PARA BELLUM ADVISORS

PRACTITIONER PAPER

Designing Long-Dated Rates Hedges That Work

Why Familiar Structures Fail and What Endures Across Regimes

By Mike Duncan, Para Bellum Advisors

Version 3, December 2025

www.parabellumadvisors.com

Executive Summary **3**
1. Why Long-Dated Rates Hedges Fail **4**
2. The Nature of Long-Dated Rate Exposure **5**
3. Why Long-Dated Rates Hedges Fail **6**
4. Why Familiar Instruments Dominate **8**
5. Instrument Behaviour Under Regime Change **9**
5.1 Instruments That Close Economics 9
5.2 Instruments That Defer Economics 9
6. Tenor Matching – Necessary but Dangerous **11**
6.1 When Tenor Matching Works 11
6.2 When Tenor Matching Becomes Dangerous 11
6.3 Tenor Matching as a Commitment Device 11
7. Rolling Structures – Tactical Tool, Structural Liability **13**
7.1 What Rolling Actually Does 13
7.2 The Illusion of Control 13
7.3 Where Rolling Makes Sense 13
8. Collateral, Liquidity, and Regime Stress **14**
8.1 Rates Hedging Is Now Capital-Intensive 14
8.2 Pro-Cyclical Margin as the Failure Mechanism 14
8.3 The "Cheap Hedge" Fallacy 14
8.4 Liquidity Stress Beats Mark-to-Market Volatility 14
8.5 Emerging Markets as the Stress Test 15
9. What Right Looks Like – Worked Examples **16**
10. A Coherent Framework for Long-Dated Rates Hedging **18**
10.1 Two Questions That Matter 18
10.2 Close Economics Where You Can 18
10.3 Use Flexibility Only Where It Is Real 18
10.4 Separate Structural and Tactical Hedging 19
10.5 Treat Collateral and Liquidity as Design Inputs 19
10.6 Governance Is Not External to the Hedge 19
11. Structural Diagnostic – Self-Assessment **20**
12. Governing Principles **21**
Conclusion: What It Means for a Hedge to Behave **22**

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. 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.

> *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.*

1. 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. And they occur even when execution is competent and policies are followed.

Between 2020 and 2023, many institutional investors learned this directly. 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.

> *Most portfolios answer the question of which risks to defer accidentally. They use familiar instruments, optimise entry pricing, prioritise flexibility, and without stating it explicitly, defer key economic risks into the future.*

A hedge behaves only if it locks economics over its full horizon, survives regime change without forcing action, and does not create liquidity stress that overwhelms the original risk. Optics, short-term comfort, and assumptions about future decision-making under pressure do not count.

The 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.

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.

> *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.

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.

> *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.*

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.

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.

Short-dated logic applied to long-dated exposure

Hedges designed using tools that work for tactical positioning – duration matching, DV01 neutrality, frequent rebalancing – 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.

Governance 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, and 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

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. 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, and governance intervenes when options are worst.

> *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.*

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.

Standard instruments offer immediate advantages unrelated to long-term behaviour: visible liquidity, established accounting and regulatory treatment, compatibility with existing systems, and 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.

> *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. Standardisation does not eliminate risk. It obscures it.*

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, and decisions that must be made under less favourable circumstances.

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.

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.*

5.1 Instruments That Close Economics

Structures that close economics do something unfashionable: they replace flexibility with certainty. Examples include 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, and 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.

5.2 Instruments That Defer Economics

Interest rate swaps fix one thing – the fixed rate – and leave everything else open: floating-rate resets, funding conditions, margin requirements, and liquidity stress. Rolling futures and short-dated instruments go further, embedding 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

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, floating legs reprice sharply, net cash outflows become persistent, margin demands increase during volatility, and 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 vs bilateral structures

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, often exhibiting lower margin velocity and fewer forced liquidity events. 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, roll costs accumulate, basis risk compounds, and 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. Mark-to-market volatility is survivable if no action is required. Deferred cashflows, margin stress, and forced decisions are not.*

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. 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.

> *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. More often, it is replaced with softer language: "strategic allocation", "long-term intent", "core exposure". None of these are contracts. They are preferences.*

6.1 When Tenor Matching Works

Tenor matching behaves well when commitment is real and enforceable. Typical examples: project finance loans held to maturity, infrastructure debt with contracted cashflows, fixed-rate bonds matched to liabilities, and asset-liability matched portfolios. In these cases, the hedge 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.

6.2 When Tenor Matching Becomes Dangerous

Problems arise when tenor matching is applied to exposures that are not contractually durable: strategic bond portfolios, credit mandates with turnover, balance-sheet duration that can be resized, or exposures embedded in multi-asset portfolios. When the exposure changes – through rebalancing, asset sales, mandate shifts, or governance decisions – the hedge no longer fits. The institution faces a set of bad choices: unwind the hedge and crystallise mark-to-market, maintain an over-hedged or misaligned position, or layer additional hedges and increase complexity. None of these outcomes were priced at inception.

6.3 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 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.*

Tenor matching: when it is appropriate and when it is not

Appropriate when: exposure has a legal or contractual end date; early exit is unlikely or economically irrelevant; governance accepts interim volatility without intervention.

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.

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. For tactical positioning, that instinct is correct. For long-dated exposure, it is usually wrong.

7.1 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.

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, and 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.

> *Rolling structures often underperform expectations without any obvious failure point. The drag is incremental, persistent, and difficult to attribute.*

By choosing a rolling structure, the institution accepts that pricing will be reset repeatedly, liquidity must be available at every roll, market access cannot be assumed, and decisions must be made under future conditions. Those costs are rarely priced explicitly. They are treated as operational details rather than structural risks.

7.2 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.

7.3 Where Rolling Makes Sense

Rolling structures are appropriate when flexibility is genuinely required: 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.

> *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.*

8. Collateral, Liquidity, and Regime Stress

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.*

8.1 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. 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.

8.2 Pro-Cyclical Margin as 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.

8.3 The "Cheap Hedge" Fallacy

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.*

8.4 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.

8.5 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. In many emerging rates markets, long-dated swaps are illiquid or unavailable, futures markets are thin or discontinuous, clearing is limited or punitive, and basis risk is structural rather than 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. Margin is often denominated in hard currency. FX and rates stress frequently coincide. As a result, hedging errors migrate rapidly into liquidity crises.

> *Institutions that operate successfully in emerging rates markets tend to match hedges closely to contractual exposures, avoid perpetual rolling structures, and accept mark-to-market volatility in exchange for economic certainty. 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.*

9. What Right Looks Like – Worked Examples

What distinguishes robust hedging is not sophistication. It is alignment. The following examples 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 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. That approach would have minimised upfront commitment and headline cost. 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.

The resulting structure accepted discomfort in reported volatility in exchange for economic certainty. Financing costs were fixed. Liquidity demands were predictable. Governance intervention became unlikely because the hedge behaved consistently with the exposure it protected. Nothing about the structure was clever. It simply reflected reality rather than aspiration.

Infrastructure: Matching Economics, Not Optics

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 result was a hedge that looked uncomfortable in some environments and uncompetitive in others. But it removed the risk of forced action under stress, which was the true objective. The hedge behaved because it acknowledged that infrastructure exposures are not traded positions. They are commitments.

Family Office: Separating Structure from Discretion

Family office balance sheets often combine long-term capital with discretionary flexibility. The common mistake is to hedge the entire balance sheet using a single approach – usually short-dated and rollable – to preserve perceived flexibility. Over time, this blurs the distinction between what must be protected and what can be adjusted.

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 separation reduced complexity rather than increasing it. Governance clarity improved. Tactical decisions no longer interfered with structural protection. The hedge behaved because it reflected the reality of the balance sheet rather than forcing a false uniformity.

> *In each case, exposure reality was prioritised over convenience. Economic commitment was made explicit. Flexibility was used sparingly and deliberately. None of the outcomes depended on forecasting rates correctly. They depended on acknowledging what could not change and designing the hedge around that fact.*

10. A Coherent Framework for Long-Dated Rates Hedging

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. 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.

10.1 Two Questions That Matter

Before instruments, pricing, or execution are discussed, any long-dated hedging programme should be able to answer the following:

> **The two questions**
>
> Which economics must be closed today?
>
> Which risks can legitimately remain flexible without threatening survival?
>
> 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 sustainable. The discipline lies in making the boundary explicit.

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

10.3 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.*

10.4 Separate Structural and Tactical Hedging

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

10.5 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 question is not whether margin can be posted today, but whether it can be sustained through prolonged volatility without forcing action 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.

10.6 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 central 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.

11. Structural Diagnostic – Self-Assessment

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 focus on structure, not outcomes. They are deliberately binary. Ambiguous answers are usually revealing.

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.

Interpretation: zero or one concern identified – the structure is likely coherent; two or three concerns – the hedge may function in benign environments but is vulnerable under regime stress; four or more concerns – failure is a question of timing rather than probability.

> *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, not to trigger reactive decision-making.*

12. 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.

Six Governing Principles

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 enforceable. 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.

Conclusion: What It Means for a Hedge to Behave

Long-dated rates hedging is often judged on how it looks today. Risk reports are clean. Sensitivities are neutral. Entry pricing appears reasonable. Governance is comfortable. Those outcomes say very little about whether the hedge will survive the environments that matter.

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 at the worst possible moment.*

That requires accepting a simple truth: uncertainty can either be resolved at inception or deferred into the future. It cannot be eliminated by report design or instrument familiarity. Structures that close economics concentrate risk early and reduce dependence on future conditions. Structures that defer economics rely on liquidity, funding, and discipline that may not exist when they are needed most.

Neither approach is inherently right or wrong. What matters is that the choice is made consciously, aligned with exposure reality, and governed accordingly. Long-dated rates hedging is not a technical exercise. It is a commitment decision.

> *When that decision is made honestly, hedging stops demanding attention. It stops generating surprises. It stops becoming the problem it was meant to solve. That is what it means for a hedge to behave.*

Further Reading

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

www.parabellumadvisors.com/insights/

About Para Bellum Advisors

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

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

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

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

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

For more information: www.parabellumadvisors.com

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

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