



PARA BELLUM ADVISORS

PRACTITIONER PAPER

Volatility as an Asset

Buying, Selling, and Mispricing It

By Mike Duncan, Para Bellum Advisors

Version 1, April 2026

www.parabellumadvisors.com

Paper 2 of 3: Volatility as an Asset

Series: Options, Volatility, and Convexity in Institutional Portfolios

Executive Summary	3
1. What Volatility Actually Is	4
1.1 The volatility risk premium	4
1.2 Realised vs implied: the gap is the trade	4
2. The Term Structure of Volatility	6
2.1 What the term structure tells you	6
2.2 Short-dated vs long-dated vol: different animals	7
2.3 Term structure trades	7
3. Skew: Structural Fact, Not Mispricing	9
3.1 What skew is	9
3.2 Why skew exists	9
3.3 What skew tells you about the cost of protection	9
3.4 Skew strategies	10
4. Volatility Regimes	11
4.1 The four regimes	11
4.2 Vol as a mean-reverting asset	11
4.3 VIX and the limits of headline vol metrics	12
5. Evaluating Whether Vol Is Mispriced	13
5.1 The starting point: implied vs realised	13
5.2 Vol percentile analysis	13
5.3 The term structure as a signal	13
5.4 Cross-asset vol signals	13
6. How Institutions Get This Wrong	15
6.1 Treating vol as a cost rather than an asset	15
6.2 Reactive buying and the timing trap	15
6.3 Short-dating the portfolio	15
6.4 Selling vol without owning the tail	15
6.5 Ignoring cross-asset vol	16
7. Practical Framework for Volatility Evaluation	17
7.1 The five questions	17
7.2 Sizing to the vol environment	17
7.3 The role of vol in a convexity budget	18
8. Conclusion	19

This paper is part of a three-paper series: Options, Volatility, and Convexity in Institutional Portfolios.

Para Bellum Advisors | www.parabellumadvisors.com

Executive Summary

Volatility is not just a risk metric. It is a tradeable asset with its own supply and demand dynamics, term structure, and mispricing cycles. Most institutions treat it as background noise – a number that describes how much a portfolio moves. That framing is expensive.

When you buy an option, you are buying volatility. When you sell one, you are selling it. The direction of the underlying is secondary. The real question is whether the volatility embedded in the option price is a fair estimate of what will actually happen. If it is too high, the seller wins. If it is too low, the buyer wins. Every option trade is ultimately a bet on that gap.

This paper develops what Paper 1 introduced. Paper 1 established that options are bundles of risk exposures and that volatility is the core driver. This paper goes further: it explains how to think about volatility as a standalone asset, how to evaluate whether it is mispriced, and how institutions systematically get this wrong in ways that cost real money.

The argument proceeds in four stages:

- Volatility has a persistent risk premium that makes selling it profitable on average – but that premium is not constant and collecting it without understanding its dynamics is how carry trades become disasters.
- The term structure of volatility contains information that most institutions ignore. Short-dated and long-dated vol behave differently and respond to different forces.
- Skew is structural, not anomalous. Understanding why downside puts are expensive tells you more about market structure than about mispricing.
- Volatility regimes are real and identifiable in broad terms. Deploying options without a view on the current regime is trading blind.

This is not a guide to volatility arbitrage. It is a framework for institutional investors who use options as risk management tools and need to understand what they are actually paying for – and when they are paying too much.

The difference between buying expensive volatility and cheap volatility, compounded across a portfolio and across years, is the difference between a hedge programme that works and one that steadily drains capital.

The practical test is simple: if you cannot explain the vol level, the term structure, the skew, and the regime in a single decision, you are probably trading blind.

1. What Volatility Actually Is

Before you can trade volatility well, you need to be clear about what it is and what it is not.

Volatility, in its realised form, is a statistical measure of how much an asset's price moves over a given period. It is calculated from historical returns – typically as the annualised standard deviation of daily log returns over a defined window. A stock that moves 1% per day in either direction has higher realised volatility than one that moves 0.3% per day. That is the maths.

Implied volatility is different. It is extracted from option prices using a pricing model, most commonly Black–Scholes or a variant. If a one-month at-the-money call on an equity index is trading at a price that implies 18% annualised volatility, then 18 is the implied vol. It is not a forecast. It is the market's current price for uncertainty over the option's life. Supply and demand set it, and models translate it.

The distinction matters because the two frequently diverge, and that divergence is where option economics live.

1.1 The volatility risk premium

Over long samples, implied volatility tends to run above realised volatility. The gap is not random. It is systematic and persistent, particularly in equity markets.

The reason is straightforward: buyers of options are willing to pay above expected value for the protection they provide. An institution buying equity puts is not just paying for expected losses; it is paying a premium for the certainty of knowing the worst-case outcome. That premium is the volatility risk premium, and it is the foundation of every short vol carry strategy. It is not stable across regimes, maturities or asset classes.

Empirically, this premium has averaged roughly 2 to 4 volatility points in equity markets over long periods – meaning implied vol has run about 2 to 4 points above realised vol on average. That sounds modest, but compounded across multiple option cycles per year, it represents a meaningful and persistent return to short vol positions.

The volatility risk premium is real. It is also not free money. It is compensation for tail risk, and it is collected steadily until the moment it is given back violently.

This asymmetry is the central structural fact of short vol strategies. The carry is positive in normal conditions. The drawdowns are sharp and rapid when conditions change.

Understanding that shape, not just the average return, is what separates a well-governed short vol programme from a disaster waiting to happen.

1.2 Realised vs implied: the gap is the trade

The gap between implied and realised vol is not fixed. It varies across market regimes, asset classes, and maturities. In calm equity markets, the gap is wide: implied vol is comfortably above realised, and short vol strategies collect premium efficiently. Going

into a stress event, realised vol spikes and frequently exceeds implied for a period. After the stress, implied vol mean-reverts, often overshooting to the downside before normalising.

Each phase of that cycle has different implications for option buyers and sellers. Buying vol in the calm phase is often expensive on a carry basis and likely to produce losses (unless the portfolio objective is explicit insurance). Selling vol at the peak of a stress event is cheap but requires both the analytical conviction and the operational capacity to do it when the market is in chaos. Most institutions can manage neither end of that trade well.

The practical implication is not that institutions should try to time the vol cycle precisely. It is that the path of returns matters as much as the end state: a trade can be right on average and still be monetised badly if the timing is wrong.

The institution should have a clear view of where in the cycle they are buying or selling, and size accordingly. Buying protection when implied vol is 30 is not the same decision as buying it when implied vol is 12, even if the underlying instrument is identical.

Implied vol below realised	Implied vol above realised
Options are cheap relative to historical moves	Options are expensive relative to historical moves
Good time to buy protection; carry headwind is low	Short vol carry is attractive; long vol is fighting the premium
Market may be underpricing uncertainty	Market is pricing in more fear than history justifies
Rare in normal conditions; common in early-cycle recoveries	The structural norm in equity markets over long periods

2. The Term Structure of Volatility

Implied volatility is not a single number. It varies across both strike (the skew, covered in Section 3) and time (the term structure). Understanding the term structure is essential for anyone making duration decisions in options – which maturity to buy, when to roll, and how to read what the market is pricing for different horizons.

2.1 What the term structure tells you

The vol term structure plots implied volatility against time to expiry for a given strike, usually at-the-money. In normal market conditions, the term structure is upward sloping: longer-dated options carry higher implied vol than short-dated ones. This reflects the basic fact that more can go wrong over a longer period.

In stressed conditions, the curve can invert sharply. Short-dated vol spikes sharply as the market prices in immediate uncertainty, while long-dated vol rises more slowly. An inverted term structure is a signal that the market is pricing a specific near-term event or crisis rather than a general increase in long-run uncertainty.

The shape of the term structure is informative in its own right. A steeply upward-sloping curve suggests the market expects current low volatility to persist in the short term but is less certain about the medium term. A flat curve suggests the market sees similar uncertainty across horizons. An inverted curve signals acute near-term stress.

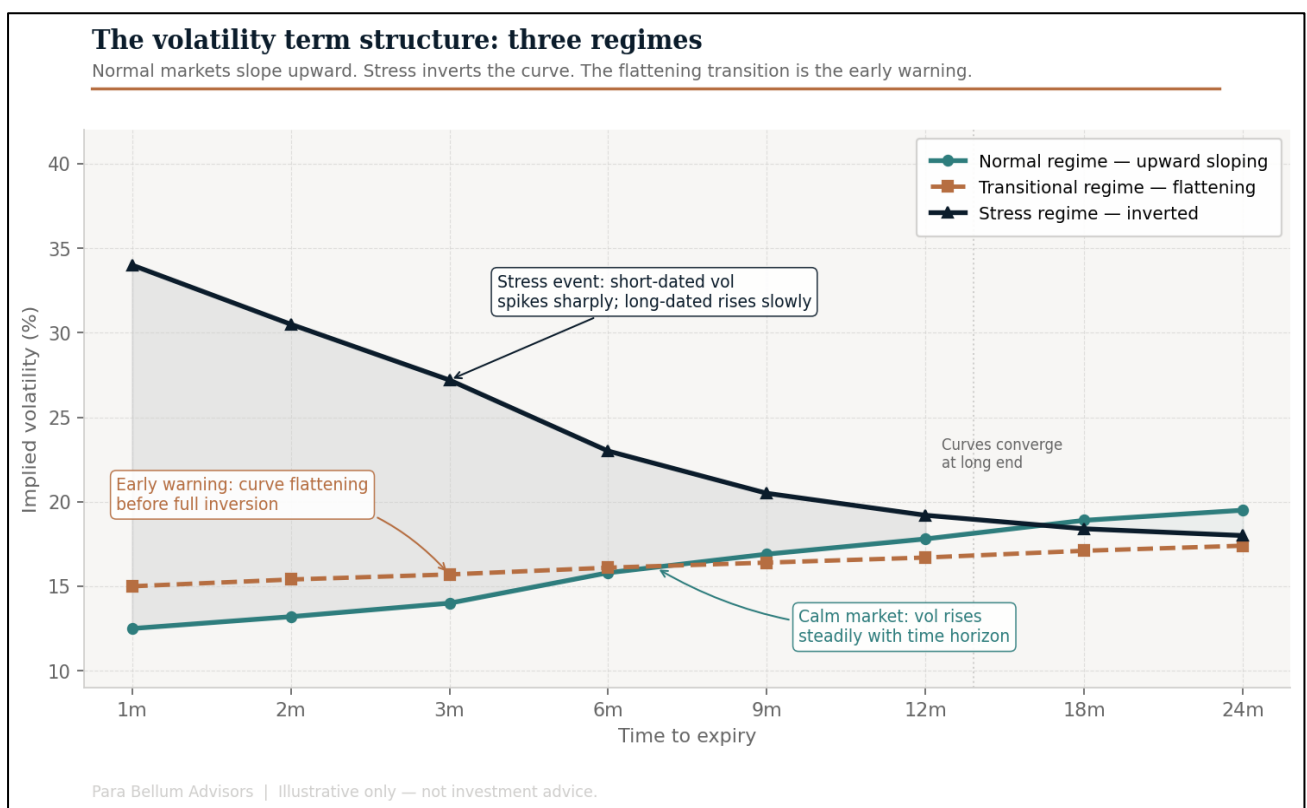


Figure 1: The volatility term structure across three regimes. Normal markets produce an upward-sloping curve. Stress inverts it as short-dated vol spikes. The transitional flattening is the early warning signal.

2.2 Short-dated vs long-dated vol: different animals

Short-dated and long-dated options respond to different forces and behave differently in portfolios. Conflating them is a common source of confusion in institutional option programmes.

	Short-dated (under 3 months)	Long-dated (over 6 months)
Primary driver	Event risk, near-term sentiment, gamma demand	Long-run uncertainty, structural macro risk
Response to stress	Spikes sharply and rapidly	Rises more slowly, reverts more slowly
Theta behaviour	Accelerates sharply into expiry	Decays slowly, more stable cost profile
Liquidity	Higher in liquid underlyings; thinner for credit and FX	Lower; bid/ask spreads wider
Best use case	Tactical protection; event hedging; gamma trading	Structural overlays; long-duration hedging programmes
Key risk	Rolls poorly; gamma exposure requires active management	Vega exposure large; slow to monetise in fast-moving markets

The practical point is that vol is not just level and shape – it is tradability. In stress, prices move and liquidity disappears together.

The implication is that an institution building a protection overlay should not simply buy the cheapest option across the curve. The maturity choice is a structural decision.

Short-dated options are better at capturing specific stress events but require active rolling and generate high theta costs.

Long-dated options are cheaper per unit of time but carry significant vega exposure and can be slow to respond in a fast-moving market.

2.3 Term structure trades

The term structure is not just descriptive. It is tradeable.

Calendar spreads – buying one maturity and selling another – allow investors to take a view on the shape of the curve without taking a directional view on vol level. Buying short-dated vol and selling long-dated vol profits if the curve inverts, which tends to happen during stress events. The reverse trade profits from a normalising or steepening curve.

For institutional portfolios, the most relevant term structure insight is this: buying long-dated protection during periods of low vol is structurally more efficient than rolling short-dated puts. The carry cost is lower per month, the theta profile is more manageable, and the position does not need to be rebuilt after every expiry. The trade-off is less gamma sensitivity and slower monetisation in a sharp but brief event.

Most institutions buy short-dated options because they feel more certain. But short-dated protection is the most expensive form of certainty available. Long-dated options, bought in calm conditions, are usually the better starting point for a structural hedge.

3. Skew: Structural Fact, Not Mispricing

The volatility skew is one of the most misunderstood features of options markets. It is consistently described by users as an anomaly, an inefficiency, or evidence that puts are overpriced. It is none of those things.

3.1 What skew is

Skew refers to the asymmetry in implied volatility across strikes. In equity markets, puts struck below the current price carry higher implied vol than calls struck above it, even for options at the same distance from spot. A put at 90% of spot might carry 22% implied vol while a call at 110% carries 16%.

Under the assumptions of Black-Scholes, all strikes at the same maturity should carry the same implied vol. They do not. The skew is the market's departure from that assumption, and it is persistent and consistent across markets and time.

3.2 Why skew exists

Skew exists because of structural supply and demand, not because the market is making a forecasting error.

On the demand side, there is persistent institutional buying of downside puts. Superannuation funds, insurers, asset managers with liability constraints, and funds with mandate-based drawdown limits all need equity protection. That demand is not discretionary or timing-sensitive – it is structural. It shows up in the market consistently, regardless of where the market is trading.

On the supply side, providing downside protection requires dealers and vol sellers to be short gamma in the tail. That is an uncomfortable position to run. Dealers charge a premium for it. Additionally, equity markets have a well-documented tendency to fall faster than they rise. Realised vol in down markets is genuinely higher than realised vol in up markets. Some of the skew reflects that asymmetry in actual market behaviour, not just demand imbalance.

3.3 What skew tells you about the cost of protection

When you buy a downside put in an equity market, you are paying skew. Part of your premium is explained by the option's time value and the ATM vol level. Part of it is the skew premium – the additional cost of sitting below spot on the vol surface where institutional demand is concentrated.

This is not a reason to avoid buying protection. It is a reason to understand what you are paying for. An institution that buys a put struck at 85% of spot is not paying the same thing as a put struck at 95%. The 85-strike option carries more skew premium, which means it requires a larger, sharper drawdown to break even in premium terms.

The skew is not telling you that downside puts are mispriced. It is telling you that downside protection is in persistent institutional demand and that providing it carries structural risk. You are paying

a market-clearing price for something that other sophisticated buyers also want.

3.4 Skew strategies

Several strategies exist specifically to monetise or trade around skew, rather than taking an outright view on vol level.

A risk reversal is a common skew trade: selling a downside put and buying an upside call at the same premium. The trade is delta-neutral and monetises the premium differential between the two strikes. It works if the underlying stays stable or rises and the skew premium narrows. It is painful if the market falls sharply, because the short put accelerates against the position exactly when it hurts most.

Put spreads are a more conservative approach to managing skew cost. Buying a put at 95% of spot and selling a put at 80% reduces premium outlay significantly by monetising the skew. The trade-off is that protection is capped below 80%.

Dispersion trading exploits the relationship between index skew and single-name skew. Index implied vol typically carries a premium over the weighted average of its constituents' implied vols, reflecting the additional demand for index-level protection. Buying single-name vol and selling index vol is a way to monetise that premium.

These are useful structures, but none of them eliminates skew risk – they just repackage it.

4. Volatility Regimes

Volatility does not move randomly over time. It tends to move in regimes. Periods of low, stable vol tend to persist. When they break, they break sharply and the transition to a high-vol regime happens faster than most participants expect.

4.1 The four regimes

For practical purposes, there are four vol regimes that matter:

Regime	Characteristics	Implication for options
Low vol, complacent	Implied vol below long-run average; skew compressed; term structure steep	Options cheap; ideal time to buy protection; short vol carry less attractive
Low vol, deteriorating	Vol starting to rise from low base; breadth narrowing; credit spreads widening	Warning signal; cost of protection rising but still manageable; buying window closing
High vol, acute stress	Vol spiked; term structure inverted; skew steep; liquidity thin	Protection expensive; monetise existing positions; new hedges costly and vega-loaded
High vol, mean-reverting	Vol elevated but falling; markets recovering; term structure normalising	Short vol carry attractive; selling elevated vol with defined risk has positive expected value

No framework can predict transitions between regimes precisely. What it can do is change the default decision when buying or selling vol. Buying protection in a low-vol, complacent regime is structurally sound even if timing is uncertain. Buying it at the peak of acute stress is structurally poor regardless of how scary the environment feels.

4.2 Vol as a mean-reverting asset

Volatility is generally mean-reverting, often more so than many other market variables.

Unlike equity prices (or other underlyings like, FX or rates etc), which can trend indefinitely, vol has a gravitational pull toward its long-run average. Extreme spikes in implied vol are reliably followed by sharp declines. Sustained periods of very low vol tend to be followed by increases, though the timing of that increase is highly uncertain.

This mean reversion is part of why short vol carry works over long periods: vol almost always comes back down after it spikes, and if you have sold vol at an elevated level, time is on your side. The practical implication is simple but frequently ignored: the most attractive time to buy volatility is often when it feels least necessary, and the best time to sell it is when it feels most dangerous.

4.3 VIX and the limits of headline vol metrics

The VIX is a useful reference but a poor substitute for a genuine vol framework. It captures one maturity, one underlying, and one strike range. It tells you nothing about the vol surface shape, the term structure, or vol in other asset classes.

More importantly, VIX tends to spike after the event rather than before it. In the days preceding sharp equity selloffs, VIX often rises modestly or not at all. The spike comes when the selloff is already underway. Using VIX as a trigger for buying protection consistently delivers the wrong timing: you buy protection when it is already expensive and the damage is already starting.

VIX often tells you what happened to vol after the market has already moved. A vol framework tells you what you should be doing about it today.

5. Evaluating Whether Vol Is Mispriced

The question most institutional option users cannot answer clearly is: is this vol cheap or expensive? That inability is itself a problem. Buying options without a view on vol pricing is like buying bonds without a view on whether yields are fair.

5.1 The starting point: implied vs realised

The most direct measure of whether vol is cheap or expensive is the current implied-to-realised spread. If one-month ATM implied vol is 18% and realised vol over the past month has been 12%, the premium is 6 points. Historically, the average premium in equity markets has been 2 to 4 points. A 6-point premium suggests options are expensive relative to history. But that conclusion still depends on whether the current regime is likely to persist or reverse.

This is a starting point, not a conclusion. Realised vol is backward-looking. Implied vol is forward-looking. The correct comparison is not realised-to-implied but expected-future-realised to implied.

5.2 Vol percentile analysis

Vol percentile analysis compares current implied vol to its historical distribution for the same underlying and maturity. If current one-month ATM implied vol is at the 20th percentile of its 5-year history, options are historically cheap. If it is at the 80th percentile, they are historically expensive.

This does not tell you what will happen next. But it does provide a calibration point for how aggressively to deploy option premium. An institution buying protection in the 80th percentile of vol history should size accordingly and have a clear view on what makes the current environment different enough to justify that premium.

5.3 The term structure as a signal

When short-dated vol starts rising while long-dated vol remains stable, it often signals that near-term event risk is being priced in. When long-dated vol rises first, it tends to reflect broader macro uncertainty. The distinction matters for structure selection – event risk calls for short-dated protection, structural uncertainty calls for longer-dated instruments.

The curve matters because it tells you whether the market is pricing immediate stress or a more persistent repricing of uncertainty.

5.4 Cross-asset vol signals

When credit spreads are widening, FX vols are elevated, and rates vol is rising simultaneously, the signal is more significant than a move in equity vol alone. But the same move in each asset class does not mean the same thing: rates vol, FX vol, and equity vol have different surfaces, different carry profiles, and different mean-reversion characteristics.

When equity vol is elevated but other markets are calm, it is more likely to be noise or positioning-driven than a genuine regime shift.

A vol spike confined to one asset class is often a positioning or local supply-demand story, while a spike that propagates across equities, credit, FX and rates is more likely to reflect a broader regime shift. The response to each should be different.

6. How Institutions Get This Wrong

The failure modes in institutional vol management are well-documented but consistently repeated. They are not failures of intelligence or sophistication. They are failures of process, governance, and incentive structure.

6.1 Treating vol as a cost rather than an asset

The most foundational error is framing option premium as a cost to be minimised rather than an asset to be priced and deployed. When the treasury or risk team views option premium as overhead, the natural response is to buy as little as possible, buy the cheapest available structure, and cut the programme when it has underperformed for two consecutive quarters.

The correct framing is that option premium is the cost of a specific portfolio characteristic – a shaped return distribution, a convex response to stress, a source of liquidity when everything else is illiquid.

6.2 Reactive buying and the timing trap

Institutions buy protection after they have been frightened. In a low-vol regime, one-month ATM implied vol on a major equity index might be 12%. After a 10% drawdown, that same vol might be 28%. The institution that waited has paid 130% more for protection after it was already needed.

This is why convexity budgets should be set in calm periods, not after the drawdown starts.

The governance implication is that protection budgets need to be set prospectively, not reactively.

6.3 Short-dating the portfolio

Rolling a monthly put programme over 12 months costs more in aggregate than buying a 12-month option, and it also embeds a recurring timing decision that institutions often get wrong.

It requires more active management, more operational capacity, and more frequent decision points where governance friction can disrupt the programme. More importantly, short-dated protection is poorly matched to institutional drawdown scenarios – credit events, policy shifts, structural market repricing – which often take months to develop.

6.4 Selling vol without owning the tail

A covered call programme framed as yield enhancement generates a line in the portfolio report that shows steady income. The tail risk that has been taken on does not appear on the same line. When the vol spike arrives and the position loses three months of premium in a week, the governance framework has no pre-agreed response because the risk was never explicitly owned.

The fix is to frame short vol strategies honestly as risk positions, size them relative to the tail exposure they carry, and have explicit criteria for when they are reduced or exited.

The problem is not selling vol per se; it is selling it without a defined tail-risk framework.

6.5 Ignoring cross-asset vol

Most institutional option programmes are single-asset. That siloed approach misses diversification opportunities and misses early warning signals. Cross-asset vol propagation is one of the most reliable indicators of a genuine regime shift. An institution that only monitors equity vol will often miss the signal until it is too late to act cheaply.

7. Practical Framework for Volatility Evaluation

This section provides a working framework for evaluating vol conditions before making option decisions. It is not a model. It is a set of questions that, answered honestly, produce a more disciplined approach to vol deployment than most institutional processes currently use.

7.1 The five questions

Before buying or selling volatility in any form, a portfolio manager should be able to answer these five questions:

Question	Why it matters
Where is current implied vol relative to its own history?	Sets the baseline for whether you are buying cheap or expensive. 20th percentile is a different decision from 80th percentile.
Where is implied vol relative to recent realised vol?	Measures the current premium you are paying or collecting. A wide gap favours sellers; a narrow or inverted gap favours buyers.
What does the term structure look like?	Identifies whether stress is near-term or structural. An inverted curve signals acute risk; a steep curve signals calm with distant uncertainty.
Are other asset class vols confirming or diverging?	Cross-asset confirmation increases conviction. Single-asset spikes are more likely to mean-revert quickly.
What specific portfolio problem is this option solving?	Without a clear answer, the option premium cannot be evaluated. Cheap vol with no clear purpose is not a bargain.

These are not questions that require a volatility desk or proprietary data to answer. They require access to standard market data and the discipline to ask them consistently before transacting.

7.2 Sizing to the vol environment

Option notional sizing should vary with vol conditions. Buying protection when implied vol is at the 20th percentile of history is an opportunity to buy more protection for the same premium budget.

Buying at the 80th percentile is a time to be selective about structure and size.

Sizing is part of the decision, not an afterthought.

7.3 The role of vol in a convexity budget

A convexity budget is most efficiently deployed when vol is low and options are cheap.

Without a vol framework, a convexity budget is just a spending target with no deployment logic.

With a vol framework, it becomes a disciplined process: deploy more when vol is cheap, deploy less when vol is expensive, adjust structure based on the shape of the term structure, and use cross-asset vol signals to calibrate timing.

A convexity budget without a vol framework is a cheque without a payee. The money is there. The discipline to spend it well is not.

8. Conclusion

Volatility is an asset. It has a price, a term structure, a smile, and a regime cycle. It can be cheap or expensive, and those conditions change over time in ways that matter for portfolio construction. Many institutional option programmes operate without a clear view on whether the volatility they are buying or selling is fairly priced, and that indifference has a cost that compounds over time.

The central arguments of this paper are these:

- The volatility risk premium is structural and persistent, but it is not constant. Its dynamics determine the economics of both long and short vol strategies.
- The term structure of volatility contains information about the nature and timing of market stress. Short-dated and long-dated options are different instruments that respond to different forces.
- Skew is not a mispricing. It is a market-clearing premium for structural demand. Understanding it changes how you evaluate the cost of protection.
- Vol regimes are identifiable in broad terms. Deploying options without a view on the current regime is trading blind.
- The consistent institutional failure modes – reactive buying, short-dating, treating premium as cost – are not technical failures. They are governance failures that a clear vol framework can address.

Paper 3 in this series, *Convexity in Portfolios: Cost, Budgeting, and When It Fails*, addresses how to translate the vol framework developed here into a portfolio-level convexity budget.

The question is not whether to buy volatility. It is whether the price, the regime and the portfolio objective justify it. That question requires a framework. Most institutions do not have one.

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/