



PARA BELLUM ADVISORS

PRACTITIONER PAPER

Why Rates Hedges Don't Behave

**Structural mistakes, familiar tools,
and deferred regime risk**

By Mike Duncan, Para Bellum Advisors

Version 3, December 2025

www.parabellumadvisors.com

Informational only
Not investment advice

Executive Summary	4
1. Why Rates Hedges Fail in Practice	5
2. Rates Hedging Is Not One Problem	6
2.1 Short-Term Risk Management vs Long-Term Economic Certainty.....	6
2.2 Volatility Reduction Is Not Risk Elimination	7
2.3 Why One Hedge Cannot Solve Both.....	7
3. Why Familiar Instruments Dominate	8
3.1 Liquidity as Comfort, Not Robustness	8
3.2 Optics, Standardisation, and Committee Safety	8
3.3 The Cost of Deferring Decisions	9
4. What Rates Hedges Actually Do.....	10
4.1 What Rates Hedges Reduce.....	10
4.2 What Rates Hedges Do Not Eliminate.....	10
4.3 Why Outcomes Diverge Over Time.....	11
5. Collateral Blindness	12
5.1 Margin as Capital, Not Plumbing.....	12
5.2 Liquidity Stress Is Ignored at Inception	12
5.3 Why Collateral Risk Only Shows Up in Crises	12
5.4 Clearing, Bilateral Structures, and Where Stress Appears.....	13
5.5 The Hidden Bet in Most Hedges.....	13
6. Futures, Swaps, and the Illusion of Control.....	14
6.1 Futures as Tactical Tools.....	14
6.2 Swaps as Deferred Decisions	14
6.3 The "Cheap Hedge" Fallacy	15
6.4 Control Versus Commitment.....	15
7. Governance vs Economics	16
7.1 Accountability Deferral as a Design Feature.....	16
7.2 The Tyranny of Mark-to-Market.....	16
7.3 Policy Rigidity and the Illusion of Control	17
7.4 Governance as a Structural Input	17
8. Failure Modes in Practice.....	18
8.1 Death by Accumulation Failure.....	18
8.2 Forced Re-Hedging During Stress	18
8.3 Liquidity-Driven Asset Sales	19
8.4 Unforeseen Events as Narrative Cover	19

9.1 Match Instrument to Objective	20
9.2 Accept Discomfort Upfront	20
9.3 Reduce Future Decision Points.....	20
9.4 Separate Structural and Tactical Hedging.....	21
Conclusion	22

Executive Summary

Rates hedging is still widely treated as a problem of measurement and control. Duration is quantified. DVO1 is neutralised. Instruments are selected for liquidity. Success is judged by the stability of reported outcomes.

That framing works – up to a point.

Most rates hedging failures do not occur because institutions misjudge the direction of interest rates. They occur even when rates evolve broadly as expected, when hedges perform to specification, and when risk reports remain clean. The failure is not predictive. It is structural.

Modern hedging frameworks prioritise what is visible, quantifiable, and defensible in governance processes. Sensitivity metrics, hedge ratios, and short-term mark-to-market volatility dominate decision-making because they can be monitored continuously and explained cleanly. The risks that actually determine whether a hedge survives across regimes – liquidity strain, collateral dynamics, funding persistence, and forced decision-making under stress – are addressed later, if at all.

As a result, hedges are designed to look correct at inception rather than to behave correctly over time.

Familiar instruments dominate not because they are structurally robust, but because they are liquid, standardised, and comfortable in committee settings. They minimise immediate discomfort and preserve the appearance of flexibility. What they do not do is eliminate uncertainty. They relocate it into future regimes, future liquidity conditions, and future governance decisions.

Hedges therefore often behave exactly as designed. The problem is what they were designed to optimise. Optics, reversibility, and short-term stability are prioritised over endurance. When conditions change – not catastrophically, but enough for structure to matter – the risks that were deferred become binding. At that point, the hedge itself becomes a source of instability.

Rates hedges don't fail simply because markets move. That exposes the failure. They fail when deferred structural risks become binding.

A rates hedge behaves only if it delivers its intended economic outcome across regimes without forcing intervention at the wrong moment. Achieving that requires accepting a simple constraint: uncertainty can either be resolved at inception or deferred into the future. It cannot be eliminated by clean reports, familiar instruments, or sensitivity metrics.

When that distinction is ignored, hedges appear stable until they are not. When it is acknowledged, rates hedging stops being a source of disruption and becomes structurally resilient.

1. Why Rates Hedges Fail in Practice

Rates hedging failures are rarely dramatic. They do not usually arrive as sudden losses, broken trades, or obvious mistakes. More often, they emerge quietly – through outcomes that drift away from expectations even though hedges appear to be in place and risk reports remain reassuring.

This creates confusion. When hedges behave badly, the instinctive explanation is that markets have moved in unexpected ways or that forecasts were wrong. Yet many disappointing outcomes occur in environments that are broadly consistent with prior expectations. Rates may rise or fall, but not disorderly. Volatility may increase, but not catastrophically. Hedges often perform broadly in line with their specifications.

The problem is not that rates hedges fail to respond to market moves. It is that they are rarely designed to govern what happens after the initial response.

Most rates hedging frameworks focus on measurable exposure at a point in time. Duration is neutralised. Sensitivities are controlled. Short-term volatility is reduced. These outcomes are observable, defensible, and easy to monitor. They create a sense of control.

What they do not address is how risk evolves once exposure persists across time. Funding costs compound. Liquidity demands change with regimes. Collateral requirements rise non-linearly under stress. Governance tolerance shifts as performance drifts or volatility returns. These dynamics sit outside traditional sensitivity metrics, yet they are precisely what determine whether a hedge remains viable.

A rates hedge behaves only if it delivers its economic outcome across regimes without forcing action at the wrong moment. A hedge fails when it becomes economically untenable to maintain or forces portfolio-level decisions under stress – even if sensitivity metrics remain stable.

As a result, many rates hedges do exactly what they are designed to do – and still fail economically. The failure lies not in execution or forecasting, but in the underlying framing of the problem. Risks that are uncomfortable to resolve at inception are deferred into future environments where they become harder, not easier, to manage.

Understanding why rates hedges do not behave therefore requires stepping away from instruments and pricing, and examining structure instead: what risks are removed, what risks are transformed, and which risks are consciously or unconsciously left to surface later.

2. Rates Hedging Is Not One Problem

Most rates hedging frameworks rest on an implicit assumption: that interest-rate risk is a single, homogeneous exposure that can be neutralised with a standard set of tools.

That assumption does not hold.

Rates hedging spans at least two distinct objectives that are routinely conflated. The first is short-term risk management – reducing near-term volatility, stabilising reported outcomes, and avoiding drawdowns that attract scrutiny. The second is long-term economic certainty – controlling financing costs, preserving liquidity across regimes, and avoiding forced decisions when conditions deteriorate.

These objectives are not additive. Structures that perform well against one often perform poorly against the other.

In practice, frameworks default to the objective that is most visible and easiest to defend: short-term stability. The cost of that choice is rarely immediate. It is paid overtime, through persistent cashflows, rising liquidity demands, and growing dependence on favourable funding and governance conditions.

Rates exposure does not fail symmetrically. Portfolios rarely break because reported volatility becomes uncomfortable. They break because cashflows persist, liquidity is consumed, and governance is forced to intervene under stress.

2.1 Short-Term Risk Management vs Long-Term Economic Certainty

Short-term rates hedging is about control. Its objective is to reduce observable sensitivity, smooth performance, and limit fluctuations over reporting horizons that matter to stakeholders. The tools that serve this objective – liquid instruments, frequent rebalancing, and volatility suppression – are well established.

Long-term rates hedging addresses a different problem. Its objective is not to suppress volatility, but to govern the economic consequences of holding exposure through multiple regimes. This includes financing costs, liquidity demands, and the risk that adverse conditions force decisions that destroy value.

Problems arise when instruments designed for short-term control are repurposed as long-term solutions. Over time, the dominant risks shift away from price sensitivity and toward persistence, funding, and governance. Most frameworks fail because they do not recognise this transition.

Short-term objective	Long-term objective
Reduce observable sensitivity	Govern economic consequences across regimes
Smooth performance over reporting horizons	Control financing costs and liquidity demands
Limit fluctuations visible to stakeholders	Avoid forced decisions when conditions deteriorate
Tools: liquid instruments, frequent rebalancing	Tools: tenor-matched structures, committed hedges

2.2 Volatility Reduction Is Not Risk Elimination

Volatility is episodic. It is often uncomfortable, but survivable if no action is required. Economic exposure is persistent. It asserts itself through cashflows, liquidity consumption, and balance-sheet pressure.

A hedge can materially reduce reported volatility while still embedding structural fragility. It can suppress mark-to-market swings while creating obligations that must be met regardless of market conditions.

The most damaging error in rates hedging is assuming that reducing volatility is equivalent to reducing risk. Once exposure persists across regimes, volatility is rarely the binding constraint. Liquidity and governance are.

2.3 Why One Hedge Cannot Solve Both

Attempting to use a single hedge to manage both short-term volatility and long-term economic forces creates trade-offs that are rarely acknowledged explicitly.

In practice, those trade-offs are resolved in favour of short-term comfort. Instruments are selected for reversibility and familiarity rather than for their ability to survive regime change without intervention.

The consequence is not immediate failure, but gradual fragility. Deferred risks accumulate. Cashflows persist. Margin demands grow. Funding becomes visible. Governance tolerance tightens.

By the time this becomes obvious, the hedge no longer behaves as protection. It behaves as a constraint.

3. Why Familiar Instruments Dominate

If these dynamics are so persistent, it is reasonable to ask why institutions continue to rely on the same hedging structures.

The answer lies in incentives rather than ignorance.

Rates hedging decisions are made in committees, reviewed by consultants, constrained by policy, and inherited by future decision-makers. In that environment, familiarity, defensibility, and reversibility carry significant weight.

Familiar instruments minimise immediate friction. They are liquid, standardised, and well understood. They fit existing systems and reporting frameworks. They reduce career risk by aligning with peer practice. What they do not do is ensure structural robustness.

The dominance of familiar instruments reflects organisational logic rather than economic optimisation. The costs of that logic are deferred in time and dispersed across future regimes and governance structures.

3.1 Liquidity as Comfort, Not Robustness

Liquidity is a genuine advantage. It allows positions to be adjusted, resized, or exited with relatively low friction in normal conditions. In uncertain environments, this flexibility feels reassuring. However, liquidity is often mistaken for resilience.

The ability to trade does not guarantee that a hedge will behave under stress. Liquidity is state-dependent. It can evaporate precisely when it is most needed. Basis can widen. Funding can tighten. Instruments that appear flexible in calm markets can become sources of instability when conditions change.

In long-dated hedging, reliance on liquidity is itself a structural bet. It assumes that markets will remain accessible and that funding will remain available when adjustments are required. That assumption is rarely made explicit, let alone governed.

3.2 Optics, Standardisation, and Committee Safety

Rates hedging decisions are often evaluated on how they look rather than how they behave. Clean risk reports, stable daily P&L, and the absence of visible losses are rewarded. Structures that produce these outcomes gain institutional support, even if they embed longer-term fragility.

Standardisation reinforces this bias. Instruments that fit consultant frameworks and peer benchmarks are easier to justify and harder to challenge. When outcomes deteriorate years later, responsibility diffuses. The original decision still appears reasonable in context.

This isn't accidental. It's rational behaviour inside flawed incentives. The result is a systematic preference for hedges that optimise comfort today at the expense of resilience tomorrow.

3.3 The Cost of Deferring Decisions

Familiar instruments often preserve the appearance of optionality. Short tenors and rollable structures allow decisions to be revisited. Nothing appears locked in. In reality, this transfers risk forward.

Decisions that could have been resolved at inception are deferred into future environments. Those environments are unlikely to be more favourable. Liquidity may be scarcer. Volatility may be higher. Governance tolerance may be lower.

By the time deferred decisions must be made, the hedge has already accumulated path dependency. What was once optional becomes compulsory.

This is the hidden cost of familiarity. It is not that familiar instruments are wrong. It is that they relocate risk into places that are harder to see and harder to manage.

4. What Rates Hedges Actually Do

Rates hedges are usually described in terms of what they remove. Duration is neutralised. Sensitivity is reduced. Exposure is taken off the table. That framing is misleading.

In practice, rates hedges do not eliminate risk. They reallocate it. Certain risks become smaller and more visible. Others are displaced into areas that are less observable, less frequently modelled, and far more likely to bind under stress.

Hedging is not subtractive. It is transformative. Whether a hedge behaves or fails depends on where the risk ends up.

4.1 What Rates Hedges Reduce

Most common rates hedges are effective at a narrow but important task: reducing near-term sensitivity to interest-rate moves. They tend to reduce DVO1 and other linear sensitivity measures, dampen short-term mark-to-market volatility, and smooth reported performance across normal market fluctuations.

These effects are real. They are not cosmetic. In many contexts – particularly for tactical positioning or transitional portfolios – they are exactly what is required. What they do not do is determine whether the hedge remains economically viable once exposure persists across time.

A hedge can look perfectly controlled on a sensitivity report and still deteriorate incrementally through cashflows, funding pressure, and governance strain.

4.2 What Rates Hedges Do Not Eliminate

There are several risks that rates hedges almost never eliminate, regardless of how cleanly they neutralise duration:

- Exposure to future rate regimes rather than point moves
- Persistent cashflow obligations tied to floating legs or carry
- Funding and liquidity requirements that scale under stress
- Collateral dynamics that are pro-cyclical by design
- Governance intervention when losses or liquidity strain become visible

These risks are not secondary. Over long horizons, they dominate outcomes. What makes them dangerous is not their size, but their invisibility. They accumulate during periods when the hedge appears to be working. By the time they surface, flexibility has usually disappeared and its now deadly.

4.3 Why Outcomes Diverge Over Time

The divergence between apparent hedge success and eventual failure reflects a simple asymmetry. Some risks are immediate and observable. They show up as price moves, P&L swings, or limit breaches. Institutions are good at responding to these.

Other risks are latent. They build through persistent carry, repeated roll costs, rising margin demands, and growing dependence on favourable funding and governance conditions. These risks do not trigger alarms early. They become visible only after time has removed the option to adjust cheaply.

As a result, hedging programmes can appear robust quarter after quarter while quietly drifting toward fragility. When failure finally occurs, it is often attributed to an external shock rather than to the structure that made the outcome inevitable.

Rates hedging failures are rarely caused by a single bad decision. They are caused by small design choices whose consequences only become visible once time has done its work.

5. Collateral Blindness

Collateral is still widely treated as a secondary consideration in rates hedging. This is no longer defensible. In modern markets, collateral is not an operational detail. It is a primary transmission channel through which hedges fail.

5.1 Margin as Capital, Not Plumbing

Post-crisis market structure has fundamentally altered the economics of rates hedging. Clearing mandates, conservative margin models, and daily variation margin have shifted risk away from counterparty credit and toward liquidity. This has improved systemic resilience. It has made individual hedges far more capital-intensive.

Initial margin immobilises capital from inception, typically posted as high-quality securities rather than pure cash, though the economic impact remains material even when securities are eligible collateral. **Variation margin** introduces cashflow volatility that scales precisely when markets are stressed. These demands are not neutral over time. They are pro-cyclical.

Treating margin as plumbing rather than capital leads to systematic underestimation of hedging cost and risk.

5.2 Liquidity Stress Is Ignored at Inception

Most hedges are approved under benign assumptions: stable funding markets, ample liquidity buffers, and orderly market functioning. Stress scenarios are often discussed abstractly, if at all. The focus remains on price sensitivity rather than on cashflow survivability.

When volatility rises and liquidity tightens, these assumptions collapse simultaneously. Margin requirements can easily increase three-to-fivefold during volatility spikes, with initial margin scaling particularly sharply for long-dated exposures as model conservatism intensifies.

At that point, the hedge itself becomes a source of strain. Margin calls accelerate. Funding costs rise. Liquidity buffers are consumed. Decisions that were comfortably deferred become unavoidable.

5.3 Why Collateral Risk Only Shows Up in Crises

Collateral risk is non-linear, regime-dependent, and invisible in calm markets. As long as volatility remains low and funding is abundant, margin feels manageable. This reinforces the belief that collateral is a secondary issue.

When conditions change, margin requirements rise sharply and persist. Liquidity stress compounds. What had been a background consideration becomes the dominant constraint.

This is why hedges are often unwound not at the point of maximum mark-to-market loss, but at the point of maximum liquidity strain. The decision is no longer economic. It is operational.

5.4 Clearing, Bilateral Structures, and Where Stress Appears

The choice between cleared and bilateral structures is often framed in regulatory or operational terms. In long-dated hedging, it is a structural decision about where stress will surface.

Cleared structures	Bilateral structures
Standardisation and counterparty safety	Greater flexibility in margin treatment through negotiated CSAs
High collateral velocity and daily liquidity demands	Concentration risk and bilateral credit exposure
Conservative margin models drive requirements	Modern CSAs can still impose significant daily variation margin
Less room for negotiation on terms	Key distinction is negotiability, not automatic reduction in velocity

Neither is universally superior. The critical question is whether the institution can sustain the form of stress the structure imposes when regimes change. A hedge that neutralises duration but overwhelms liquidity capacity does not behave.

5.5 The Hidden Bet in Most Hedges

Most rates hedges embed an unstated assumption: future liquidity will be available when it is needed. This is a bet, not a fact. It is rarely articulated, rarely priced, and rarely governed.

When that bet fails, the hedge fails with it. Collateral is the mechanism through which this failure occurs.

6. Futures, Swaps, and the Illusion of Control

Many rates hedging instruments create a powerful sense of control. Duration is neutralised. Positions can be adjusted. Exit appears possible. Reports look tidy. This sense of control is often illusory.

Different instruments distribute risk across time, liquidity, and governance in very different ways. When horizons extend, those differences matter far more than entry pricing or headline sensitivity. What looks flexible in a dashboard can become rigid in practice.

6.1 Futures as Tactical Tools

Bond futures are among the most effective instruments available for managing interest-rate exposure. They are liquid, transparent, and operationally efficient. For tactical positioning, they are hard to beat.

Problems arise when futures are used as structural hedges. A futures hedge does not exist unless it is maintained. Maintaining it requires rolling. Each roll is a new trade, executed under the conditions of the day. Over time, this introduces structural features that are rarely acknowledged at inception:

- Exposure to roll timing and curve shape
- Cumulative roll costs that only look small one period at a time
- Basis risk that changes character across interest rate regimes
- Dependence on continuous market access

Most futures-based "long-term" hedges do not fail on a bad roll. They fail because the institution eventually realises it has been running an unacknowledged active strategy for years – and no longer has the appetite, liquidity, or governance tolerance to keep doing so.

Rolling a futures hedge does not turn it into a long-dated solution. It turns it into a commitment to keep deciding under whatever conditions the future delivers.

Futures are excellent instruments for what they were designed to do. For tactical positioning, short-term hedging, and situations where flexibility genuinely has value, they are often the right choice. The error lies in repurposing them as structural solutions for long-dated exposures simply because they are familiar and liquid.

6.2 Swaps as Deferred Decisions

Interest rate swaps are often described as instruments that "lock in" rates. This description is incomplete. A swap fixes one element of the exposure – the fixed rate – while leaving others open:

- Floating-rate resets remain exposed to future regimes
- Funding costs are not fixed
- Margin requirements vary with volatility
- Liquidity demands scale non-linearly under stress

In effect, a swap resolves one decision at inception and defers several others into the future. Many institutions only discover what they did not fix when rates move and stay there. The swap continues to neutralise duration exactly as designed, while persistent net cash outflows quietly become the dominant risk.

Nothing breaks. Reports remain clean. The hedge simply becomes expensive to live with.

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

6.3 The "Cheap Hedge" Fallacy

Many rates hedges persist for one simple reason: they look cheap on day one. They may offer tight spreads, low upfront cost, favourable carry in the prevailing environment, and familiarity that reduces execution friction. What this framing ignores is lifecycle cost.

For long-dated hedges, the dominant costs are rarely incurred on day one. They arise through sustained negative carry as regimes shift, margin funding during extended volatility, liquidity drag that only appears under stress and forced restructuring when governance tolerance is exhausted.

Hedges that appear cheap at entry are often those that defer the largest costs into the future – precisely where institutions have the least flexibility to respond. Cheap hedges are rarely cheap to live with.

6.4 Control Versus Commitment

The appeal of futures and swaps lies in their apparent reversibility. Positions can be adjusted. Trades can be unwound. Nothing appears final. This reversibility feels like control. It is often the opposite.

True control in long-dated hedging does not come from the ability to change one's mind. It comes from removing the need to do so. The hedges that survive stress are rarely the most flexible ones. They are the ones that leave the fewest decisions to be made when conditions deteriorate.

A hedge that requires ongoing discretion to remain viable is fragile by definition. The more a hedge relies on future judgement, the less it can be relied upon to behave.

7. Governance vs Economics

Rates hedging failures are often blamed on markets. More often, they are caused by governance. This is not an indictment of governance processes themselves. It is an observation about how incentives, reporting structures, and decision-making authority interact with long-dated risk.

The central problem is that governance tends to reward deferral, while economic resilience requires commitment.

7.1 Accountability Deferral as a Design Feature

Many commonly used hedging structures diffuse responsibility by design. Short tenors, rolling instruments, and incremental adjustments spread decisions across time and across people. No single decision appears decisive. Each roll, rebalance, or extension can be justified on its own terms.

This diffusion is often described as prudent governance. In reality, it defers accountability. When outcomes deteriorate, there is no clear point at which a wrong decision was made. Each step was reasonable in context. Responsibility is diluted across committees, reporting cycles, and market conditions.

A hedge that relies on continuous governance discipline across multiple regimes is fragile by definition. It assumes that future decision-makers will have the same risk tolerance, the same incentives, and the same willingness to accept discomfort as those who initiated the hedge. That assumption rarely holds.

7.2 The Tyranny of Mark-to-Market

Mark-to-market volatility exerts disproportionate influence over hedging decisions because it is immediate, visible, and reportable. Daily P&L moves trigger attention. Breaches of volatility thresholds prompt discussion. Losses that appear on reports invite explanation.

Cashflow risk, by contrast, is deferred. Funding strain accumulates gradually. Liquidity stress only becomes visible once buffers are exhausted. This asymmetry skews behaviour.

Hedges that generate smooth mark-to-market outcomes are favoured even if they embed persistent cash outflows. Structures that produce volatile marks, but stable economics are viewed with suspicion, even when they are more robust over time.

Mark-to-market volatility is often survivable. Liquidity stress is not. No CIO ever lost their job for suppressing volatility. Many have for liquidity events.

7.3 Policy Rigidity and the Illusion of Control

Formal hedging policies are designed to impose discipline. They define hedge ratios, permissible instruments, and rebalancing rules. In doing so, they reduce discretion and improve consistency. Problems arise when policies substitute rules for judgement.

Policies that mandate specific instruments or hedge ratios without accounting for regime change implicitly assume that the future will resemble the past. They also assume that compliance is synonymous with risk management.

In practice, rigid policies can force institutions to maintain hedges that are economically misaligned simply because unwinding them would breach process rather than reduce risk. When stress arrives, governance often intervenes abruptly, suspending policies or overriding them entirely. By that point, the hedge has already become a liability.

7.4 Governance as a Structural Input

Governance is often treated as something external to hedge design – a constraint applied after the structure is chosen. In practice, governance is a core structural input. Hedging outcomes are shaped not only by markets, but by who has authority to act, under what conditions, and with what tolerance for interim discomfort.

Effective rates hedging acknowledges this explicitly at inception. It recognises that governance tolerance is not static, that incentives change as performance deteriorates, and that leadership, committee composition, and risk appetite will evolve over time. A hedge that relies on sustained institutional resolve, stable personnel, or consistent risk tolerance across regimes is fragile by construction.

Robust hedges are designed to survive governance pressure rather than assume it away. Designing for governance weakness is not cynical. It is a realistic acknowledgement of how institutions actually behave under stress.

8. Failure Modes in Practice

When the structural, behavioural, and governance dynamics described so far combine, the same failure modes appear repeatedly across institutions and cycles. They are not edge cases. They are the predictable consequences of deferring economic resolution into future regimes.

8.1 Death by Accumulation Failure

The most common failure mode is not dramatic. A hedge is put in place. It neutralises duration. Reports look clean. Over time, small costs accumulate negative carry persists as regimes shift, roll costs compound, margin funding absorbs liquidity, and governance tolerance erodes gradually. No single period looks catastrophic. Performance drifts. Returns lag peers.

Eventually, it becomes deadly, the cumulative drag becomes impossible to ignore. At that point, restructuring is discussed – usually under less favourable conditions than those that existed at inception.

The Accumulation Pattern in Practice

An institutional investor maintains a 10-year receiver swap, rolled forward every two years to preserve approximate duration neutrality against a long-dated liability portfolio. Initially, the hedge performs well. Over a five-year period, however, several dynamics compound.

The swap is repeatedly reset at progressively less favourable fixed rates as the yield curve shifts. Negative carry persists as the floating leg resets higher while the economic value of protection deteriorates. Periodic mark-to-market losses trigger governance discussions, though no single period appears catastrophic. By year six, cumulative opportunity cost relative to peers becomes visible in performance reports.

The decision to unwind occurs not because the hedge failed mechanically, but because sustained underperformance exhausted institutional patience. The restructuring happens at precisely the point when starting fresh is least attractive. Nothing broke. The hedge simply revealed what it always was: a structure optimised for comfort at inception rather than endurance across regimes.

8.2 Forced Re-Hedging During Stress

As rates move sharply and volatility rises, margin calls accelerate. Liquidity buffers are consumed. Governance attention intensifies. At precisely the point when optionality is most valuable, it disappears.

Hedges are resized, restructured, or unwound not because the original thesis was wrong, but because liquidity constraints or governance pressure leave no alternative. Losses are crystallised. Protection disappears. The hedge fails at the moment it was meant to matter most.

8.3 Liquidity-Driven Asset Sales

In more severe cases, hedging structures interact with the broader portfolio in destructive ways. Margin calls tied to rates hedges force the sale of unrelated assets. Liquid positions are sold to fund derivative cashflows. Portfolio composition shifts for reasons unrelated to investment conviction.

At this point, the hedge has inverted its purpose. It is no longer protecting the portfolio. It is dictating its evolution.

This outcome is often described as an unforeseen liquidity event. In reality, it is the mechanical consequence of structures that deferred liquidity risk rather than addressing it.

8.4 Unforeseen Events as Narrative Cover

When hedges fail, post-mortems cite exceptional circumstances: unprecedented volatility, unexpected regime shifts, temporary dislocations, extraordinary market conditions. These explanations are comforting. They absolve structure.

In most cases, nothing truly unforeseen occurred. Rates moved. Volatility rose. Liquidity tightened. Governance intervened. These are not tail events. They are features of markets.

What failed was not market forecasting ability. It was structural foresight – the failure to recognise at inception how the hedge would behave as exposures persisted across regimes and how governance would respond under stress.

9. From Measurement to Structural Intent

Because rates hedging failures are systematic, the solution cannot lie in better forecasting or marginal optimisation. It must lie in making structural intent explicit. Rates hedging behaves better when institutions are clear about what they are trying to achieve and honest about the trade-offs involved.

9.1 Match Instrument to Objective

Rates hedging behaves poorly when instruments are selected before objectives are made explicit. Too often, familiar tools are chosen first and expected to satisfy multiple, competing goals simultaneously. This ambiguity guarantees disappointment.

Before any instrument is selected, the intended outcome must be clearly defined. Reducing short-term volatility, locking long-term economics, preserving flexibility, and managing liquidity under stress are distinct objectives. Each implies different trade-offs and different failure modes. No single hedge can optimise all of them at once.

Most structural failures arise not because the wrong instrument was used, but because the objective was never resolved. Matching instrument choice to a clearly articulated objective is the foundation of whether the hedge will behave at all.

9.2 Accept Discomfort Upfront

Robust hedges often look uncomfortable at inception. They may introduce mark-to-market volatility. They may appear expensive relative to familiar alternatives. They may reduce optionality. This discomfort is the price of removing uncertainty.

Discomfort deferred is not discomfort avoided. It is discomfort magnified and encountered under worse conditions.

9.3 Reduce Future Decision Points

The most resilient hedges minimise the need for future intervention. They are designed so that outcomes do not depend on favourable timing, liquidity demands are predictable, and governance does not need to act under stress.

A hedge that requires constant attention is not stable. It is an ongoing negotiation with the future.

9.4 Separate Structural and Tactical Hedging

Structural hedging exists to protect enduring exposures. It should be boring, stable, and largely immune to short-term noise. Tactical hedging exists to manage transitory risk. It is active by nature and sensitive to timing.

Combining the two in a single structure guarantees confusion. Separating them clarifies intent and improves outcomes.

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.

Conclusion

Most rates hedges behave exactly as designed. The problem is that they are designed to be comfortable rather than correct.

They prioritise liquidity over structure, optics over economics, and flexibility over endurance. These choices make sense in calm markets. They fail when regimes change.

Rates hedging does not break because rates move. It breaks because decisions that should have been made upfront are deferred until they can no longer be made safely.

A hedge behaves only if it delivers the intended economic outcome across regimes, without forcing action at the worst possible moment. That requires recognising a simple constraint: uncertainty can either be resolved at inception or deferred into the future. It cannot be eliminated by sensitivity metrics, familiar instruments, or clean reports.

When acknowledged, rates hedging stops being a source of disruption. It demands minimal intervention and generates no surprises.

A hedge that runs cleanly across regimes was designed to do so from the start. It is evidence that the structure was aligned with its intent from inception.

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.

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/