Practitioner Paper

Designing Long - Dated Rates Hedges That Actually Work

Why Familiar Structures Fail and What Endures Across Regimes

Version 2, December 2025

By Mike Duncan, Para Bellum Advisors

www.parabellumadvisors.com

Executive Summary

Most long-dated rates hedges fail even when interest rates behave broadly as expected.

They fail because they are designed to look correct at inception rather than to survive regime change, liquidity stress, and governance intervention over time. Duration is neutralised, reports are clean, and execution is competent – yet outcomes deteriorate as cashflows persist, margin demands grow, and flexibility disappears.

The failure is structural, not predictive.

This paper examines long-dated rates hedging through the lens of behaviour rather than pricing. It shows that the decisive distinction is whether a hedge closes economics at inception or defers them into future regimes. Instruments that defer economics accumulate exposure to funding conditions, liquidity availability, and repeated decision-making under stress. Instruments that close economics concentrate risk early but reduce dependence on future intervention.

Common practices such as tenor matching and rolling structures are not flawed in themselves. They become dangerous when applied to exposures that are not contractually durable, or when flexibility is assumed rather than real. In these cases, hedges introduce termination risk, persistent cash outflows, and forced action at precisely the wrong point in the cycle.

Collateral and liquidity are central to this dynamic. Modern rates hedging is capital-intensive by design. Margin is pro-cyclical, funding is not guaranteed, and liquidity stress – not mark-to-market volatility – is what ultimately breaks hedges in practice.

Emerging markets provide a clear stress test. They do not change the logic of rates hedging, but they remove the buffers that allow poor structure to persist. Where liquidity is fragile and instrument choice is constrained, the consequences of deferring

economics appear faster and more violently. The disciplines enforced in emerging markets are the same ones that long-dated hedging requires everywhere.

The paper concludes with a coherent framework and a structural self-assessment designed to help institutions identify latent fragility before it becomes binding. It does not prescribe trades or offer optimisation tactics. Its purpose is to reset how long-dated rates hedging is understood, designed, and governed.

Long-dated rates hedging behaves only when uncertainty is removed rather than postponed, when liquidity and governance are treated as design inputs, and when structures are built to survive the environments that inevitably arrive.

Table of Contents

1. Introduction – Why Long-Dated Rates Hedges Fail

Long-dated rates hedging is usually treated as a technical problem.

Measure duration. Neutralise DV01. Pick a liquid instrument. Rebalance when conditions change.

That approach works tolerably well for short horizons. It breaks down once exposures persist across years, regimes, and governance cycles.

Most long-dated hedging failures are not the result of being wrong about interest rates. They occur even when rates behave broadly as expected. They occur even when hedges "work" according to standard risk metrics. They occur even when execution is competent and policies are followed.

The failure is structural.

At its core, every long-dated rates hedge makes a choice: close the economics of the exposure at inception or defer them into future regimes. Instruments that close economics fix the outcome when the trade goes on. Instruments that defer economics leave material risks open–cashflows that reset, funding that compounds, decisions that must be revisited. That distinction determines whether a hedge behaves or fails once conditions change.

Between 2020 and 2023, many institutional investors learned this the hard way. Receive-fixed swaps locked in historically low rates, exactly as intended. Duration risk was neutralised. Reports looked clean. Yet cashflows turned punitive as floating legs reset higher. Liquidity stress replaced interest-rate risk as the dominant problem.

Nothing unexpected happened. The hedges did what they were designed to do.

What failed was the assumption that locking a rate was the same thing as locking the economics.

This paper starts from a simple observation: long-dated rates hedging is not about suppressing mark-to-market volatility. It is about deciding which risks are eliminated today and which are left open for future regimes, future liquidity conditions, and future governance to resolve.

Most portfolios answer that question accidentally.

They use familiar instruments. They optimise entry pricing. They prioritise flexibility. Without stating it explicitly, they defer key economic risks into the future.

Those deferred risks are not theoretical. They show up as persistent cash outflows, margin stress, forced unwinds, and governance interventions at precisely the wrong point in the cycle.

A hedge that looks neutral on a risk report can still behave badly over time.

In this paper, a hedge *behaves* only if it delivers the intended economic outcome over its full horizon, survives regime change without forcing action, and does not create liquidity stress that overwhelms the original risk.

That definition deliberately excludes optics. It excludes short-term comfort. It excludes the assumption that future decisions will be made rationally under pressure.

This paper rests on a single distinction: whether a hedge closes economics or defers them. Everything that follows – instrument choice, tenor matching, rolling structures, collateral behaviour, governance stress, and emerging-market failure – is an application of that distinction.

The paper does not attempt to predict interest rates. It does not recommend a universal "best" instrument. It does not focus on pricing optimisation or execution tactics.

Instead, it examines why familiar approaches to long-dated rates hedging fail so consistently, how different instruments actually behave once time and regime change are introduced, and what structural alignment looks like when the objective is endurance rather than convenience.

The aim is not complexity. It is honesty.

Long-dated rates hedging fails not because markets are unpredictable, but because too many hedges are designed to look sensible today rather than survive the environments that inevitably arrive later.

2. The Nature of Long-Dated Rate Exposure

Long-dated rate exposure is often described in terms of duration. That is convenient, measurable, and largely insufficient.

Duration tells you how much a portfolio moves when rates change today. It says very little about how long that exposure persists, how difficult it is to reverse, or what happens when conditions change over time.

For short-term positioning, that distinction barely matters. For long-dated exposure, it is everything.

What makes long-dated rate exposure dangerous is not volatility. It is persistence.

Many institutional exposures do not mature neatly or unwind easily. They are embedded in balance sheets, contractual cashflows, leverage structures, and funding arrangements that extend well beyond any single investment decision. They survive strategy reviews, committee changes, and market cycles.

In practice, long-dated rate exposure often has no clear end point. It endures.

This creates a fundamental mismatch with the way hedging decisions are typically framed. Risk is assessed at a point in time. Reports focus on sensitivity. Hedges are judged on how neutral the portfolio looks today. Little attention is paid to how long the exposure will exist, or how many regimes the hedge will have to survive.

That gap matters because long-dated exposure fails in a different way.

Portfolios rarely blow up because mark-to-market volatility becomes uncomfortable. They fail because cashflows become unsustainable, liquidity is consumed at the wrong moment, or governance is forced to intervene when options are worst.

A hedge can look perfectly neutral on a risk report and still drain cash year after year. It can reduce DV01 and still amplify funding stress. It can meet every policy requirement and still become a structural liability.

The key distinction is between cashflow certainty and mark-to-market neutrality.

Mark-to-market volatility is an accounting outcome. It can be tolerated if no action is required. Cashflow stress is an economic reality. Once it dominates, something must be done.

Long-dated exposures are vulnerable precisely because they span multiple environments. Interest-rate regimes change. Liquidity conditions tighten and loosen. Funding costs move. Margin models shift and the people responsible for the hedge change as well.

Time magnifies small design errors.

A decision that looks reasonable at inception can become fragile years later, not because the original logic was wrong, but because the hedge was never designed to survive routine change. Assumptions that were implicit become binding. Deferred risks surface. Flexibility disappears.

This is why applying short-dated hedging logic to long-dated exposure does not scale. The longer the horizon, the more the outcome depends on what was left open rather than what was fixed.

Long-dated rates hedging is therefore not primarily a problem of precision. It is a problem of commitment.

The central question is not how much the portfolio moves when rates change, but how long the exposure persists, how credible the intent to hold really is, and how difficult it would be to change course once conditions deteriorate.

Ignore those questions, and duration metrics provide a false sense of control.

By the time the weakness becomes visible, the problem is no longer rates. It is liquidity, governance, and forced decision-making under stress.

3. Why Long-Dated Rates Hedges Fail

Most long-dated rates hedging failures follow a familiar pattern. They do not begin with a bad market call, a policy breach, or poor execution. They begin with a design choice that looks reasonable in isolation and fragile over time.

A hedge is approved for sound reasons. Duration risk is neutralised. The instrument is liquid and familiar. Entry pricing looks acceptable. Governance is comfortable.

For a time, everything works.

Then conditions change – not catastrophically, but enough for structure to matter. Cashflows drift. Margin demands rise. Funding becomes visible. What had been secondary becomes dominant.

Eventually, the hedge itself becomes the problem.

This pattern is consistent across institutions and cycles. It reflects a small number of structural mistakes that are easy to overlook at inception and difficult to correct later.

The anatomy of failure

Most long-dated hedging failures share a common structure. They result from a small number of design choices that look reasonable at inception and become fragile over time.

Short-dated logic applied to long-dated exposure

Hedges are often designed using tools that work for tactical positioning:

- duration matching
- DV01 neutrality
- frequent rebalancing
- reliance on liquid, standardised instruments

These approaches manage short-term sensitivity. They do not manage long-term behaviour.

Once exposure persists across years, the hedge accumulates exposure to funding conditions, liquidity regimes, and repeated decision points. The risk profile changes even if the hedge does not.

Optionality mistaken for flexibility

Short tenors and rollable structures are often justified as prudent flexibility. In practice, they transfer risk from markets to future governance.

Decisions that could have been settled at inception are deferred. Those decisions rarely become easier with time. They become more constrained by liquidity, optics, and stress.

Regime change treated as exceptional

Many long-dated hedges implicitly assume that:

- liquidity will remain available
- funding will remain manageable
- volatility will be episodic

These assumptions are rarely explicit. For long-dated exposure, they are rarely true.

Regime change is not a tail risk. It is the base case.

Collateral and funding treated as secondary

Collateral dynamics are often addressed after the trade:

- margin behaviour is under-modelled
- funding costs are incompletely captured
- liquidity buffers are assumed rather than tested

This works until collateral becomes binding. By then, structure is fixed and choices are poor.

Unpriced termination risk

Hedges are frequently longer than the exposure's enforceable life. When exposure changes, the hedge does not adjust quietly. It must be actively unwound.

That termination risk is rarely priced at inception.

The result

When these elements combine, hedges fail in predictable ways:

- persistent cash outflows dominate returns
- margin calls accelerate during stress
- liquidity buffers are consumed at the wrong time
- governance intervenes when options are worst

At no point does the hedge "stop working". It does exactly what it was designed to do.

The failure lies in what the design deferred.

Most long-dated hedging failures are not forecasting errors. They are the consequence of deferring economic and governance decisions into future environments that turn out to be hostile.

When those decisions arrive, the hedge no longer behaves as protection. It behaves as a constraint.

4. Why Familiar Instruments Dominate

If these failure modes are so consistent, the obvious question is why institutions continue to use the same structures.

The answer is not ignorance. It is incentive.

Long-dated hedging decisions are made in environments where familiarity, defensibility, and ease of governance matter as much as structural fitness. Instruments are chosen in committees, reviewed by auditors, benchmarked by consultants, and inherited by successors.

Familiar instruments win in that setting.

The familiarity premium

Standard instruments offer immediate advantages unrelated to long-term behaviour:

- visible liquidity
- established accounting and regulatory treatment
- compatibility with existing systems
- peer comparability

These attributes make hedges easier to approve and easier to defend. They do not make them more robust.

When outcomes deteriorate years later, the original decision often still looks reasonable in hindsight. Accountability diffuses. That is part of the appeal.

Standardisation as camouflage

Standardisation creates the impression that structures are interchangeable. Hedges that neutralise duration are treated as equivalent even when their cashflow, liquidity, and governance consequences differ materially.

When outcomes diverge, explanations focus on "market conditions" rather than structure.

Standardisation does not eliminate risk. It obscures it.

Optionality as institutional comfort

Short-dated and rollable instruments preserve the appearance of choice. Each roll can be justified. Conditions can be reassessed. Nothing appears locked in.

From an organisational perspective, this feels prudent. Structurally, it transfers risk forward.

Future decision-makers inherit:

- accumulated roll costs
- ongoing exposure to funding conditions
- decisions that must be made under less favourable circumstances

This asymmetry is rational for individuals. It is expensive for institutions.

Liquidity masks fragility

In benign environments, familiar instruments perform well. Liquidity is deep. Funding is cheap. Margin demands are manageable.

This reinforces confidence in the structure. Weakness remains invisible.

When conditions tighten, fragility appears quickly. By then, alternatives are costly and unwinds crystallise losses. The hedge becomes a focal point for intervention.

Why this persists

Institutions will continue to favour:

- standardisation over customisation
- liquidity over economic certainty
- peer validation over structural analysis
- deferral over commitment

These preferences are human and organisational. They are not irrational.

They become dangerous when long-dated exposure is hedged as though it were a short-term problem.

The task of effective rates hedging is not to eliminate these preferences, but to make their consequences explicit – before time and regime change do it instead.

5. Instrument Behaviour Under Regime Change

Long-dated rates hedging fails when instruments are judged on how they price at inception rather than how they behave once conditions change.

Two hedges can have identical duration, identical accounting treatment, and identical entry pricing, yet produce radically different outcomes over time. The difference is not technical. It is structural.

Some instruments close the economics of an exposure when they are put on. Others leave the economics open and defer resolution into the future. That distinction determines whether a hedge behaves or fails once regimes change.

Close economics versus deferring them

An instrument closes economics when the cashflow outcome is fixed at inception. Mark-to-market will move as conditions change, but the underlying economics do not. If the hedge is held to maturity, future regimes are largely irrelevant.

An instrument defers economics when material components of the outcome remain exposed to future conditions. Cashflows reset. Funding costs evolve. Margin requirements change. Decisions must be revisited.

That difference is often blurred in practice because both types of instrument can neutralise duration and look similar on a risk report. Over time, their behaviour diverges sharply.

Instruments that close economics

Structures that close economics do something unfashionable: they replace flexibility with certainty.

Examples include cash-settled bond forwards, forward-starting structures, and genuinely tenor-matched bilateral arrangements where the full economic exchange is fixed up front.

Once executed:

- the financing cost is known
- the carry is embedded
- future rate regimes do not change the cashflow outcome

The hedge may show mark-to-market volatility, sometimes severe. That volatility is uncomfortable but largely irrelevant if the structure is held to settlement alongside the exposure.

The defining feature is not pricing. It is the absence of future decision points.

These instruments force a choice at inception. They demand clarity about commitment and intent. In exchange, they remove uncertainty rather than deferring it.

Instruments that defer economics

Many commonly used rates instruments fall into the opposite category.

Interest rate swaps fix one thing – the fixed rate – and leave everything else open:

- floating-rate resets
- funding conditions
- margin requirements
- liquidity stress

Rolling futures and short-dated instruments go further. They embed repeated repricing, basis exposure, and roll decisions as a permanent feature of the hedge.

These instruments are not defective. They are effective tools for tactical positioning and short-term risk management. Problems arise when they are used to hedge exposures that persist for many years.

In those contexts, the hedge does not eliminate uncertainty. It accumulates it.

The low-rate swap trap

A common real-world example illustrates the point.

During extended low-rate environments, institutions enter long-dated receive-fixed swaps to “lock in” attractive levels. Duration risk is neutralised. Entry pricing looks compelling. Reports show clean hedging outcomes.

When rates reset higher and remain there, the structure reveals its true behaviour:

- floating legs reprice sharply
- net cash outflows become persistent
- margin demands increase during volatility
- funding stress replaces interest-rate risk

The hedge continues to neutralise duration. Economically, it has become a liability.

This is not a forecasting error. It is the predictable result of using an instrument that defers financing exposure in a long-dated setting.

Cleared versus bilateral structures

The distinction between cleared and bilateral instruments becomes critical as horizons extend.

Cleared structures optimise liquidity and standardisation. They impose daily variation margin and conservative initial margin. Liquidity demands increase precisely when volatility rises.

Bilateral structures rely more on credit and negotiated collateral terms. They often exhibit lower margin velocity and fewer forced liquidity events, at the cost of headline pricing and operational simplicity.

Neither is inherently superior. The question is where stress appears when it inevitably arrives.

In long-dated hedging, shifting risk from rates to funding is not always an improvement.

Bond futures and the illusion of substitution

Bond futures are excellent tools for managing short-term duration exposure. They are liquid, cheap, and efficient. They were never designed to provide economic certainty over long horizons.

When futures are rolled repeatedly to maintain a long-dated hedge, the structure quietly changes character:

- roll costs accumulate
- basis risk compounds
- outcomes become path-dependent

What looks like a hedge becomes an active strategy, whether acknowledged or not.

If a hedge requires perpetual rolling to exist, it is not aligned with a long-dated exposure by design.

The core distinction revisited

The decisive difference across rates instruments is not complexity or pricing. It is whether the structure closes the economics of the exposure or leaves them open to future regimes, funding conditions, and governance decisions.

Mark-to-market volatility is survivable if no action is required.
Deferred cashflows, margin stress, and forced decisions are not.

Long-dated rates hedging behaves only when this distinction is made explicitly and respected in the structure.

6. Tenor Matching – Necessary but Dangerous

Tenor matching is one of the most intuitive ideas in rates hedging.

If an exposure lasts ten years, hedge it for ten years.
If cashflows are fixed, match the hedge to the cashflows.

In the right context, this logic is not just sensible – it is optimal. When exposure is contractual, stable, and genuinely long-dated, tenor matching can remove uncertainty almost entirely.

The problem is not tenor matching itself.
The problem is the assumption it relies on.

The assumption no one states explicitly

Tenor matching only works if one condition holds:

- the exposure will still exist, in the same form and size, for the full life of the hedge.

That assumption is rarely tested. When it proves false, the hedge becomes more rigid than the exposure it was meant to protect. More often, it is replaced with softer language:

- “strategic allocation”
- “long-term intent”
- “core exposure”

None of these are contracts. They are preferences.

When a hedge is longer than the exposure’s actual commitment, the hedge becomes the most rigid object in an otherwise flexible portfolio. What was meant to reduce risk quietly introduces termination risk.

When tenor matching works

Tenor matching behaves well when commitment is real and enforceable.

Typical examples include:

- project finance loans held to maturity
- infrastructure debt with contracted cashflows
- fixed-rate bonds matched to liabilities
- asset–liability matched portfolios

In these cases, the hedge does exactly what it should do. It fixes financing economics, removes path dependency, and makes future rate regimes largely irrelevant.

Mark-to-market volatility may exist, sometimes significantly. But if the hedge is held alongside the exposure, that volatility has no economic consequence.

This is tenor matching at its best: boring, precise, and effective.

When tenor matching becomes dangerous

Problems arise when tenor matching is applied to exposures that are not contractually durable.

Common examples include:

- strategic bond portfolios
- credit mandates with turnover
- balance-sheet duration that can be resized
- exposures embedded in multi-asset portfolios

In these cases, the hedge assumes a stability that the portfolio itself does not possess.

When the exposure changes – through rebalancing, asset sales, mandate shifts, or governance decisions – the hedge no longer fits. At that point, the institution faces a set of bad choices:

- unwind the hedge and crystallise mark-to-market
- maintain an over-hedged or misaligned position
- layer additional hedges and increase complexity

None of these outcomes were priced at inception.

Tenor matching as a commitment device

In practice, a long-dated hedge does more than hedge rates. It commits the organisation to holding the exposure.

That commitment is often implicit rather than explicit. Governance may not realise that a decision taken for hedging purposes has materially reduced future flexibility.

When governance tolerance changes – a new CIO, a new board, a shift in strategy – the hedge becomes a focal point for intervention. Not because it is wrong, but because it is rigid.

This is why many long-dated hedges are unwound at precisely the wrong time. The problem is not the hedge itself. It is the mismatch between economic commitment and governance reality.

Mark-to-market is not the risk

Tenor matching is often criticised because of mark-to-market volatility. This misses the real issue.

Mark-to-market volatility is survivable if no action is required.
Forced action is not.

A short-dated hedge expires quietly.
A long-dated hedge must be actively unwound.

The longer the tenor, the larger the consequences when that moment arrives.

A practical rule of thumb

Tenor matching is appropriate when:

- the exposure has a legal or contractual end date
- early exit is unlikely or economically irrelevant
- governance accepts interim volatility without intervention

Tenor matching is dangerous when:

- exposure size or existence is discretionary
- portfolio composition can change materially
- governance comfort depends on short-term optics

In those cases, shorter-dated or layered structures may be economically inferior on paper but structurally safer in practice.

The core insight

Tenor matching does not eliminate risk. It concentrates it in time.

When the hedge matures alongside the exposure, that concentration disappears.
When it does not, it explodes.

Long-dated rates hedging behaves only when tenor matching reflects actual commitment, not aspiration.

7. Rolling Structures – Tactical Tool, Structural Liability

Rolling structures are often presented as a sensible compromise in long-dated rates hedging.

They avoid committing to a long tenor. They preserve flexibility. They are liquid, familiar, and easy to adjust. When uncertainty exists about holding periods or future conditions, rolling can feel like the cautious choice.

For tactical positioning, that instinct is correct.

For long-dated exposure, it is usually wrong.

What rolling actually does

A rolling hedge does not eliminate risk. It reintroduces it repeatedly.

Each roll forces the hedge to be repriced at prevailing market conditions. Carry, curve shape, liquidity, and funding assumptions are reset. What looks like continuity is in fact a sequence of new decisions, each one exposed to the environment of the day.

Over short horizons, this hardly matters. Over long horizons, it dominates outcomes.

Rolling structures therefore behave less like hedges and more like ongoing strategies, whether that is acknowledged or not.

The accumulation problem

The weakness of rolling structures is not visible in any single roll. It emerges through accumulation.

Over time:

- roll costs compound
- basis risk builds quietly
- outcomes become sensitive to timing
- cumulative cashflows diverge from intent

None of these effects are dramatic in isolation. Together, they can overwhelm the original exposure the hedge was meant to control.

This is why rolling structures often underperform expectations without any obvious failure point. The drag is incremental, persistent, and difficult to attribute.

Flexibility is not free

Rolling is often justified as flexibility. That flexibility has a cost.

By choosing a rolling structure, the institution accepts that:

- pricing will be re-set repeatedly
- liquidity must be available at every roll
- market access cannot be assumed
- decisions must be made under future conditions

Those costs are rarely priced explicitly. They are treated as operational details rather than structural risks.

In benign environments, this feels harmless. When liquidity tightens or volatility persists, it becomes binding.

The illusion of control

Rolling structures create an illusion of control because nothing appears locked in. The hedge can always be adjusted, resized, or abandoned.

In reality, this means the hedge is never finished.

Decisions that could have been resolved at inception are deferred indefinitely. Each roll postpones commitment rather than removing uncertainty. Over time, the hedge becomes more dependent on future conditions than the exposure it was meant to protect.

That inversion is rarely intentional.

Where rolling makes sense

Rolling structures are appropriate when flexibility is genuinely required.

Examples include:

- transitional portfolios
- temporary exposures
- discretionary allocations
- tactical overlays

In these contexts, accepting roll risk is the price of optionality.

Problems arise when rolling is used to hedge exposures that are durable, contractual, or structurally embedded. In those cases, the cost of perpetual optionality often exceeds the benefit.

The hard boundary

A simple test clarifies the issue:

If the hedge must be rolled forever to exist, it is not aligned with a long-dated exposure.

Rolling is a valid tactical tool. It is a poor substitute for structural hedging.

Long-dated rates hedging behaves only when repeated repricing is the exception, not the foundation.

8. Collateral, Liquidity, and Regime Stress – Where Structures Break

Most rates hedging discussions focus on price, sensitivity, and hedge ratios. Very few start with collateral and liquidity.

That ordering is backwards.

In modern markets, long-dated rates hedging is capital-intensive by design. Margining, funding, and liquidity are no longer operational details. They are structural risks that determine whether a hedge survives stress or becomes the trigger for forced action.

When hedges fail in practice, it is rarely because interest rates moved “too far”. It is because the hedge could not be funded through the move.

Rates hedging is now capital-intensive

Post-crisis market structure has fundamentally changed the economics of derivatives.

Clearing mandates, conservative margin models, and daily variation margin have shifted risk from counterparty credit to liquidity. This has improved systemic resilience, but it has made individual hedges far more demanding to run.

For long-dated hedges, this matters enormously.

A hedge that looks cheap on entry pricing can consume large amounts of liquidity over time. Initial margin immobilises capital from day one. Variation margin introduces pro-cyclical cashflows that peak precisely when volatility rises and funding is most constrained.

None of this shows up in duration metrics.

Pro-cyclical margin is the failure mechanism

Margin is not neutral over time.

When volatility rises, margin requirements increase. When rates move sharply, variation margin flows accelerate. When liquidity tightens, funding those flows becomes more expensive or more difficult.

These dynamics interact. They do not cancel out.

A long-dated hedge that defers economics into floating cashflows or repeated repricing is therefore implicitly a bet on future liquidity conditions. That bet is rarely stated, let alone approved.

In benign environments, margin feels manageable. In stressed environments, it becomes the dominant driver of outcomes.

The myth of the “cheap” hedge

Hedges are often judged on headline pricing: spreads, carry, and initial execution cost.

This framing misses the largest cost component in long-dated structures: lifecycle liquidity.

A hedge that defers economics may appear cheaper at inception because it leaves funding and carry exposed. Over time, that exposure compounds. Cash outflows persist. Liquidity buffers are consumed. Optionality disappears.

What looked like a cheap hedge turns out to be expensive to live with.

Uncleared structures and collateral velocity

The distinction between cleared and bilateral structures is often reduced to regulatory or operational preference. In long-dated hedging, it is a structural choice.

Cleared structures prioritise standardisation and counterparty safety. They impose daily margining and high collateral velocity. Liquidity must be available continuously.

Bilateral structures rely more on credit support arrangements and negotiated terms. Margin flows are often less frequent and less volatile, but credit and concentration risks increase.

Neither structure is universally superior. The question is where stress appears first, and whether the organisation is equipped to absorb it.

Long-dated hedging fails when collateral velocity overwhelms liquidity tolerance.

Liquidity stress beats mark-to-market volatility

Mark-to-market volatility is uncomfortable. Liquidity stress is decisive.

A hedge can survive extreme mark-to-market swings if no action is required. It cannot survive sustained cash outflows if funding capacity is finite.

This is why hedges are often unwound not at the point of maximum loss, but at the point of maximum liquidity strain. By then, the decision is no longer economic. It is operational.

Once that line is crossed, the hedge stops behaving as protection and starts dictating portfolio actions.

Collateral is still underpriced in decisions

Despite these realities, collateral and funding remain underweighted in hedge design.

They are often treated as:

- operational issues to be managed later
- treasury problems, not investment decisions
- temporary inconveniences

In long-dated hedging, this is a mistake.

Collateral is not a side effect of the hedge. It is part of the hedge.

Ignoring it does not eliminate the risk. It simply delays its recognition until conditions are least forgiving.

The connecting thread

Across instruments, tenors, and structures, the same pattern appears.

Hedges that defer economics accumulate exposure to liquidity, funding, and governance stress. Hedges that close economics concentrate risk at inception and reduce dependence on future conditions.

Collateral is where that difference becomes visible.

9. Emerging Markets as the Stress Test

Emerging markets do not require a different framework for rates hedging. They remove the buffers that allow poor structure to persist.

The disciplines that remain optional in developed markets become binding constraints in emerging contexts. This distinction does not render emerging markets exceptional; rather, it serves to provide greater clarity.

In developed markets, deep liquidity, broad instrument availability, and resilient funding markets can mask fragility for years. Problems emerge slowly, often attributed to bad luck or temporary dislocation.

Emerging markets offer no such protection.

In many emerging rates markets:

- long-dated swaps are illiquid or unavailable
- futures markets are thin or discontinuous
- clearing is limited or punitive
- basis risk is structural, not episodic

Instrument choice becomes unavoidable. Mismatches cannot be papered over.

Rolling structures that function tolerably in developed markets fail quickly in emerging contexts. Roll liquidity cannot be assumed. Pricing gaps widen sharply under stress. Market access can disappear altogether.

Collateral dynamics are harsher as well. Margin is often denominated in hard currency. Funding sources are constrained. FX and rates stress frequently coincide.

As a result, hedging errors migrate rapidly into liquidity crises.

In emerging markets, a poorly designed rates hedge does not merely underperform. It can force asset sales, trigger capital support, or destabilise the entire portfolio.

Ironically, this fragility enforces better behaviour.

Institutions that operate successfully in emerging rates markets tend to:

- match hedges closely to contractual exposures
- avoid perpetual rolling structures
- accept mark-to-market volatility in exchange for economic certainty
- be explicit about which risks are closed and which are not

These are precisely the disciplines that remain optional in developed markets.

Emerging markets do not change the logic of rates hedging. They enforce it sooner.

This is why practitioners who succeed in emerging rates markets often bring superior hedging discipline to developed markets. They have already learned what cannot be deferred.

10. What Right Looks Like – Worked Examples

Up to this point, the paper has focused on why long-dated rates hedges fail and how familiar structures create predictable fragility. It is equally important to show that alternatives exist – and that they are not exotic or clever.

What distinguishes robust hedging is not sophistication. It is alignment.

The following examples are deliberately high-level. They focus on exposure reality, structural intent, and behavioural outcomes rather than instruments or execution details. In each case, the objective was the same: remove uncertainty rather than defer it.

Private credit – fixing what is actually held

The exposure in this case was a portfolio of senior secured loans with long contractual lives and limited secondary liquidity. While turnover existed at the margin, the economic reality was that the portfolio would be held through multiple cycles.

The temptation was to hedge using short-dated or rollable instruments to preserve flexibility and minimise headline cost. That approach was initially favoured internally because it kept the hedge "easy to adjust".

Instead, the hedge was designed around what could not change. The contractual nature of the loans, the absence of realistic exit optionality, and the portfolio's tolerance for interim mark-to-market volatility were treated as primary inputs.

The resulting structure accepted discomfort in reported volatility in exchange for economic certainty. Financing costs were fixed. Liquidity demands were predictable.

The main resistance did not come from markets. It came from internal discomfort with volatility that was economically irrelevant but visually confronting.

The hedge behaved because it reflected reality rather than preference.

Infrastructure – matching economics, not optics

Infrastructure exposures often look deceptively simple from a duration perspective. Cashflows are long-dated, inflation-linked, and relatively stable. The instinct is to hedge with liquid instruments that track benchmarks cleanly.

The challenge is that infrastructure assets are rarely as flexible as the portfolios that hold them. Exit options are limited. Refinancing risk is real. Governance tolerance for restructuring under stress is low.

In this case, the hedging approach prioritised certainty over elegance. Rather than optimising for short-term reporting or peer comparison, the structure focused on aligning hedge maturity and economics with the asset's financing profile.

The hedge looked uncompetitive in some environments and uncomfortable in reports, which created internal pressure to 'revisit' it more than once. But it removed the risk of forced action under stress, which was the true objective.

Those questions never coincided with liquidity stress – which was the point.

The hedge behaved because it was designed for the asset, not the optics.

Family office – separating structure from discretion

Family office balance sheets often combine long-term capital with discretionary flexibility. This creates a unique challenge for rates hedging: parts of the exposure are durable, others are optional.

The common mistake is to hedge the entire balance sheet using a single, short-dated, rollable structure to preserve perceived flexibility. Over time, this blurs the distinction between what must be protected and what can be adjusted.

In this case, the solution was not a single hedge, but a separation of intent.

Structural exposures that were unlikely to change were hedged in a way that closed economics and reduced reliance on future liquidity. Discretionary exposures were left more flexible, accepting that their risk profile could change with circumstances.

This required accepting that not all rate risk would be treated uniformly, which initially created discomfort and debate.

That debate disappeared once stress arrived and decisions did not have to be made under pressure.

The hedge behaved because structure and discretion were no longer confused.

The common thread

Across these examples, the details differ, but the logic is consistent.

In each case:

- exposure reality was prioritised over convenience
- economic commitment was made explicit
- flexibility was used sparingly and deliberately
- liquidity and governance constraints were treated as design inputs

None of the outcomes depended on forecasting rates correctly. They depended on acknowledging what could not change and designing the hedge around that fact.

What "right" looks like in long-dated rates hedging is not perfection. It is coherence.

A hedge behaves when it removes uncertainty instead of shifting it forward, and when it remains aligned with the exposure it protects long after the conditions at inception have faded.

11. A Coherent Framework for Long-Dated Rates Hedging

By this point, the problem should be clear. Long-dated rates hedging fails when structures are chosen for convenience rather than endurance, and when economic decisions are deferred into future environments that are unlikely to be forgiving.

In practice, most institutions only realise which economics they failed to close after the first forced unwind. By then, the framework is obvious – and useless.

What follows is not a checklist or a recipe. It is a way of organising decisions so that hedging choices remain coherent as time, regimes, and governance change.

Two questions that matter

The framework begins with two questions. They must be answered before instruments, pricing, or execution are discussed:

1. **Which economics must be closed today?**
2. **Which risks can legitimately remain flexible without threatening survival?**

Everything else follows from how these questions are answered.

Most failures occur because these questions are either not asked or answered inconsistently.

Closing everything is neither necessary nor desirable. Leaving everything open is rarely survivable. The discipline lies in making the boundary explicit before markets or governance force it.

Close economics where you can

When exposure is contractual, durable, and unlikely to change, deferring economics adds no value.

In these cases, the role of the hedge is not to preserve optionality. It is to eliminate uncertainty. That means accepting structures that concentrate risk at inception in exchange for stability over time.

This often looks conservative. It can look uncomfortable. It frequently looks sub-optimal on short-term metrics. That discomfort is the cost of not having to make decisions later under worse conditions.

Closing economics early reduces dependence on future liquidity, future funding conditions, and future decision-makers. It replaces path dependency with certainty.

Use flexibility only where it is real

Flexibility is valuable when it reflects genuine discretion.

Where exposures can be resized, exited, or reallocated without forcing losses elsewhere in the portfolio, maintaining optionality makes sense. Rolling structures, shorter tenors, and tactical overlays belong here.

Problems arise when flexibility is assumed rather than earned.

If governance is unlikely to tolerate significant mark-to-market swings, or if liquidity is finite under stress, then apparent flexibility is illusory. In those cases, leaving economics open does not preserve choice. It postpones constraint.

Flexibility that cannot be exercised safely is not flexibility. It is deferred risk.

Separate structural and tactical hedging

One of the most common sources of confusion in long-dated rates hedging is the attempt to use a single structure to serve two purposes.

Structural hedging exists to protect enduring exposures. It should be boring, stable, and largely immune to short-term market noise.

Tactical hedging exists to manage transitory risk. It is inherently active, responsive, and sensitive to timing.

When these two objectives are combined into one hedge, neither is achieved properly. Tactical adjustments undermine structural protection. Structural rigidity constrains tactical response. The result is usually a hedge that looks clever early and fragile later.

A coherent framework separates them explicitly. Each is governed, measured, and reviewed on its own terms.

Treat collateral and liquidity as design inputs

Collateral and funding are not operational afterthoughts. They are part of the hedge.

Any framework that evaluates hedging structures without modelling liquidity demand under stress is incomplete. The relevant question is not whether margin can be posted today, but whether it can be sustained through prolonged volatility without forcing asset sales elsewhere in the portfolio.

If a hedge relies on future liquidity being abundant, that reliance should be acknowledged and approved as a risk in its own right, not discovered mid-cycle.

Governance is not external to the hedge

Governance is often treated as an exogenous constraint: something to manage around rather than design for.

In long-dated hedging, governance is a structural input.

Who can intervene? Under what conditions? How much volatility is tolerated before action is required? How do incentives change when leadership turns over?

A hedge that depends on continuous institutional discipline across multiple cycles is fragile by definition. A hedge that remains robust even when governance becomes conservative is far more likely to behave.

The end state

A coherent long-dated rates hedging framework does not optimise for elegance, peer comparison, or short-term optics.

It aims for clarity about what is fixed and what is not, alignment between exposure reality and hedge behaviour, and resilience under regime change with minimal reliance on future intervention.

When these conditions are met, hedging becomes less visible. It stops demanding attention. It ceases to generate recurring debate.

That is not a failure of risk management. It is usually a sign that the hard decisions were made early, rather than deferred.

12. Structural Diagnostic – Self-Assessment

By this stage, most readers will already have an intuition about whether their rates hedging programme is robust or fragile. This section is designed to make that intuition explicit.

It is not a remediation guide. It is not a playbook. It is a way to identify whether the hedge is likely to behave when conditions become uncomfortable.

The diagnostic focuses on structure, not outcomes. A hedge that has performed well to date can still be structurally fragile. A hedge that has been painful can still be coherent.

The questions below are deliberately binary. Ambiguous answers are usually revealing.

A five-question structural screen

1. Does the hedge close the economics of the exposure, or does it defer them?

If material cashflows, funding costs, or roll decisions remain exposed to future conditions, the hedge is deferring economics rather than eliminating them.

2. Is the hedge tenor aligned with contractual commitment, not aspirational holding period?

If the hedge assumes the exposure will persist because “that is the intent”, termination risk exists by design.

3. Could the hedge survive prolonged stress without forced action?

This is not about whether margin can be posted today, but whether liquidity could be sustained through extended volatility without asset sales or governance intervention.

4. Does the structure rely on continuous market access to function?
If rolling, refinancing, or repositioning is required to maintain the hedge, market access is a structural dependency.

5. Would the hedge remain acceptable if reviewed by a different governance group?
If the hedge only works under the current risk appetite, reporting tolerance, or leadership, it is fragile.

Interpreting the answers

The diagnostic is not scored to create false precision. Its purpose is to identify concentration of fragility.

- **Zero or one concern identified**
  The structure is likely coherent. Ongoing monitoring is appropriate, but no immediate structural action is indicated.
- **Two or three concerns identified**
  The hedge may function in benign environments but is vulnerable under regime stress. The cost of inaction should be examined explicitly.
- **Four or more concerns identified**
  Failure is a question of timing rather than probability. The hedge is likely to force action under adverse conditions.

These outcomes are not judgments. They are indicators.

A critical warning

Identifying structural fragility does not imply that immediate action is required.

Unwinding or restructuring long-dated hedges impulsively can be more damaging than living with an imperfect structure. Transition risk is real, often underestimated, and rarely symmetric.

The purpose of this diagnostic is to surface hidden assumptions and deferred risks. It is not to trigger reactive decision-making.

Any response should be deliberate, sequenced, and informed by a clear understanding of costs, constraints, and alternatives.

Why this matters

Most long-dated hedging failures do not occur because institutions were unaware of risk. They occur because structural weaknesses were tolerated until circumstances removed all good options.

This diagnostic is intended to shift that recognition earlier, when choices still exist.

13. Governing Principles

Long-dated rates hedging does not fail because the rules are unclear. It fails because the wrong rules are applied.

The principles below are not best practice. They are not aspirational. They are the conditions under which hedging survives contact with time, regime change, and governance reality.

They are deliberately blunt.

1. Close economics where commitment is real.
If an exposure is contractual, durable, and unlikely to change, deferring economics adds risk rather than reducing it.

2. Do not hedge aspiration.
Tenor matching only works when holding periods are contractual, not aspirational. Intent is not a substitute for commitment.

3. Flexibility that cannot be exercised safely is not flexibility.
Rolling structures and short tenors only preserve choice if liquidity and governance allow that choice to be used under stress.

4. Separate structural protection from tactical positioning.
One hedge cannot reliably serve both purposes. Combining them guarantees confusion and poor outcomes.

5. Treat collateral and funding as part of the hedge.
If a structure cannot be funded through stress, it does not behave, regardless of how well it neutralises duration.

6. Design for the regime you will face, not the one you are in.
Long-dated hedges must survive rate resets, volatility spikes, and leadership changes. Anything less is incomplete.

These principles are not restrictive. They are clarifying.

They do not dictate instruments. They dictate honesty.

14. Conclusion – What It Means for a Hedge to Behave

Long-dated rates hedging is usually judged on how it looks at inception.

Risk reports are clean. Sensitivities are neutral. Entry pricing appears reasonable. Governance is comfortable.

None of that determines whether the hedge will survive.

Hedges fail not because interest rates move unexpectedly, but because liquidity tightens, funding costs persist, and deferred decisions collide with stress and governance intervention. By the time that happens, the structure has already done its damage.

A hedge behaves only if it delivers the intended economic outcome across regimes, without forcing action when options are worst.

That outcome is determined at inception. Uncertainty can be resolved early, or it can be deferred into future environments. It cannot be eliminated by familiarity, reporting optics, or instrument choice alone.

Structures that close economics concentrate risk upfront and reduce dependence on future liquidity, funding conditions, and decision-making. Structures that defer economics rely on discipline, access, and tolerance that often disappear when they are needed most.

Neither approach is universally right. What matters is whether the choice is deliberate and aligned with exposure reality rather than convenience.

Long-dated rates hedging is not a technical optimisation exercise. It is a commitment decision.

When economics are closed where commitment is real, and flexibility is used only where it can actually be exercised, hedging stops drawing attention. It stops forcing debate. It stops becoming the problem it was meant to solve.

That is what it means for a hedge to behave.

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, collateral, and balance-sheet risk management from a structural and implementation perspective.

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.