PARA BELLUM ADVISORS

PRACTITIONER PAPER

Convexity in Portfolios

Cost, Budgeting, and When It Fails

By Mike Duncan, Para Bellum Advisors

Version 1, April 2026

www.parabellumadvisors.com

Paper 3 of 3: Convexity in Portfolios

Series: Options, Volatility, and Convexity in Institutional Portfolios

Executive Summary

Most institutional portfolios have hedges. Very few of these institutions have a convexity budget.

The difference is not semantic. A hedge is a position designed to offset a specific risk. A convexity budget is a structural portfolio decision: how much of the portfolio's return should be allocated to shaping the distribution of outcomes rather than maximising expected return. One is reactive. The other is deliberate.

This distinction matters because the two approaches produce very different outcomes in practice. An institution that hedges reactively buys protection after it becomes frightened, pays elevated volatility for the privilege, and cuts the programme when it underperforms for two consecutive quarters. An institution with a convexity budget deploys protection systematically when it is cheap, sizes it relative to a portfolio objective, and measures it against that objective rather than against the cost of premium in any given period.

Papers 1 and 2 in this series established the conceptual foundation. Options are bundles of risk exposures, not strategies. Volatility is an asset with its own pricing dynamics, term structure, and regime cycle. This paper closes the series by addressing the portfolio-level question: how do you build, size, and govern a convexity programme that actually works?

The central arguments are:

- Convexity is a legitimate portfolio allocation, not a cost to be minimised. It changes the shape of the return distribution in ways that have real value for institutions managing to specific constraints.
- Sizing a convexity budget requires a clear portfolio objective. Without one, the budget is arbitrary and ungovernable.
- Monetisation is where most programmes fail. An in-the-money hedge (especially a deep-in-the-money hedge) that is never converted into cash (i.e. realised) is not protection. It is narrative comfort.
- Governance is the binding constraint. The best-designed programme will fail if the institution cannot act at the speed that markets require.

This paper also addresses directly the question that boards and investment committees most commonly ask: how do we know whether the convexity budget is working?

That question has a specific answer, and it is not the P&L of the option portfolio in isolation.

> ***A hedge that works but is never monetised is worth the same as a hedge that was never bought. The instrument is not the protection. The realised cash is.***

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

1. Reframing Convexity as a Portfolio Allocation

The language used to describe option programmes shapes how they are governed, measured, and ultimately whether they survive long enough to do their job.

When option premium is described as a cost, the natural institutional response is to minimise it. Cheaper structures, shorter maturities, smaller notional. Each of those choices reduces the effectiveness of the programme in exactly the scenarios it was designed for.

When convexity is described as an allocation, the response changes. An allocation has a purpose. It has a benchmark. It has a framework for evaluating whether it is working. It does not get cut when it underperforms in a quiet market any more than a fixed income allocation gets cut because equities returned more in a given year.

1.1 What convexity does to a portfolio

Convexity changes the shape of the return distribution. A portfolio without convexity has returns that are approximately normally distributed around its expected return – symmetric, with both upside and downside roughly proportional to the portfolio's risk level.

Adding long convexity through options changes that shape. The left tail – large negative returns – is compressed. The cost of that compression shows up as a small but persistent drag on performance in calm periods. In stress periods, the convexity position accelerates in value while the portfolio is falling, creating a return profile that is asymmetric in the portfolio's favour.

That asymmetry has real institutional value, beyond the P&L of the option position itself. For a superannuation fund that cannot afford a 25% drawdown without triggering a contribution call, the value of convexity is not what it earns on its own; it is what it prevents the fund from having to do. For a private credit fund that cannot meet redemptions during a credit freeze, convexity that generates cash during that freeze is worth more than its mark-to-market gain at any point in time.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

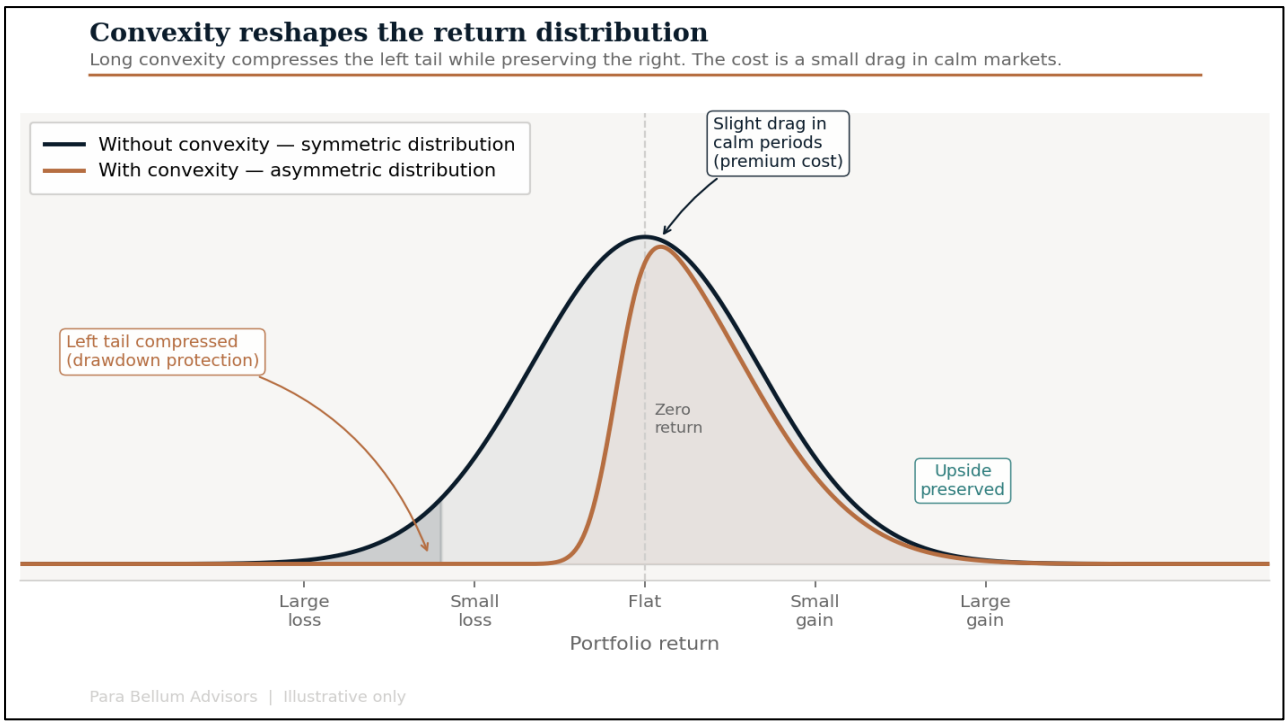

Figure 1: Convexity reshapes the return distribution. The left tail is compressed through drawdown protection. The right tail is preserved. The cost is a small but persistent drag in calm markets from option premium.

> ***The value of convexity is not measured by what it earns. It is measured by what it prevents and by the options it creates. That is a different calculation from option P&L, and most performance frameworks do not make it.***

1.2 The allocation decision

Treating convexity as an allocation requires three decisions that most institutions have not made explicitly.

1. The portfolio objective the convexity is serving. Protection against a specific drawdown threshold. Liquidity creation during a defined stress scenario. Governance cover for a board that needs to demonstrate tail risk management. Each objective implies a different structure, a different maturity, and a different sizing.

2. The return the portfolio is willing to give up in normal conditions to fund the allocation. A convexity budget that costs 80 basis points per year in calm markets will drag on performance in years when markets are quiet. The institution needs to have explicitly decided that drag is acceptable.

3. How the convexity gain will be used when it is realised. Without a pre-committed answer to this question, the convexity budget is incomplete regardless of how well the instruments are designed.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

2. Sizing the Convexity Budget

There is no universal right answer to how large a convexity budget should be. The right size depends on the portfolio's specific objectives, risk constraints, liability structure, and the vol environment at the time of deployment.

2.1 Starting from the portfolio objective

The sizing question begins with a specific scenario: what event does this budget need to defend against, and what does defence mean in that context?

For an equity-heavy portfolio with a drawdown tolerance of 15%, defence means generating enough convexity gain to keep net drawdown below 15% in a 25 to 30% equity market correction. That is a quantifiable target. It implies a specific notional of protection, which implies a premium budget, which defines the allocation.

For a private credit fund with quarterly redemption gates, defence means generating enough liquidity to meet redemptions without forced asset sales during a credit dislocation. For a family office with concentrated equity exposure, defence may mean generating enough cash during a correction to selectively add to core positions at dislocated prices. Each objective produces a different budget number.

2.2 The 2 to 5 percent range

In practice, institutional convexity allocations for portfolios with genuine tail risk management objectives typically sit between 2 and 5 percent of portfolio capital on an annual premium basis.

Allocation level	Profile	What it achieves
1 to 2%	Minimum meaningful allocation	Partial offset in severe stress; primarily governance cover; limited liquidity creation
2 to 3%	Standard institutional range	Material drawdown reduction in a significant stress event; meaningful liquidity creation; defensible to boards
3 to 4%	Higher conviction or constrained mandate	Significant tail protection; substantial liquidity in stress; visible impact on portfolio outcomes over a full cycle
4 to 5%+	Explicit crash protection mandate	Dominant impact in severe stress; high carry cost in calm conditions; requires strong governance to sustain

These numbers assume efficient deployment across the vol cycle. An institution that buys all of its protection at elevated volatility will get significantly less convexity per dollar of premium than one that deploys systematically when vol is cheap.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

2.3 Notional vs premium budgeting

A premium budget specifies how much to spend on option premium per year. In a low-vol environment, the premium budget buys more notional. In a high-vol environment, it buys less. That is the correct behaviour: buying more protection when it is cheap and less when it is expensive.

Notional budgeting forces the same notional regardless of cost – buying the most expensive protection at the worst time.

> ***Budget in premium, not notional as the primary budgeting rule. A premium budget that deploys more protection when vol is cheap and less when it is expensive is doing the right thing automatically.***

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

3. Instrument Selection and Programme Construction

Once the budget and objective are defined, the question becomes how to construct the programme. Instrument selection, maturity, strike, and rebalancing mechanics all affect how efficiently the budget translates into actual portfolio protection.

3.1 Matching instrument to objective

Different instruments provide different types of convexity with different cost profiles and liquidity characteristics. The choice of instrument should follow from the objective, not from familiarity or administrative convenience.

Instrument	Convexity type	Best suited for	Key risk
Equity index puts	Equity tail	General equity drawdown protection; governance cover	Skew cost; basis to active equity books
Put spreads	Defined-range equity	Cost-efficient protection within a defined drawdown range	Gap risk below the short strike
Variance swaps	Realised vol spike	Pure vol exposure; crisis correlation often better than puts	Mark-to-market volatility; complex governance
Payer swaptions	Rates tail	Rate spike protection; duration-sensitive portfolios	Premium cost; path dependency
Credit index CDS	Credit tail	Credit book protection; liquid in most conditions	Basis to single-name book; roll cost
Cross-asset overlays	Multi-asset tail	Diversified tail exposure; FX convexity in EM stress	Basis risk; less intuitive governance

Most institutional programmes are too concentrated in a single instrument type. An equity-only convexity programme protects against equity corrections but may leave the portfolio exposed to the scenarios – credit freezes, rates shocks, FX dislocations – that are most likely to produce genuine liquidity stress.

3.2 Strike selection

Deep out-of-the-money puts – struck at 70 to 80% of spot – are cheap in premium terms but require a severe correction before they generate significant payoff. Near-the-money puts – struck at 90 to 95% of spot – cost significantly more but provide protection from the first meaningful drawdown. Put spreads sit between the two extremes and are the most common institutional structure because they reduce premium cost materially while maintaining protection across the most likely drawdown range.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

3.3 Maturity and roll mechanics

Three-to-six month options are the institutional workhorse. They balance theta cost with convexity sensitivity, roll at a cadence that allows regular review, and are liquid enough across major equity indices to execute without significant market impact.

The roll decision – when and how to replace expiring options – is more important than most programmes acknowledge. Rolling mechanically at expiry means buying at whatever vol is prevailing on the roll date. Rolling into the best available maturity and structure based on current vol conditions produces materially better outcomes over time.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

4. The Governance Framework

The most technically well-designed convexity programme will fail if the governance structure cannot support it. Governance failure takes three forms: 1) the programme is cut during a quiet period; 2) it is maintained but not monetised when the stress event occurs; or 3) it is monetised at the wrong time because there was no pre-committed trigger framework.

4.1 The governance principles

Three principles govern a well-designed convexity programme.

1. the investment case must be decided in advance and separated from performance evaluation.
2. monetisation triggers must be pre-committed and delegated. The decision to monetise a convexity gain cannot wait for an investment committee meeting.
3. the programme must be evaluated against its objective, not against its P&L in isolation.

4.2 The speed problem

A credit event can develop from the first spread widening to peak stress in days. An equity correction of 15 to 20% can happen in a couple of days. A vol spike from 15 to 40 can occur in 48 hours.

An investment committee that meets monthly cannot respond at that speed. The solution is pre-commitment: defining in advance the conditions under which specific actions are taken and delegating the authority to take them.

Trigger type	Signal	Action
Portfolio drawdown	Net portfolio loss exceeds 8%	First monetisation tranche: realise 30% of hedge gains; proceeds to cash
Portfolio drawdown	Net portfolio loss exceeds 12%	Second tranche: realise additional 30%; evaluate reload options
Portfolio drawdown	Net portfolio loss exceeds 18%	Final tranche: realise remaining position; maintain residual hedge for second-leg risk
Hedge gain ratio	Hedge gains equal 50% of contemporaneous portfolio loss	First tranche regardless of drawdown level
Market structure	Vol term structure inverts and VIX above 35	Accelerate monetisation timeline; do not wait for drawdown triggers
Policy intervention	Emergency central bank action or exchange halt	Immediate monetisation of liquid positions; decay clock has started

4.3 Delegated authority

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

The CIO, or a designated portfolio manager, should have explicit delegated authority to execute monetisation tranches up to a defined size without committee approval, subject to immediate notification and next-meeting reporting. Without it, the programme will consistently fail at the moment it matters most.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

5. Monetisation: Where Most Programmes Actually Fail

Design is the easy part. Monetisation is where the programme either works or it does not.

A convexity position that moves deeply in the money during a market stress event has created value. But that value is unrealised. It sits on the portfolio as a mark-to-market gain that will decay as rapidly as markets stabilise, policy intervenes, or volatility mean-reverts. The window between peak convexity value and the beginning of that decay is often measured in days, sometimes hours.

> ***Marks do not fund margin calls, redemptions, or operating expenses. Only cash does. A hedge that is never monetised is narrative comfort, not capital protection.***

5.1 Why convexity decays after stress

When a market stress event is acute, implied volatility rises sharply and the convexity position generates large gains from both the intrinsic value of the options and from vega expansion. But three forces begin working against the position almost immediately: 1) policy intervention compresses vol rapidly; 2) mean reversion works against the position even without policy intervention; and 3) the underlying may recover before the institution has acted.

5.2 The monetisation doctrine

A monetisation doctrine is a pre-committed set of rules governing when, how much, and in what tranches a convexity gain is converted to cash. It removes the discretion that allows governance processes to create delays, and it replaces heroic timing decisions with procedural execution.

The core mechanics are straightforward. Monetise in tranches rather than all at once, because no one knows where the trough is and tranching reduces regret. Direct proceeds in a specific order: operating cash and near-term obligations first, then defensive reallocation, then selective re-risking into dislocated positions. Maintain a small residual hedge after each tranche. Reload the convexity programme at the vol levels prevailing after the event.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

5.3 What to do with the proceeds

Priority	Rationale
1. Operating cash and near-term obligations	Secure the institution's ability to function regardless of market conditions. Redemptions, margin calls, operating expenses.
2. Reduce leverage in stressed positions	Deleverage pro-cyclical exposures that amplify losses if stress continues. This is risk management, not market timing.
3. Defensive reallocation	Move to short-duration, high-quality assets that preserve capital and provide re-risking optionality.
4. Pre-vetted selective re-risking	Deploy capital into dislocated positions identified before the stress event. Requires a pre-approved opportunity list.
5. Reload convexity	Rebuild the protection programme at post-stress vol levels. The next event does not announce itself in advance.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

6. Measuring Whether the Programme Is Working

The performance question that most investment committees ask is the wrong one. 'Did the option portfolio make money?' is not the right measure of a convexity programme. The right question is: 'Did the portfolio meet its drawdown objective, and did the convexity programme contribute to that outcome?'

6.1 The correct performance metrics

Metric	What it measures	Benchmark
Net portfolio drawdown vs objective	Did the portfolio stay within its defined tolerance during stress events?	The institution's stated drawdown constraint, not peer comparison
Monetisation capture rate	What percentage of peak convexity value was converted to cash?	Internal target, typically 60 to 80% of peak mark
Cost per unit of protection	How much premium was spent per unit of drawdown reduction achieved?	Historical benchmarks; vol-adjusted comparison across periods

If the portfolio stayed within its drawdown constraint during a stress event, the programme worked, regardless of whether the option P&L was positive in isolation.

6.2 The carry cost conversation

The carry cost of a convexity programme is the price of maintaining a specific portfolio characteristic: the assurance that the portfolio will behave in a defined way during a stress event. The conversation becomes easier when the carry cost is stated as a percentage of the portfolio's drawdown tolerance rather than as an absolute premium figure.

6.3 Attribution and reporting

Convexity programme performance should be reported separately from the core portfolio, with attribution showing premium spend, option P&L, monetisation events and proceeds, and the net portfolio drawdown versus the objective for each period.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

7. Case Study: The Difference a Doctrine Makes

The following example is illustrative but reflects realistic dynamics observed across institutional portfolios during stress events. It shows the same portfolio, the same convexity instruments, and the same market conditions – with and without a monetisation doctrine.

7.1 The scenario

A balanced institutional portfolio of USD 1 billion holds a rolling equity protection programme covering 40% of the equity sleeve. The programme consists of three-month put spreads struck at 90 to 75% of spot, costing approximately 80 basis points per year of the equity sleeve. A stress event produces a 22% equity correction over six weeks, with implied vol moving from 14 to 48 at the peak.

At the trough, the put spreads are worth approximately USD 60 million in aggregate, against an equity book loss of USD 180 million – covering roughly 33% of the equity loss. Policy intervention arrives six days after the trough. Over the following two weeks, equity markets recover 8% and implied vol drops from 48 to 24.

7.2 Without a monetisation doctrine

The investment committee convenes ten days after the trough, by which time markets have partially recovered and vol has already compressed significantly. No decision is made at the first meeting. At the second meeting, three weeks after the trough, the equity recovery is 12% and vol is at 19. The put spreads are now worth 18 million dollars.

The committee monetises the remaining position, realising 18 million dollars. Total monetisation: 18 million dollars on a peak mark of 60 million dollars. Capture rate: 30%. The programme worked. The governance framework did not.

7.3 With a monetisation doctrine

Pre-committed triggers are in place. The first tranche executes automatically when the portfolio crosses the 8% drawdown threshold, twelve days into the selloff. The second tranche executes four days later. The market structure trigger fires the day after policy intervention. Total monetisation: 58 million dollars on a peak mark of 60 million dollars. Capture rate: 97%. Net portfolio drawdown: approximately 8% after hedge contribution.

> ***The instruments were identical. The market was identical. The difference was entirely in whether the institution had decided in advance what to do when the hedge worked.***

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

Side-by-side comparison: same instruments

Stage	Timing	Without doctrine	With doctrine	What made the difference
Portfolio drawdown crosses -8%	**Day 12**	No action. Committee not yet convened.	First tranche executed automatically. $18m to cash.	Pre-committed trigger. No committee required.
Drawdown reaches -12%	**Day 16**	No action. Committee meeting scheduled for following week.	Second tranche: $18m. Leverage reduced in two pro-cyclical positions.	Delegated authority. Executed same day.
Policy intervention. Vol drops 48 to 24.	**Day 22**	Committee convenes. Debate on second-leg risk. No consensus.	Market structure trigger fires. Balance monetised: $22m. Proceeds to pre-vetted positions.	Market structure trigger removed discretion. Action taken in hours, not days.
Second committee meeting	**Day 35**	Remaining position monetised at compressed vol: $18m. Vol at 19; peak was 48.	Programme reload underway at post-stress vol levels.	Window had closed. Most of the value had already been given back.
Total proceeds		$18m	$58m	**$40m difference. Identical instruments. Identical market.**
Net portfolio drawdown after hedge		~15%	~8%	**7 percentage points of additional protection delivered.**
Capture rate		30%	97%	**The difference was entirely in governance, not instruments.**

The table makes the point that is easy to miss when the two scenarios are described in prose: every dollar of difference in the outcome's column traces directly to a governance decision, not a market one.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

8. Conclusion: What a Complete Programme Looks Like

Across three papers, this series has built from first principles to a complete framework for how institutional investors should think about options, volatility, and convexity.

Paper 1 established that options are bundles of risk exposures, not strategies. Paper 2 developed the case that volatility is a standalone asset with its own pricing dynamics and regime cycle. This paper has closed the series with the portfolio-level framework.

A complete convexity programme has five components, each of which must be in place for the programme to work:

1. A clear portfolio objective that the convexity is serving. Not 'reduce risk' but a specific drawdown constraint, a specific liquidity requirement, or a specific governance obligation.
2. A premium budget calibrated to that objective and managed against the vol environment. Budget in premium, not notional. Deploy more when vol is cheap, less when it is expensive.
3. An instrument selection and construction process that matches the convexity type to the tail scenario being hedged. Ladder across instruments and maturities.
4. A pre-committed monetisation doctrine with delegated authority. Define triggers before the event, not during it. Tranche the monetisation. Direct proceeds in a specific order. Maintain a residual hedge. Reload after the event.
5. A performance framework that measures the programme against its objective rather than against option P&L in isolation. Report over a full market cycle. Address the carry cost conversation directly and proactively.

The institutions that get this right do not have more sophisticated instruments than those that get it wrong. They have better governance, clearer objectives, and the discipline to act on pre-committed plans when markets are moving and the instinct is to wait.

Most institutional portfolios have hedges. Very few have a complete convexity programme. The gap between the two is not technical. It is structural.

> ***The question is not simply whether to own convexity. It is whether the institution is prepared to convert it into something useful when it matters. That preparation happens before the crisis, not during it.***

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.

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.

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/

© Para Bellum Advisors, 2026. All rights reserved. | Informational only – not investment advice.