



PARA BELLUM ADVISORS

PRACTITIONER PAPER

Designing FX Overlays That Behave

**A practical framework for
managing carry, path dependency,
and long-horizon FX risk**

By Mike Duncan, Para Bellum Advisors

Version 2, December 2025

www.parabellumadvisors.com

Informational only
Not investment advice

Executive Summary	4
1. FX Hedging Is Not One Problem	5
1.1 The Core Distinction	5
1.2 Why the Distinction Matters	5
1.3 The Exposure Classification Framework	5
2. Why Rolling FX Forwards Dominate	6
2.1 When Rolling Forwards Are Optimal	6
2.2 The Governance Trap	6
3. What Rolling Forwards Do	7
3.1 Capabilities and Limitations	7
3.2 The Compounding Effect: A Worked Example	7
3.3 Path Dependency in Practice	7
3.4 Quarterly Smoothing vs Long-Term Divergence	8
4. Hedge Ratios and the Myth of Neutrality	9
4.1 The Fundamental Asymmetry	9
4.2 Why 100% Hedge Ratios Persist	9
4.3 Optimal Hedge Ratios Are Context-Dependent	9
5. Carry Can be the Dominant Driver	10
5.1 Why Carry Matters More Than Volatility	10
5.2 The Base Currency Effect	10
5.3 Quantifying the Effect	10
5.4 Why Carry Is Ignored	10
6. Tenor Matching: When It Works, When It Fails	11
6.1 When Tenor Matching Works Perfectly	11
6.2 When Tenor Matching Fails Badly	11
6.3 The Hidden Risks of Long-Dated Hedges	11
7. FX Forwards vs Cross-Currency Swaps	13
7.1 What Cross-Currency Swaps Do	13
7.2 Why CCS Look Expensive	13
7.3 When CCS Are the Right Tool	13
7.4 When CCS Are the Wrong Tool	14
8. Options: The Missing Dimension	15
8.1 What Options Do in FX Overlays	15
8.2 Option Structures for FX Overlays	16
8.3 Strike and Tenor Selection	20

8.4 Why Options Are Underused.....	20
8.5 When Options Make the Most Sense.....	20
9. Mandatory FX Hedging Regimes.....	21
9.1 The Governance vs Economics Tension.....	21
9.2 Strategies Within Mandatory Regimes.....	21
9.3 MAS and APRA Considerations.....	21
9.4 Documenting the Trade-Offs.....	23
10. A Unified FX Overlay Framework.....	24
10.1 The Three-Layer Framework.....	24
10.2 How the Layers Interact: An Illustrative Example.....	24
10.3 Determining Layer Allocation.....	25
10.4 Portfolio Sizing Examples.....	25
11. Implementation.....	26
11.1 Implementation Sequence.....	26
11.2 Common Mistakes in Implementation.....	26
11.3 Troubleshooting Guide.....	27
12. Case Study: Misaligned vs Aligned Hedging.....	28
12.1 The Setup.....	28
12.2 Fund A: Conventional Approach (Misaligned).....	28
12.3 Fund B: Aligned Approach (Three-Layer Framework).....	29
12.4 The Governance Conversation.....	29
Conclusion.....	30

Executive Summary

FX hedging is routinely treated as a single problem with a single solution. It is not.

Most portfolios combine indefinite FX exposures with finite hedging instruments, then rely on convention: typically rolling short-dated FX forwards. Operationally, this works. Economically, it often fails.

Rolling FX forwards persist not because they are optimal, but because they are liquid, familiar, and governance-friendly. Cross-currency swaps are often dismissed as expensive or complex, despite being the correct tool for genuinely long-dated exposures. Options are frequently excluded altogether, even though they are the only instruments that introduce convexity and control into FX overlays.

The result is predictable:

- FX risk that is shaped in the short term but uncontrolled over time
- Carry costs that quietly compound
- Path dependency that surfaces only after years of underperformance
- Governance frameworks that reward comfort over coherence

This paper reframes FX hedging as two distinct problems, not one. The first is hedging indefinite, evolving exposures, where flexibility matters more than precision. The second is hedging genuinely long-dated exposures, where certainty matters more than liquidity. Each requires a different structure.

The objective is not to eliminate FX risk. It is to match the hedge to the nature of the exposure. Most portfolios do not.

1. FX Hedging Is Not One Problem

FX exposure is often discussed as if it were homogeneous. It is not. Some FX exposures are finite. Others are indefinite. Some terminate naturally. Others persist for as long as the portfolio exists. Treating them all the same is the root cause of most FX hedging failure.

1.1 The Core Distinction

Consider two positions. Position A is a five-year USD-denominated corporate bond held by an Australian super fund. The bond matures on a known date. Principal and coupons are contractually defined. The FX exposure has a natural endpoint. Position B is a strategic allocation to US equities with no defined exit date. The position will be rebalanced, assets will turn over, managers may change, but the offshore exposure persists indefinitely.

Both create FX risk. But they are fundamentally different problems. Position A can be hedged to a known terminus. A five-year cross-currency swap matches the exposure precisely. There is no structural mismatch. Position B cannot. Any fixed-tenor hedge creates a decision point: what happens when it expires? If the exposure continues but the hedge does not, FX risk resurfaces. If the hedge is rolled, path dependency returns.

Yet in practice, both are often hedged using the same structure: rolling short-dated FX forwards. This is convenient. It is also structurally incoherent.

1.2 Why the Distinction Matters

FX risk is path dependent. What matters is not just where the currency ends up, but how it gets there. Short-dated hedges reset that path repeatedly. Long-dated hedges fix it. Options reshape it. If you do not distinguish between these objectives, you end up optimising for the wrong thing, usually liquidity and optics rather than economics.

1.3 The Exposure Classification Framework

Before selecting any instrument, classify the exposure. The first step in any FX hedging programme is not instrument selection. It is exposure classification. Get this wrong, and everything downstream fails.

Exposure Type	Characteristics	Natural Hedge Structure
Finite, Certain	Known maturity, fixed cashflows, stable size	Tenor-matched CCS or long-dated forwards
Finite, Uncertain	Probable but not certain exit date	Rolling forwards with option overlay
Indefinite, Stable	Perpetual exposure, strategic allocation	Rolling forwards, potentially with options for path management
Indefinite, Volatile	Active trading, frequent rebalancing	Short-dated rolling forwards only

2. Why Rolling FX Forwards Dominate

Despite years of criticism, rolling short-dated FX forwards remain the default hedging tool for institutional portfolios. This is not accidental, nor is it evidence of collective incompetence. Rolling forwards dominate because they optimise for operational survival.

They offer deep liquidity in most currency pairs, tight bid-offer spreads at 1–3 month tenors, simple hedge accounting treatment, minimal termination risk, easy resizing as portfolios evolve, and familiar documentation and operational processes.

From a governance perspective, they are hard to argue against. They are familiar, auditable, and fit neatly into committee-driven risk frameworks. They do not require multi-year commitments that might outlive the CIO who approved them.

2.1 When Rolling Forwards Are Optimal

For portfolios with uncertain holding periods, active rebalancing, manager turnover, evolving mandates, or governance structures that cannot tolerate mark-to-market volatility, rolling forwards are robust. That robustness is often mistaken for optimality.

2.2 The Governance Trap

Rolling forwards survive because they make bad outcomes less visible. Carry drag accumulates slowly. Path dependency only surfaces after years. Each individual roll looks sensible. Accountability is diffused across time. A long-dated hedge that loses money is immediately visible. Forty rolling hedges that collectively underperform are not.

This creates an institutional bias toward short-dated structures that persists regardless of economic merit. Rolling forwards are excellent at reducing short-term FX volatility. They are poor at delivering long-term FX certainty. Conflating the two is where problems begin.

Rolling forwards survive because they are resilient. Not because they solve the right problem.

3. What Rolling Forwards Do

To evaluate FX hedging properly, we need to be precise about what rolling forwards deliver and what they do not.

3.1 Capabilities and Limitations

Rolling short-dated FX forwards reduce short-term FX volatility, smooth periodic reporting outcomes, limit drawdown amplification during acute FX shocks, provide flexibility to adjust hedge ratios as exposures change, and avoid large upfront commitment to long-dated structures.

They do not neutralise long-term FX risk, lock in base-currency returns, remove path dependency, control cumulative carry drag, or guarantee outcomes over multi-year horizons. Each roll resets the hedge at prevailing spot and forward points. Over time, this introduces compounding uncertainty, not certainty.

3.2 The Compounding Effect: A Worked Example

Consider an Australian super fund with a persistent USD 100 million equity exposure, hedged using rolling 3-month FX forwards over 10 years. Assume AUD/USD spot starts at 0.6500, the interest rate differential averages 2.5% per annum with USD rates higher, the fund rolls the hedge 40 times over the decade, and each roll incurs 5 basis points in transaction costs.

Year 1: forward points cost 250 bps. Transaction costs: 20 bps. Total drag: 270 bps on the hedged exposure. Over 10 years, cumulative carry cost is approximately 27% of the exposure, cumulative transaction costs are 200 bps, and total economic cost is roughly 29% before any FX movement.

Compare to a 10-year cross-currency swap executed at inception. The carry cost is embedded in the fixed swap spread at approximately 25–28% over the period, transaction costs are executed once at 15–20 bps, and mark-to-market volatility is high but economically irrelevant if held to maturity. The CCS looks more expensive on day one. Over the full horizon, it is structurally cheaper, but only if the exposure persists.

The "rolling forwards are cheaper" narrative survives because the comparison is always point-in-time, never lifecycle.

3.3 Path Dependency in Practice

Path dependency means your outcome depends on the route taken, not just the destination. Take USD appreciating 20% over 10 years in a straight line versus USD appreciating 20% but with a 30% spike in year 5 before settling.

With rolling forwards, in the linear case you reset at moderately unfavourable rates throughout. In the volatile case, you reset at extremely unfavourable rates during the spike, locking in losses, then reset again at better rates later but the damage is done. With a long-dated hedge, path does not matter. The rate is fixed from day one.

Rolling forwards remove your ability to ignore the path. If you have strong views on long-run FX but no view on the path, rolling forwards systematically work against you.

3.4 Quarterly Smoothing vs Long-Term Divergence

A balanced fund with 30% offshore equities and a 50% hedge ratio hedged using rolling 3-month forwards illustrates the divergence clearly. In year 1, quarterly FX volatility drops from 8.2% unhedged to 4.1% hedged. Reporting outcomes are smoothed and governance is satisfied. Over 10 years, cumulative carry drag is -18% on the hedged portion, while the unhedged portion gains 5% from FX appreciation. The net result is the hedged portfolio underperforming by 11.5% versus unhedged, despite identical asset returns.

The hedge worked every quarter. The outcome still disappointed. If the mandate is to stabilise quarterly returns and prevent extreme drawdowns, rolling forwards do exactly what they are meant to do. If the mandate is to deliver currency-neutral returns over multi-year horizons, they never will.

Rolling forwards shape volatility. They do not eliminate FX exposure.

4. Hedge Ratios and the Myth of Neutrality

There is no universally correct FX hedge ratio. Observed institutional hedge ratios span a wide range: 0–30% for growth-oriented portfolios with long horizons; 30–60% for balanced mandates seeking moderate volatility reduction; 60–80% for liability-aware investors matching near-term cashflows; and 100% almost exclusively for governance or regulatory reasons.

4.1 The Fundamental Asymmetry

Here is the core problem that most hedge ratio discussions ignore. The expected long-run FX return is approximately zero. The expected hedge cost is not zero. This asymmetry is not a market inefficiency. It is structural.

Currency markets exhibit no persistent directional edge. Over long periods, spot rates oscillate around purchasing power parity with no systematic drift. The expected FX return of an unhedged position is, on average, zero. But FX hedges are not free. They crystallise interest rate differentials (carry), bid-offer spreads, roll costs, and funding costs in stressed markets. These costs are real, systematic, and one-directional.

A portfolio can be fully hedged and still experience significant FX-driven underperformance relative to expectations. When this happens, the hedge is blamed for not working, even though it did exactly what it was designed to do. The hedge eliminated FX volatility. It also eliminated the possibility of FX gains that might have offset carry drag.

4.2 Why 100% Hedge Ratios Persist

Full hedging survives because of governance simplicity ("we hedge all FX risk" is easier to explain than a calibrated ratio), accountability avoidance (if the ratio is 100%, FX outcomes cannot be blamed on CIO discretion), regulatory pressure, and the widespread misunderstanding that 100% hedged equals 0% FX impact, which is wrong.

4.3 Optimal Hedge Ratios Are Context-Dependent

The appropriate hedge ratio depends on time horizon (longer favours a lower optimal ratio due to carry drag), liability structure, governance tolerance for FX volatility, carry environment, and base currency. There is no single answer. But there are wrong answers, and they usually involve ignoring carry entirely.

5. Carry Can be the Dominant Driver

Most FX hedging discussions fixate on volatility. This is the wrong variable. Over long horizons, carry can dominate outcomes.

5.1 Why Carry Matters More Than Volatility

Currencies do not offer persistent directional returns. Interest rate differentials do. Forward points compound relentlessly, whether anyone is paying attention. Volatility is episodic. Carry is systematic. A 3% interest rate differential costs 3% per year, every year, regardless of whether spot moves or not. Over 10 years, that is 30% of the exposure, more in compounding terms.

5.2 The Base Currency Effect

The same hedging strategy produces radically different outcomes depending on your base currency. Consider an Australian super fund versus a Japanese pension fund, both hedging USD equity exposure. For the Australian fund (AUD base), USD rates are typically higher, forward points favour USD (discount to spot), hedging costs money every roll, and the 10-year cumulative drag runs between -20% and -30%. For the Japanese fund (JPY base), USD rates are significantly higher, forward points heavily favour USD, hedging earns positive carry, and the 10-year cumulative gain runs between +15% and +25%.

Same exposure. Same hedge structure. Opposite economic outcome. This is why you cannot import hedge ratios or hedging strategies from overseas peers without adjusting for your base currency environment.

5.3 Quantifying the Effect

Take an AUD investor hedging USD exposure with spot at AUD/USD 0.6500, USD 3-month rate at 5.5%, and AUD 3-month rate at 4.0%. The annualised forward points cost is -1.5%. Over 10 years at the current differential: year 1 carry cost is 1.5% of exposure, compounding over 10 years gives approximately 16% cumulative drag, add transaction costs of around 2%, and total structural cost is roughly 18%.

For a JPY investor hedging the same USD exposure with JPY rates at 0.25%, the annualised forward points gain is +5.25%. Year 1 carry gain is 5.25%, compounding over 10 years gives approximately 65% cumulative gain, net of transaction costs roughly 63%. The same USD equity portfolio, hedged identically, delivers an 81 percentage point difference in outcome purely from carry.

5.4 Why Carry Is Ignored

Carry is invisible in conventional risk reporting. VaR models focus on volatility, not systematic drift. Performance attribution isolates spot moves but buries forward points in "other". Hedge accounting treats carry as an offset, not a cost. Governance committees see quarterly hedge P&L, not cumulative lifecycle economics.

Ignoring carry does not make it disappear. It just makes the outcome surprising later.

6. Tenor Matching: When It Works, When It Fails

Tenor matching is conceptually appealing: hedge the exposure for as long as it exists. It works extremely well for instruments with defined cashflows and maturities. Bonds, loans, project finance, anything with a contractual endpoint. It fails when applied indiscriminately to exposures that do not terminate naturally.

6.1 When Tenor Matching Works Perfectly

An insurance company holds USD 200 million in investment-grade corporate bonds with average duration of 7 years and a liability matching requirement for AUD cashflows in 5–8 years. A 7-year cross-currency swap converts all USD fixed cashflows and principal to AUD at known rates: perfect tenor match. FX risk is eliminated entirely for the life of the assets. Carry is locked in at inception. No roll risk, no path dependency. Mark-to-market movements are economically irrelevant if held to maturity. This is tenor matching at its best. The exposure has a natural terminus. The hedge matches it.

6.2 When Tenor Matching Fails Badly

A pension fund holds AUD 400 million in global equities with an expected holding period described as long term (undefined), hedged at 70% using 10-year FX forwards locked at 0.6800. By year 5, a new CIO is appointed, and a strategic asset allocation review reduces global equities from 35% to 25%. The fund needs to reduce the hedge by AUD 100 million equivalent.

Spot has moved to 0.7200, meaning the hedge is deeply in-the-money. Unwinding crystallises large MTM gains, triggering tax consequences, performance distortion, CSA implications, and potential collateral posting requirements. The alternatives are equally bad: keep the hedge and be over-hedged (governance breach), pay termination costs and destroy value, or enter offsetting trades adding complexity, basis risk, and ongoing cost.

The hedge did exactly what it was supposed to do. The problem was upstream: applying tenor matching to an exposure that was not fixed.

6.3 The Hidden Risks of Long-Dated Hedges

Termination risk arises when exposure changes through manager decisions, mandate shifts, or asset sales, leaving a hedge that no longer fits. Mark-to-market volatility on a 10-year FX hedge can swing 20–40% in value. Even if economically irrelevant for a hold-to-maturity position, this creates balance sheet noise, audit questions, board nervousness, and pressure to act.

Credit and collateral consumption are also material. Long-dated hedges consume more credit lines and CSA capacity than rolling shorts. In stressed markets, counterparties may demand additional collateral, refuse to extend further hedges, or force early termination. Liquidity is also constrained outside core pairs: rolling 3-month EUR/AUD forwards are liquid with tight spreads, whereas 10-year EUR/AUD forwards carry wide spreads, limited capacity, and few counterparties.

If your exposure is...	Then tenor matching is...
Contractually defined with fixed maturity	Optimal
Strategically stable but no fixed exit	Risky unless governance can tolerate MTM and termination risk
Subject to active management	Wrong
Uncertain or evolving	Dangerous

Tenor matching is not wrong. It is context-dependent.

7. FX Forwards vs Cross-Currency Swaps

Cross-currency swaps are often described as expensive, complex, or overkill. They are simply misapplied as often as they are underused.

7.1 What Cross-Currency Swaps Do

A cross-currency swap converts a foreign-currency asset into a synthetic domestic-currency asset for the life of the swap. It locks FX conversion at inception, fixes the funding spread for the entire tenor, aligns hedge and asset cashflows if structured correctly, eliminates FX path dependency, and removes the need for active hedge management. For genuinely long-dated, stable exposures, this is exactly what you want.

7.2 Why CCS Look Expensive

CCS appear expensive because they are compared to the wrong benchmark. You compare a 10-year CCS quote today against a 3-month FX forward quote today. The CCS looks expensive. But you are not buying one 3-month hedge. You are buying forty of them.

A 10-year cross-currency swap executed at inception has an all-in cost locked for 10 years, a transaction cost of 15–20 bps paid once, minimal operational cost (set and forget), and carry embedded in the swap spread. Rolling 3-month FX forwards for 10 years have an unknown all-in cost depending on forward points at each roll, cumulative transaction costs of 200 bps (5 bps multiplied by 40 rolls), ongoing operational cost for execution, confirmation, settlement, and accounting, and carry paid continuously compounding over time.

Numerical example for an AUD investor hedging a USD 100 million bond portfolio: the 10-year CCS costs 500 bps over the decade plus 20 bps transaction cost for an all-in of 520 bps. Rolling 3-month forwards cost 1,500 bps in average carry, 200 bps in transaction costs, and 100 bps in conservative re-strike drag for an all-in of 1,800 bps. The CCS is one-third the cost. But it looked more expensive on day one because all the costs were visible and upfront.

CCS concentrate cost upfront. Rolling forwards distribute it invisibly over time. When you annualise the full economics, CCS are often cheaper.

7.3 When CCS Are the Right Tool

Cross-currency swaps suit stable, long-dated exposures with clear intent to hold corporate bonds with known maturity, infrastructure project cashflows, strategic offshore property holdings, multi-year loan commitments. They require meaningful size (minimum USD 50–100 million to justify documentation and pricing), and tolerance for mark-to-market noise, meaning governance must understand that MTM volatility does not equal economic loss if held to maturity.

7.4 When CCS Are the Wrong Tool

Avoid cross-currency swaps when holding period is uncertain (active equity allocation subject to SAA review, manager mandates that may change, assets that might be sold opportunistically), when exposure size is volatile (equities where market movements change notional, actively managed portfolios with high turnover), when governance cannot tolerate MTM volatility, or when currency pair liquidity is limited (exotic or emerging market currencies, pairs with limited long-dated market depth).

CCS also require ISDA Master Agreement with Credit Support Annex, collateral posting arrangements, CSA valuation and margin call processes, systems to handle swap cashflows and valuations, and regular mark-to-market reporting. This is material operational overhead. For a large fund hedging a stable bond portfolio it is worthwhile. For a smaller fund with volatile offshore exposure, it probably is not.

CCS are not superior. They are specific.

8. Options: The Missing Dimension

Forwards and CCS are linear instruments applied to non-linear portfolios. This creates a fundamental mismatch. Options introduce convexity. They fundamentally change the payoff profile of your FX overlay, providing asymmetric protection against tail events, flexibility to time carry drag, reduction in forced re-hedging, and breaking of pro-cyclical dynamics. They are not a replacement for forwards or CCS. They are a control layer.

8.1 What Options Do in FX Overlays

Tail risk protection: a put option protects against extreme adverse FX moves while preserving upside if FX moves favourably. This is particularly valuable for portfolios that cannot tolerate drawdown amplification from FX shocks, exposures where governance requires protection but not complete hedging, and situations where carry drag makes full forwards economically painful.

Timing flexibility around carry: options allow you to defer the decision about whether to hedge. A 1-year put option with a 5% out-of-the-money strike means you are protected if FX deteriorates more than 5%, you benefit if FX improves, you pay only a known premium upfront, and you avoid 12 months of negative carry from rolling forwards. This is particularly useful when carry is expensive, but long-run FX direction is uncertain.

Reduction of forced re-hedging: rolling forwards create forced decision points every quarter. If your hedge expires and FX has moved against you, you must re-strike at unfavourable levels. Options eliminate this. The protection remains regardless of spot movements.

Breaking pro-cyclicality: during FX stress, linear hedges force bad behaviour. Forwards expiring during volatility spikes require re-striking at terrible levels. Margin calls on CCS force collateral posting at the worst time. Governance panic leads to over-hedging after drawdowns. Options break this cycle. Protection is already in place. No forced action required during stress.

8.2 Option Structures for FX Overlays


Vanilla puts

Buy 1-year AUD puts (USD calls) at 5-10% out-of-the-money. Cost is typically 1.5-3% of notional varying with volatility and strike. Use case: downside protection for strategic offshore allocations, an alternative to 100% forward hedge when carry is expensive, governance-friendly insurance framing. Example: portfolio of AUD 100 million in USD equities, concern about potential AUD weakness without wanting to pay negative carry. Buy AUD puts struck at 0.6200 (5% OTM) for 2% cost. Maximum loss is 7% from FX (5% plus 2% premium), unlimited upside.

Vanilla put

Buy AUD put at 0.6200 (5% OTM) for 2% premium. Spot = 0.6500.

Break-even (0.6370)



Vanilla put – the copper line flattens above 0.6200 at -2% (the premium cost), which correctly shows that upside participation is retained but you always pay the premium. The unhedged line and the hedged line converge as AUD strengthens, making the cost of protection visible.


Collars (zero-cost structures)

Buy put and sell call at equidistant strikes. Cost is zero or near-zero premium. Use case: portfolios that want protection but cannot justify premium expense, substituting for partial forward hedges, accepting range-bound outcomes rather than paying carry. Example: buy AUD put at 0.6200, sell AUD call at 0.6800, net premium approximately zero. Protected below 0.6200, capped above 0.6800, no ongoing carry drag between strikes. This is economically superior to a 50% rolling forward hedge if you believe spot will stay within a range.

Collar (zero-cost)

Buy put at 0.6200, sell call at 0.6800. Net premium $\approx$ zero. Spot = 0.6500.

P&L Protection range (0.6200–0.6800)



Zero Cost Collar — the characteristic flat band between 0.6200 and 0.6800 is clear: no P&L drag in that range, full protection below the put strike, gain capped above the call strike. The zero-cost nature shows up clearly.


Seagull structures

Buy put, sell call (OTM), and sell far OTM put. Cost is reduced premium or zero. Use case: need some protection but premium budget is tight, comfortable accepting risk beyond extreme tail moves. Example: buy AUD put at 0.6300 (cost 2.5%), sell AUD call at 0.7000 (receive 1.0%), sell AUD put at 0.5500 (receive 0.5%). Net cost 1.0%. Protected between 0.5500 and 0.6300, naked below 0.5500, capped above 0.7000. If catastrophic moves below 0.5500 are unacceptable, this structure is wrong. If they are tolerable, option costs have been cut by 60%.

Seagull structure

Buy put 0.6300, sell call 0.7000, sell put 0.5500.

Net cost 1.0%. Spot = 0.6500.



Seagull — the most complex shape: protected slope between 0.5500 and 0.6300, flat cap above 0.7000, and the line diving again below 0.5500 where the short put leaves the portfolio naked. The danger zone below 0.5500 is unmistakable.


Knock-out options

Buy put with knock-out barrier above current spot. Cost is 30–50% cheaper than a vanilla put. Use case: seeking tail protection only, comfortable with option terminating if FX improves temporarily. Example: buy AUD put at 0.6200 with knock-out barrier at 0.7000. If spot ever touches 0.7000, option terminates. Premium 1.0% versus 2.5% for the vanilla put. Risk: if AUD spikes to 0.7000 temporarily and then falls to 0.5800, the position is unprotected.

Knock-out put

Buy put at 0.6200 with knock-out barrier at 0.7000.

Premium 1.0% vs 2.5% vanilla. Spot = 0.6500.



Knock-out put — two scenarios shown: the copper line (barrier intact, put active) and the platinum dashed line (barrier breached at 0.7000, position reverts to unhedged). The fade of the copper line above 0.70 signals that territory is where the option ceases to exist.

8.3 Strike and Tenor Selection

Strike selection depends on objective. At-the-money provides maximum protection at maximum cost and is rarely justified. A 5% OTM strike gives balanced protection and cost and is most common for strategic hedging. A 10–15% OTM strike covers tail risk only and is appropriate when governance requires something, but carry is expensive. For tenor, 3–6 months suits tactical protection during known volatility events, 1 year is standard for strategic overlays balancing cost and coverage, and 2-plus years is usually too expensive unless volatility is suppressed and tail risk conviction is high.

8.4 Why Options Are Underused

The objection that options are expensive collapses on examination. A 2% option premium for 1-year protection costs less than 2 years of negative carry on a 100% forward hedge. The objection that options expire worthless is precisely the point: you are paying for protection you hope not to need. The objection that options are complex does not hold either: vanilla puts are simpler than rolling quarterly forwards with dynamic hedge ratio adjustments. The objection that options require market timing is no truer than deciding hedge ratios or roll timing for forwards. Options can be mechanically rolled at 6–12 month intervals without any directional view.

8.5 When Options Make the Most Sense

Options are particularly valuable when carry drag from forwards is large (200 or more bps annually), governance requires downside protection but not certainty, the portfolio is prone to pro-cyclical behaviour, the exposure is indefinite but not suitable for long-dated forwards, or regulatory hedge ratios create forced trading at bad times.

Most portfolios avoid options not because they are unsuitable, but because they require decision-making. That discomfort is often mislabelled as prudence.

9. Mandatory FX Hedging Regimes

When hedging is compulsory, economics take a back seat. Some portfolios operate under regulatory or policy mandates that require minimum FX hedge ratios regardless of cost or market conditions. Common examples include insurance prudential standards (APRA CPG 230 in Australia), pension fund regulations with defined hedge ratio minimums, investment mandates with explicit FX hedging requirements, and sovereign wealth fund policies that prohibit FX discretion. In these environments, the objective shifts from optimisation to damage control.

9.1 The Governance vs Economics Tension

Mandatory hedging creates a structural problem. Policy requires an 80% hedge ratio. Economics says the current carry drag of -250 bps per annum suggests a 40% ratio. The institution hedges 80% and underperforms by approximately 100 bps per annum indefinitely. Changing the policy is difficult, requiring board approval, regulatory discussion, and often multi-year review cycles. What can be changed is how the mandate is implemented.

9.2 Strategies Within Mandatory Regimes

Using options to reduce pro-cyclicality: if a fund must maintain an 80% hedge ratio, using options instead of forwards for a portion (for example, 60% forwards and 20% put options at 5% OTM) still satisfies the policy while avoiding forced re-striking during volatility spikes, reducing margin call pressure, preserving some FX upside, and lowering cumulative transaction costs.

Extending tenor where possible: if policy specifies hedge ratio but not tenor, switching from 3-month to 12-month forwards reduces roll frequency from 4 times to once per year, cuts transaction costs by 75%, and reduces operational burden and re-strike risk.

Collar structures: some governance frameworks accept zero-cost collars as satisfying hedge ratio requirements. Policy reads as 80% hedged, implementation covers 80% with collars (puts bought, calls sold), governance is satisfied, and economics are materially better than 80% forwards because there is no carry drag within the collar bounds.

Advocating for dynamic hedge ratios: push for policy amendments that allow hedge ratios to vary with carry environment, volatility regime, and time horizon. This requires education and persistence but can save millions annually.

9.3 MAS and APRA Considerations

MAS Notice 124 requires Singapore insurers to apply appropriate risk management for FX exposures but does not mandate specific hedge ratios. Capital treatment under RBC2 creates indirect incentives to hedge, as unhedged FX exposure consumes capital. The strategy is to focus on capital-efficient hedging using longer-dated instruments and options to reduce VaR without paying excessive carry.

APRA CPG 230 requires Australian insurers to identify, measure, and manage FX risk but does not prescribe hedging levels. Internal policies often interpret this as requiring high hedge ratios. The strategy is to ensure internal policies explicitly permit hedge ratio

flexibility based on economic conditions and to document the trade-offs clearly for the board.

Mandatory regimes often require hedge effectiveness testing under accounting standards (AASB 9, IFRS 9). Options can fail effectiveness tests more easily than forwards because their value behaves non-linearly, creating a perverse outcome where forwards pass effectiveness tests even when economically damaging, and options may fail even when economically superior. Structuring options carefully to maintain 80% or higher correlation with hedged items addresses this.

9.4 Documenting the Trade-Offs

When governance forces suboptimal hedging, the responsibility is to implement the policy as instructed, document the economic cost explicitly, quantify alternatives annually, and advocate for policy evolution. Example memo language: current policy requires 80% FX hedging. At prevailing interest differentials (USD +250 bps versus AUD), this costs approximately 200 bps annually on the hedged portion. Alternative structures (60% forwards plus 20% options) would satisfy the policy intent while reducing annual cost by approximately AUD 1.5 million.

In mandatory hedging regimes, options become more valuable, not less. They are often the only remaining lever to reduce pro-cyclicality and manage carry without breaching policy.

10. A Unified FX Overlay Framework

A coherent FX overlay recognises three layers, not one. Most portfolios run only the middle layer, then wonder why outcomes disappoint.

10.1 The Three-Layer Framework

Layer 1: Structural FX (Certainty Layer)

Purpose: Lock in FX outcomes for genuinely long-dated, stable exposures

Instruments: Cross-currency swaps, long-dated forwards

Exposure types: Fixed-maturity bonds, project finance, infrastructure cashflows

Typical allocation: 10–25% of offshore exposure

Layer 2: Flexible FX (Volatility Management Layer)

Purpose: Reduce short-term volatility for indefinite exposures

Instruments: Rolling forwards (3–12 month tenors)

Exposure types: Strategic equity allocations, alternative investments, undefined holding periods

Typical allocation: 50–70% of offshore exposure

Layer 3: Convexity and Control (Protection Layer)

Purpose: Tail risk protection, carry management, pro-cyclicality reduction

Instruments: Options: puts, collars, seagulls

Exposure types: Any exposure where downside matters more than average outcome

Typical allocation: 10–30% of offshore exposure (can overlap with Layer 2)

10.2 How the Layers Interact: An Illustrative Example

Total offshore exposure of USD 500 million. Layer 1 covers USD 150 million (30%) via 7-year CCS to lock returns on the bond portfolio. Layer 2 covers USD 250 million (50%) via 6-month rolling forwards to reduce equity volatility. Layer 3 covers USD 100 million (20%) via 1-year put options for tail protection.

Effective hedge ratio is 80% (Layers 1 and 2 combined) for day-to-day volatility, with 100% protection against extreme moves through Layer 3. The bond portfolio is FX neutral. The equity portfolio has volatility smoothed with tail protection. Carry drag is minimised on the 20% with options-only coverage. Operational complexity is manageable with two instruments and clear allocation.

10.3 Determining Layer Allocation

Four questions drive the allocation. What portion of your FX exposure has contractual maturity dates? That is your Layer 1 allocation. What portion of the remaining exposure requires quarterly reporting smoothness? That is your Layer 2 allocation. What portion cannot tolerate adverse FX moves greater than 15%? That is your Layer 3 allocation, which can overlap with Layer 2.

Finally, what is your base currency carry environment? Large negative carry favours increasing Layer 3 and reducing Layer 2. Large positive carry favours increasing Layer 2 and reducing Layer 3. Neutral carry supports a balanced allocation.

10.4 Portfolio Sizing Examples

Orientation	Layer 1 (Structural)	Layer 2 (Flexible)	Layer 3 (Convexity)	Effective Hedge Equivalent
Conservative	30%	50%	20% (overlaps L2)	~80% hedged
Balanced	20%	40%	30% (partial overlap)	~60% hedged
Growth-oriented	10%	20%	30% (tail risk only)	~40% hedged

The key is intentional mismatch between layers based on exposure characteristics, not applying one structure to everything.

11. Implementation

11.1 Implementation Sequence

Phase 1 (Month 1): Classify exposures. Map all offshore holdings to maturity profile, separate contractual endpoints from indefinite exposures, and quantify current hedge ratio by exposure type.

Phase 2 (Months 2–3): Implement Layer 1. Execute CCS for bond portfolios and other fixed-maturity exposures, document hold-to-maturity intent for accounting, and set up CSA processes if not already in place.

Phase 3 (Months 3–6): Optimise Layer 2. Extend roll tenor from 3-month to 6–12 month where possible, implement a dynamic hedge ratio policy (requires governance approval), and establish operational processes for quarterly reviews.

Phase 4 (Months 6–12): Introduce Layer 3. Start with simple vanilla puts at 5% OTM on 12-month tenors, establish option pricing and execution relationships, educate governance on option payoffs and accounting treatment, and consider zero-cost collars if premium budget is constrained.

Phase 5 (Ongoing): Annual review of Layer 1 structures (nothing should change unless assets mature). Quarterly review of Layer 2 hedge ratios (adjust for carry and volatility). Annual roll of Layer 3 options (reassess strikes and notionals).

11.2 Common Mistakes in Implementation

Mistake	Symptom	Fix
Running only Layer 2 for everything	Chronic carry drag, governance frustrated	Introduce Layer 1 for bonds and Layer 3 for tail risk
Over-allocating to Layer 1 based on strategic intent	Large CCS positions on equity portfolios, forced unwinds after SAA reviews	Reserve Layer 1 for contractual maturities only
Treating Layer 3 as optional or nice-to-have	Carry drag erodes returns, governance resists options as expensive	Compare option premium to cumulative carry drag over 5–10 years: options are often cheaper
Static hedge ratios across all layers	80% hedged in all conditions regardless of carry	Layer 2 should flex 40–80% based on carry; Layer 3 provides tail protection when Layer 2 is reduced

11.3 Troubleshooting Guide

Problem	Likely Cause	Fix
Persistent underperformance despite being fully hedged	Ignoring carry drag from rolling forwards	Reduce hedge ratio or introduce options to reduce carry bleed
Large hedge P&L swings creating governance panic	Long-dated hedges on uncertain exposures	Shorten tenor or use collars instead
Forced to re-hedge at terrible levels after FX spike	Rolling forwards expiring during volatility	Introduce put options for tail protection
Hedge ratio breaches during market stress	Static hedge ratios with volatile equity exposure	Move to dynamic hedge ratios or use options
Audit committee questions why hedge lost money	MTM volatility confused with economic loss	Improve governance education on hedge accounting vs economic hedging
Cumulative transaction costs exceeding 2% over 5 years	Rolling 3-month forwards too frequently	Extend roll tenor to 6-12 months
Unable to execute desired hedge size	Illiquid currency pair with long tenor	Use rolling forwards plus options instead of CCS

12. Case Study: Misaligned vs Aligned Hedging

Two Australian super funds with identical offshore exposures illustrate the economics of misaligned versus aligned hedging.

12.1 The Setup

Both funds have total FEV of AUD 2 billion, an offshore allocation of 35% (AUD 700 million) consisting of USD 300 million equities (strategic), USD 150 million bonds (5–7 year maturity), and USD 100 million alternatives (10-plus year lockup), an AUD base currency, and a policy hedge ratio of 70%. The carry environment has USD rates at 5.5% and AUD rates at 4.0%, with forward points cost of approximately 150 bps per annum.

12.2 Fund A: Conventional Approach (Misaligned)

Fund A hedges 70% of the entire offshore exposure with rolling 3-month FX forwards, applies the same structure to equities, bonds, and alternatives, fixes the hedge ratio at 70% regardless of market conditions, and uses no options.

Year	Spot Move	Hedge P&L	Carry Cost	Transaction Costs	Net FX Impact
1	–5%	+3.5%	–1.5%	–0.2%	+1.8%
2	+8%	–5.6%	–1.5%	–0.2%	–7.3%
3	+2%	–1.4%	–1.5%	–0.2%	–3.1%
4	–3%	+2.1%	–1.5%	–0.2%	+0.4%
5	+6%	–4.2%	–1.5%	–0.2%	–5.9%

Cumulative 5-year result: spot moved +8%, hedge P&L was –5.6%, carry drag was –7.5%, transaction costs were –1.0%, and net FX impact was –14.1%. The hedge did exactly what it was supposed to do, reducing the spot impact from +8% to approximately +3%. But carry drag and transaction costs quietly consumed –11.5%, resulting in net underperformance of –14.1% versus unhedged. Governance never questioned it because each quarterly hedge looked fine in isolation.

2.3 Fund B: Aligned Approach (Three-Layer Framework)

Fund B implements Layer 1: USD 150 million bonds hedged at 100% via 7-year CCS, locked at inception with no ongoing management. Layer 2: USD 200 million equities hedged at 50% via 6-month rolling forwards, with the hedge ratio flexing between 40–60% based on carry conditions. Layer 3: USD 250 million combined exposure protected at 20% via 1-year put options at 5% OTM, rolled annually. Effective hedge ratio is 68%, comparable to Fund A for governance purposes.

Year	Spot Move	Layer 1 (CCS)	Layer 2 (Forwards)	Layer 3 (Options)	Net FX Impact
1	–5%	+5.0%	+1.5%	+0.5%	+7.0%
2	+8%	–8.0%	–1.5%	–1.8% (premium)	–11.3%
3	+2%	–2.0%	–1.2%	–1.8%	–5.0%
4	–3%	+3.0%	+0.8%	+0.2%	+4.0%
5	+6%	–6.0%	–1.8%	–1.8%	–9.6%

Cumulative 5-year result: spot moved +8%, Layer 1 (CCS) contributed –8.0% (locked, no surprises), Layer 2 (forwards) contributed –1.0% (lower notional, longer rolls), Layer 3 (options) cost –4.7% net premium over 5 years, and net FX impact was –5.7%.

Fund B achieved a similar hedge ratio to Fund A (68% versus 70%) with dramatically better economics. Layer 1 eliminated path dependency on bonds. Layer 2 reduced carry drag by hedging less notional with longer tenor. Layer 3 cost premium but provided asymmetric protection during the Year 4 drawdown.

Net difference: Fund B outperformed Fund A by 8.4% over 5 years purely from hedge structure, with identical underlying assets.

12.4 The Governance Conversation

Fund A board meeting, year 5: "We are fully hedged as per policy, but we have underperformed the unhedged benchmark by 14%. Why is the hedge not working?" CIO response: "The hedge is working as designed. It is reducing volatility. The underperformance is from carry costs." Board: "We never agreed to pay 14% for hedging." Reality: they did. They just did not realise it.

Fund B board meeting, year 5: "We are hedged at 68% as per policy. We have underperformed the unhedged benchmark by 5.7%, which is consistent with our carry drag analysis presented at inception." Board: "Understood. Is there any way to reduce that drag?" CIO: "Yes. We could reduce Layer 2 allocation slightly and increase Layer 3 options. Cost of approximately 2% per annum in premium, reduction in carry drag of approximately 3% per annum. Net benefit 1% per annum improvement." Board: "Approved for next year."

The difference is not intelligence. It is intentional design.

Conclusion

FX hedging fails not because portfolios choose the wrong instruments, but because they expect one instrument to solve every FX problem.

Match the hedge to the exposure. Finite exposures deserve term structures. Indefinite exposures require flexibility. Match the structure to the intent: if you want certainty, use CCS; if you want flexibility, use forwards; if you want convexity, use options. Match governance to reality: full hedging is not neutral, carry is not optional, and path dependency is not a nuisance. It is the thing you are managing.

Rolling forwards are excellent tools. They are not universal solutions. Cross-currency swaps are not expensive. They are expensive to unwind early. Options are not complex. They are decision-forcing.

The three-layer framework works because it recognises that portfolios contain multiple FX problems requiring multiple FX solutions. Most portfolios run one structure for everything, then blame the hedge when it underperforms.

The hedge did what it was designed to do. The design was wrong.

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/