expected drawdown if scenario occurs: 18%
- Probability of breaching drawdown threshold: 8%

With hedge:

- Cost: 120bp annually (reduces base returns to 5.3%)
- Expected hedge payout if scenario occurs: 60% of losses
- Expected drawdown with hedge: 12%
- Probability of breaching threshold: 3%

Cost-benefit:

- Reduced drawdown risk: 8% → 3% (5 percentage points)
- Cost: 120bp annually = 12% of total return over 10 years
- Trade-off: Give up 1.2% in returns to reduce drawdown probability by 5 percentage points

Decision: Is 1.2% of returns worth 5pp lower drawdown probability?

This is an answerable question. It forces explicit trade-offs. It reveals whether hedging is worth the cost.

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.

Reframe effectiveness around:

- Probability of achieving objectives (with vs without hedge)
- Cost per unit of risk reduction (bp spent per % reduction in bad outcomes)
- Trade-offs made explicit (return give-up vs drawdown protection)
- Stress-tested under actual historical scenarios (not theoretical models)

This is how you decide if a hedge is worth having. Not correlation statistics – outcome probabilities and explicit costs.

13. Conclusion: The Credit Hedge Illusion

Credit hedging persists not because it works, but because it satisfies governance requirements, provides career cover, and creates the appearance of risk management. For most institutional portfolios – particularly private credit, infrastructure, insurance general accounts, and long-duration multi-asset strategies – credit hedging as commonly practised is expensive defensive activity that transfers accountability while leaving economic risk substantially intact.

The core problems are structural, not solvable:

- CDS hedges spread volatility, not default losses
- Correlation assumptions break precisely when you need them
- Liquidity evaporates under stress when hedge monetization matters most
- Roll costs and basis risk compound over multi-year periods
- Counterparty risk and operational complexity create new vulnerabilities
- Misalignment between hedge horizons and loss realization timelines
- Fundamental impossibility of hedging illiquid exposures with liquid derivatives

The alternatives:

- Portfolio construction that avoids risk (diversification, structural seniority, conservative underwriting)
- Accepting unhedged exposure with adequate capital and monitoring
- Insurance-style protection for specific binary risks
- Improving credit surveillance and workout capabilities
- Building loss reserves instead of paying ongoing hedge premiums

The path forward requires:

- Clarity about what hedges actually achieve versus what they cost
- Willingness to challenge institutional inertia ("we've always hedged")
- Courage to stay unhedged when economics support it
- Reframing credit risk management as portfolio construction, not derivative overlays
- Measuring hedge effectiveness by outcomes achieved, not accounting correlation
- Accepting that sometimes the right hedge is no hedge

This paper won't be popular. It challenges conventional wisdom, questions standard practice, and suggests that billions spent annually on credit hedging might be wasted. But after 30 years structuring credit protection across insurance companies, pension funds, asset managers, and family offices in Asia-Pacific, I've watched the same patterns repeat: sophisticated investors paying for protection that doesn't work, suffering losses the hedge was meant to prevent, and continuing the programs because stopping would require admitting the previous decade was mistaken.

The credit hedge illusion persists because questioning it is uncomfortable. The institutions that break this pattern – that honestly assess whether their hedges transfer economic risk or just satisfy governance optics – will make better decisions and achieve better outcomes.

What this means for your portfolio:

If you currently hedge credit risk, work through the diagnostic in Section 10. If you can't answer "yes" to four of the five questions, your hedge is probably incorrect. Consider the alternatives in Sections 7 and 11.

If you're considering credit hedging, work through the portfolio construction checklist in an associated paper ***Before You Hedge Credit: A 10-Question*** Diagnostic (contact me if you would like a copy)**.**

If you can't check eight of the boxes, you're not ready to hedge – you need to fix portfolio construction first.

If you're being pressured to hedge by governance, regulators, or consultants, use the frameworks in Section 11 to have a real conversation about costs, benefits, and alternatives.

The goal isn't to never hedge credit risk. It's to hedge only when it transfers genuine economic risk at justifiable cost, and to have the integrity to admit when it doesn't.

14. Further Reading and Practitioner Resources

Additional practitioner papers and CIO Briefs published by Para Bellum Advisors are available at:

www.parabellumadvisors.com/insights.

These materials address rates, FX, credit, collateral, and balance-sheet risk management from a structural and implementation perspective.

15. About Para Bellum Advisors

Para Bellum Advisors is an independent advisory firm specialising in derivatives, collateral, and balance-sheet efficiency for institutional investors.

The firm works with lean investment teams managing complex, long-dated portfolios across FX, rates, credit, equity, and volatility risk. Its focus is not product distribution or transaction volume, but structure: how hedges are designed, how capital is consumed, and how portfolios behave under stress.

Para Bellum Advisors' work is grounded in practitioner experience across trading, structuring, and portfolio management within banks, asset managers, and insurance balance sheets. The objective is not theoretical optimisation, but durable improvement in capital efficiency, liquidity resilience, and realised outcomes.

Further information is available at **www.offers.parabellumadvisors.com**

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

Para Bellum Advisors – Disclaimer

This paper is provided for informational purposes only. It does not constitute investment advice, financial product advice, or a recommendation to transact. It is not tailored to any institution's objectives, financial position, risk appetite, or regulatory constraints.

All examples are illustrative. Markets move, assumptions change, and outcomes will differ. Past performance is not a guide to future results. Any views expressed reflect Para Bellum Advisors' judgement at the time of writing and may change without notice.

Institutions should obtain independent advice and conduct their own analysis before making any investment, hedging, or risk-management decision.