Para Bellum Advisors

PRACTITIONER PAPER

Quantitative Investment Strategies

From Systematic Exposure to Structural Reality

By Mike Duncan, Para Bellum Advisors

Version 1, April 2026
www.parabellumadvisors.com

Executive Summary

Quantitative investment strategies (QIS) have become one of the most widely distributed products in institutional portfolios.

They are marketed as systematic, diversified, transparent, liquid, and cost-efficient. That framing is not wrong. It is incomplete in ways that matter most when markets stop behaving.

QIS is not just a strategy. It is a structure and that structure introduces risks that are largely absent from standard presentations:

- Counterparty dependency
- Pricing captivity
- Conditional liquidity
- Balance sheet transformation
- Model opacity

These risks do not show up in backtests. They surface in stress.

Most allocators believe they are buying factor exposure. In reality, they are also buying a set of structural dependencies they have not fully examined.

Key conclusions:

1. QIS is bilateral OTC exposure to a dealer, not market exposure to a strategy
2. Liquidity is conditional on dealer balance sheet, not market depth
3. Pricing is model-driven and not independently observable
4. Path dependency drives outcome dispersion between investors in identical strategies
5. Non-replicable strategies cannot be fully risk-managed
6. All-in costs are materially higher than fee schedules suggest
7. Structural differences explain why identical strategies produce different outcomes

The central point is simple: the structure determines the outcome more than the signal.

Why this matters now?

This is not a static problem. It is intensifying. Three structural shifts have made QIS more fragile than it appears:

1. **Crowding:** more capital is running more similar signals than at any point in the history of the product category
2. **Balance sheet constraints:** dealer capacity to intermediate risk is more limited and more cyclical than in the pre-2008 environment
3. **Strategy convergence:** differentiation between providers has narrowed, increasing the likelihood of synchronised behaviour under stress

These forces do not show up in backtests. They define outcomes in real markets.

> *In QIS, implementation is the strategy. The signal is secondary.*

1. The Gap Between Promise and Practice

Quantitative investment strategies arrived in institutional portfolios as a solution to several problems at once. Active managers were failing to justify fees. Factor investing was academically legitimate but operationally cumbersome. Portable alpha was appealing in theory but difficult to execute at scale. QIS offered a way to package systematic exposures into accessible, apparently transparent structures with competitive cost profiles.

The pitch is coherent. Factor premia are real. Systematic approaches reduce behavioural bias. Bank distribution networks make execution straightforward. The documentation is polished. The risk statistics look reassuring.

What the pitch does not say is that most QIS –whether from banks or major asset managers– are still bilateral OTC contracts with a dealer, priced by a dealer model, and exitable only on terms set by that dealer. The factor exposure is real. The structural dependency is also real. Most allocators think they are buying the first and not the second.

That gap is what this paper addresses.

1.1 How QIS grew

The QIS market expanded materially in the decade following 2008, driven by three conditions that converged. The collapse of traditional alpha in active management pushed allocators toward systematic approaches with more honest cost profiles. Post-GFC regulation increased the cost of balance sheet for dealer proprietary trading, accelerating the repackaging of trading strategies into distributed client products. A prolonged low-rate environment created demand for yield and diversification that traditional asset classes could not supply.

The result was rapid growth across momentum, carry, value, low volatility, and multi-factor strategies. By the early 2020s, QIS had become a standard component of sophisticated institutional portfolios, often sitting in alternatives allocations or portable alpha sleeves. Growth created complexity. As strategies multiplied, the distance between the underlying exposure and the investor's understanding of what they owned increased. That distance is now one of the primary sources of structural risk in institutional portfolios.

1.2 Two lenses

The most useful way to think about QIS is through two simultaneous lenses:

- **The strategy lens:** what you think you are buying. Factor exposures, return streams, portfolio characteristics. This is what the marketing material addresses.
- **The structural lens:** what you are actually exposed to. Counterparty dependency, pricing mechanics, liquidity conditions, balance sheet linkages. This is what the marketing material does not address.

Both lenses are legitimate. The strategy lens is not wrong. The structural lens is simply absent from most allocator due diligence. This paper is primarily written from the second lens, because that is where the risks most likely to cause damage in practice actually live.

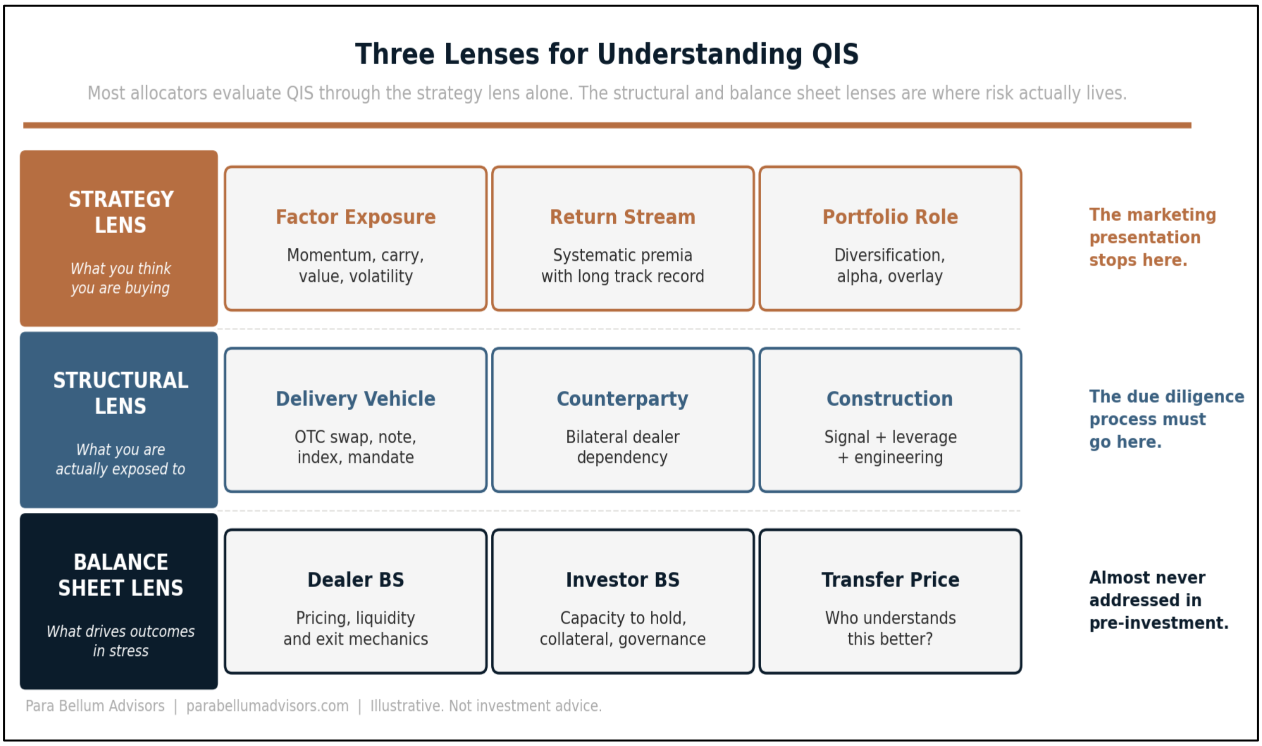

Figure 1: Three lenses for understanding QIS – strategy, structure, and balance sheet.

2. What QIS Really Is

Quantitative investment strategies (QIS) are systematic, rules-based investment products distributed primarily by investment banks that package exposure to well-documented factor premia – momentum, carry, value, low volatility, and quality – into derivative structures accessible to institutional investors.

Unlike active funds, QIS products follow pre-defined, algorithmically executed signals with no discretionary portfolio management. Unlike passive index funds, they target return premia that are distinct from broad market beta, typically using leverage to deliver economically meaningful return levels from small raw factor premia.

The defining structural feature of QIS is the delivery mechanism: exposure is accessed via bilateral OTC contracts – swaps, structured notes, or index-linked derivatives – with a single dealer counterparty who prices, executes, and provides liquidity for the strategy. It is this structural characteristic, more than the underlying factor logic, that determines how QIS behaves in practice.

At its most basic: QIS is a set of packaged factor exposures, delivered via derivative structures and distributed by investment banks. The underlying factors are real. Momentum, carry, value, volatility, and quality premia have long track records in academic literature and live markets. The delivery mechanism is where the complexity begins.

Factor exposures are not new. What QIS adds is engineering. Signals are combined, rebalanced, weighted, and then wrapped into an index structure or a swap or a structured note that an institutional investor can access via a standard ISDA agreement.

2.1 The three components

Every QIS product is a combination of three things, whether or not they are presented that way:

- **Beta:** systematic exposure to one or more factor premia. This is the return source the strategy claims to access. It is real, but it is not unique. The same factor exposure is often available through simpler or cheaper means.
- **Leverage:** factor strategies typically require leverage to deliver useful return levels from small raw premia. The leverage is either explicit, embedded in the derivative structure, or both. It is rarely disclosed in plain terms.
- **Engineering:** construction choices that determine how the signal is translated into a portfolio. Rebalancing frequency, signal weighting, execution methodology, turnover controls. These choices have a larger impact on outcomes than the underlying factor in most cases.

2.2 The delivery structures

The choice of delivery structure is one of the most consequential decisions in QIS, and it is frequently treated as an administrative detail rather than a risk decision.

- **Index-linked structures:** the QIS strategy is published as an index. The investor accesses it via a total return swap referencing that index. Pricing is publicly available in theory; in practice the index methodology is controlled by the dealer and can be changed. Exit is via unwinding the swap at dealer mark (with the original dealer).

- **OTC swaps:** bespoke exposure negotiated directly with the dealer. More flexible, less transparent. Pricing, construction, and exit are entirely bilateral.
- **Structured notes:** QIS exposure embedded in a note, often with capital protection or enhanced coupons. Introduces issuer credit risk on top of strategy risk. Liquidity is typically the worst of the three delivery structures.
- **Managed accounts:** the investor owns the underlying positions. Highest transparency, highest operational burden. True secondary market exit is possible if positions are listed. Still depends on dealer for strategy construction and signal.

> ***QIS is beta plus leverage plus engineering, wrapped for institutional consumption. Know what you are actually paying for.***

3. Anatomy of a QIS Strategy

Most QIS presentations begin and end with the signal. The factor model is explained, the academic literature is cited, and the backtest is shown. Everything that happens after the signal is generated is treated as implementation detail. It is not. Construction choices have a larger effect on realised returns than most allocators appreciate.

3.1 The five layers

Signal

The factor model that generates the investment signal. Momentum rank across equities, cross-sectional carry in FX, value spreads in credit. The signal is typically the best-understood component of the strategy because it is the most disclosed. It is also the component that has been most extensively curve-fitted in the design process.

Portfolio construction

How the signal is translated into a portfolio. Risk weighting, factor diversification, correlation assumptions, position sizing. This layer often contains the most significant discretionary choices in the entire product, embedded silently in the methodology. Vol-weighting is standard. Its effect is to increase position size when volatility is low and reduce it when volatility is high. In practice this means the strategy is most exposed precisely when a regime shift is beginning, before realised vol has risen to trigger the scaling reduction.

Execution

How trades are executed against the signal. Rebalancing frequency, execution window, crossing internal flow, interaction with dealer inventory. This is where paper returns meet real markets and the gap opens. Execution costs in QIS are systematically understated. Backtests use mid-market prices. Live execution does not. For strategies with high turnover or exposure to less liquid underlyings, the difference between backtest and live performance is dominated by execution cost, not signal quality.

Financing

How the positions are funded. This includes the explicit swap spread or funding cost embedded in the derivative, but also the collateral requirement, margin dynamics, and any embedded leverage cost. Financing is the layer most frequently ignored in pre-investment analysis and most frequently cited post-investment as explaining the performance shortfall.

Rebalancing

How often the portfolio reconstitutes, what triggers reconstitution, and what happens at rebalance points. High-frequency rebalancing increases turnover cost. Low-frequency rebalancing allows factor drift that may create unintended exposures between rebalance dates. In stress conditions, rebalancing mechanics can become the primary driver of returns. Vol targeting strategies that rebalance into rising volatility systematically sell into drawdowns. The mechanical predictability of this behaviour invites front-running by informed market participants.

3.2 Construction choices matter more than the signal

Two strategies with identical signals can produce materially different outcomes based solely on construction choices. Rebalancing frequency, turnover controls, execution methodology, and volatility targeting parameters are all within the dealer's discretion in a standard QIS product. When evaluating a QIS product, the appropriate question is not only what the factor exposure is, but what are the construction choices and who controls them.

THE QIS RISK STACK

Every QIS strategy sits on top of a layered risk structure. Most allocator analysis focuses on layer 1. Most realised losses occur in layers 3 to 6.

Layer 1 – Signal risk

Does the factor work? Is the signal real, persistent, and non-deceptive? This is where almost all pre-investment analysis is concentrated.

Layer 2 – Construction risk

How the signal is translated into a portfolio. Vol targeting parameters, rebalancing rules, position sizing. Construction choices drive outcomes as much as the signal itself.

Layer 3 – Execution risk

How trades are actually implemented. Turnover costs, slippage, market impact, execution timing. The gap between backtest and live performance is almost entirely here.

Layer 4 – Funding risk

The cost of leverage, collateral, and financing embedded in the derivative structure. Frequently omitted from cost analysis. Compounds materially on levered strategies.

Layer 5 – Counterparty risk

Dependency on the dealer for pricing, liquidity, and marks. The dealer controls entry, ongoing valuation, and exit. This is not market risk. It is relationship risk.

Layer 6 – Exit risk

The cost and feasibility of unwinding the position under stress. Unknown until the moment it matters most. In crowded strategies, exit costs are a structural cost, not a market cost.

> ***Most QIS products are sold at the signal level. They should be evaluated at the exit layer.***

4. Implementation Reality

There is a consistent pattern in QIS outcomes. The backtest looks good. The live performance is weaker. The explanation offered is usually factor regime, market conditions, or bad luck. In most cases, the actual explanation is implementation cost.

4.1 Collateral and funding drag

Every derivative position requires collateral management. The collateral posted against a swap generates a return, typically the risk-free rate or a spread above it. What is not always credited correctly is the opportunity cost of holding liquid assets as collateral rather than deploying them in higher-returning instruments. Additionally, margin calls during drawdowns force collateral top-ups at precisely the moment when liquidity is most valuable. This is a structural feature of derivative-based strategies that tends to worsen drawdown dynamics relative to what the backtest suggests.

4.2 Roll costs, turnover, and slippage

QIS strategies with high turnover quietly bleed performance in ways that are very difficult to see in standard reporting. Each rebalance trades at market impact. Each position roll across swap tenors has a spread cost. Each futures roll in underlying contracts has a basis cost that varies with market conditions and dealer balance sheet. Individually, these costs look small. In aggregate, for a strategy that rebalances weekly or more frequently, they can consume a significant portion of the theoretical factor premium.

4.3 Balance sheet pricing

The price at which a QIS position is executed, held, and unwound is set against the dealer's balance sheet, not against an independent market. Dealers price QIS positions to reflect their cost of hedging, their capital charges, their internal transfer pricing, and their marginal cost of balance sheet usage. These inputs fluctuate. In benign conditions, dealer balance sheets are well supplied and pricing is reasonable. In stress conditions, dealer balance sheets contract. The price at which the dealer is willing to unwind your position reflects that cost.

> ***You are paying for the dealer’s stress, not just the strategy’s performance.***

5. QIS as Balance Sheet Transformation

The most useful reframing of QIS for an institutional investor is this: you are not simply buying a strategy. You are entering into a balance sheet transaction.

> ***MENTAL MODEL: QIS AS A TRANSMISSION PIPE***
>
> *A QIS investment is best understood not as a portfolio, but as a transmission mechanism.*
>
> *Return flows through a controlled pipe:*
>
> *Signal → Construction → Execution → Financing → Dealer Balance Sheet → Investor Outcome*
>
> *At each stage:*
>
> - *costs are added*
> - *assumptions are embedded*
> - *control shifts away from the investor*
>
> *The investor observes the output but does not control the mechanism that produces it.*
>
> *Two investors can access the same signal and receive different outcomes because they are connected to different pipes, with different levels of friction, pricing, and balance sheet constraint.*
>
> ***Implication:***
>
> *You are not just buying returns.*
>
> *You are buying how those returns are delivered.*
>
> ***In QIS, the pipe matters more than what flows through it.***

Every QIS trade involves a transfer of balance sheet usage between investor and dealer. The economics of that transfer – not just the strategy signal – determine the realised outcome.

The investor provides capital, collateral, and credit capacity. The dealer provides pricing, execution, hedging, and the intermediary function of connecting the strategy to underlying markets.

The economics of the trade depend on how that balance sheet usage is priced, and who has the better understanding of that pricing at the point of execution.

5.1 Dealer balance sheet drives pricing, liquidity, and exit

Three of the most important variables in a QIS investment are not driven by the strategy. They are driven by the dealer's balance sheet:

- **Pricing:** entry and exit prices are generated by dealer models. A healthy, well-capitalised dealer with a clean balance sheet prices things one way. A dealer under balance sheet pressure prices them another way. The strategy has not changed. The price has.
- **Liquidity:** the advertised liquidity in most QIS products is daily or weekly. This liquidity is conditional on dealer balance sheet availability. In practice, liquidity is reliably available when you do not need it and conditionally available when you do.
- **Exit mechanics:** when you want to exit a QIS position, you are negotiating with the same entity that sold it to you. There is no secondary market. There is no independent price. The dealer knows your position size, your likely urgency, and the cost of their hedge.

5.2 Investor balance sheet drives capacity to hold

On the other side, the investor's balance sheet determines the ability to hold through adversity and extract the full cycle return. An investor with a robust liquidity position, stable liabilities, and delegated authority to act quickly can hold through drawdowns and exit when the strategy has recovered rather than when pressure is at its worst. An investor with tight liquidity constraints, quarterly governance cycles, and no pre-authorised exit protocol will tend to exit at the worst point of the drawdown.

This is not a theoretical risk. It is a consistent empirical pattern in institutional QIS experience. The same strategy produces different outcomes for different investors because their balance sheets and governance structures create different capacity to hold through adversity.

> ***QIS is a transfer of balance sheet usage, not just risk exposure. Understanding how that transfer is priced matters more than understanding the factor model.***

6. Counterparty Dependency and Structural Fragility

The counterparty structure of QIS is the single most underappreciated source of risk in the product category. Most institutional investors would not accept a portfolio where a single entity controlled pricing, liquidity, and exit for a material allocation. Yet that is precisely the structure most QIS arrangements create.

6.1 Bilateral exposure, not market exposure

When you invest in a listed equity or a physically settled bond, your exposure is to the issuer and the market. Other market participants create a secondary market. Prices are observable. Exit is separable from any single counterparty. When you invest in a QIS product via an OTC swap or index-linked note, your exposure is bilateral. You are exposed to one dealer. That dealer prices your position, provides your liquidity, and controls your exit.

6.2 Counterparty concentration

The concentration problem is exacerbated by two structural features of institutional QIS practice. First, the QIS market is dominated by a small number of major dealers. Most institutional QIS books have exposure to two or three of the same counterparties, so the diversification at the strategy level frequently coexists with concentration at the counterparty level. Second, the same dealer often provides multiple QIS strategies to the same client across different mandates. The portfolio looks diversified. The credit risk is aggregated in ways that may not be captured by standard netting arrangements.

6.3 Pricing captivity

Pricing captivity is a structural consequence of bilateral OTC exposure. The dealer who sold you the trade is the only entity who can price it for you. They are using their model. You do not have access to the model. For governance purposes, most institutional investors accept dealer marks as the basis for performance reporting and investment committee review. This means the governance review of a QIS position is based entirely on numbers provided by the entity that profits from the position continuing. This is not fraud. It is a structural misalignment of interests that governance frameworks rarely acknowledge.

6.4 Liquidity illusion

Daily liquidity in QIS is conditional liquidity. A QIS strategy that draws down 12% during a risk event and offers daily liquidity will, in principle, allow the investor to exit. In practice, the dealer is marking the position, setting the exit price, and managing their own balance sheet constraints simultaneously. The investor exits at the mark the dealer sets, at the speed the dealer's systems allow, against a backdrop where the dealer's own risk appetite has contracted. The GFC, the March 2020 event, and the 2022 rates shock all produced episodes where dealer-intermediated liquidity contracted sharply precisely when investors most wanted to exit.

6.5 Replicability and valuation opacity

A useful test for any QIS product: can you replicate it outside the dealer wrapper? Replicability requires an observable signal, scalable execution at comparable cost, and reproducible construction. Many QIS products fail one or more of these tests.

If you cannot replicate the strategy, you cannot independently risk-manage it. Independent risk management does not require actual replication. It requires the ability to model, stress-test, and validate the position without relying on the dealer's numbers. Where that capability does not exist, governance is incomplete.

6.6 Close-out reality

When you exit a QIS position, the economics are determined by the dealer's model mark at exit, any bid-offer spread applied to the unwind, any breakage costs embedded in the structure, the dealer's cost of unwinding their hedge, and any market impact from the unwind itself. None of these are visible to the investor before the exit is executed. In stress, all of these costs move adversely simultaneously. The investor is paying for the stress environment through every layer of the cost stack at once

> ***If you cannot replicate it, you cannot fully risk-manage it. Governance frameworks that rely on dealer marks to monitor positions have a structural hole in them.***

7. The Hidden Risk Stack

Beyond the counterparty structure, QIS strategies carry a set of risks embedded in the strategy design and construction. These do not appear as counterparty risk on a standard risk report. They appear, eventually, as performance shortfall.

7.1 Crowding and the transmission mechanism

Factor crowding is a well-recognised risk in the QIS literature. What is less well understood is the mechanism through which crowding transmits into losses. Similar signals produce similar positioning across many strategies. Similar positioning leads to synchronised rebalancing when signals shift. Synchronised rebalancing creates concentrated demand for liquidity in the same instruments at the same time. Dealers absorb that demand against their balance sheets. Balance sheet strain increases the cost of liquidity. Pricing gaps widen. The strategies that need to exit most urgently face the worst exit economics.

This mechanism is not visible in historical backtests because backtests do not model the behaviour of other market participants. Most papers stop at noting that crowding exists. The more important point is how it becomes loss, and that the mechanism runs through dealer balance sheets.

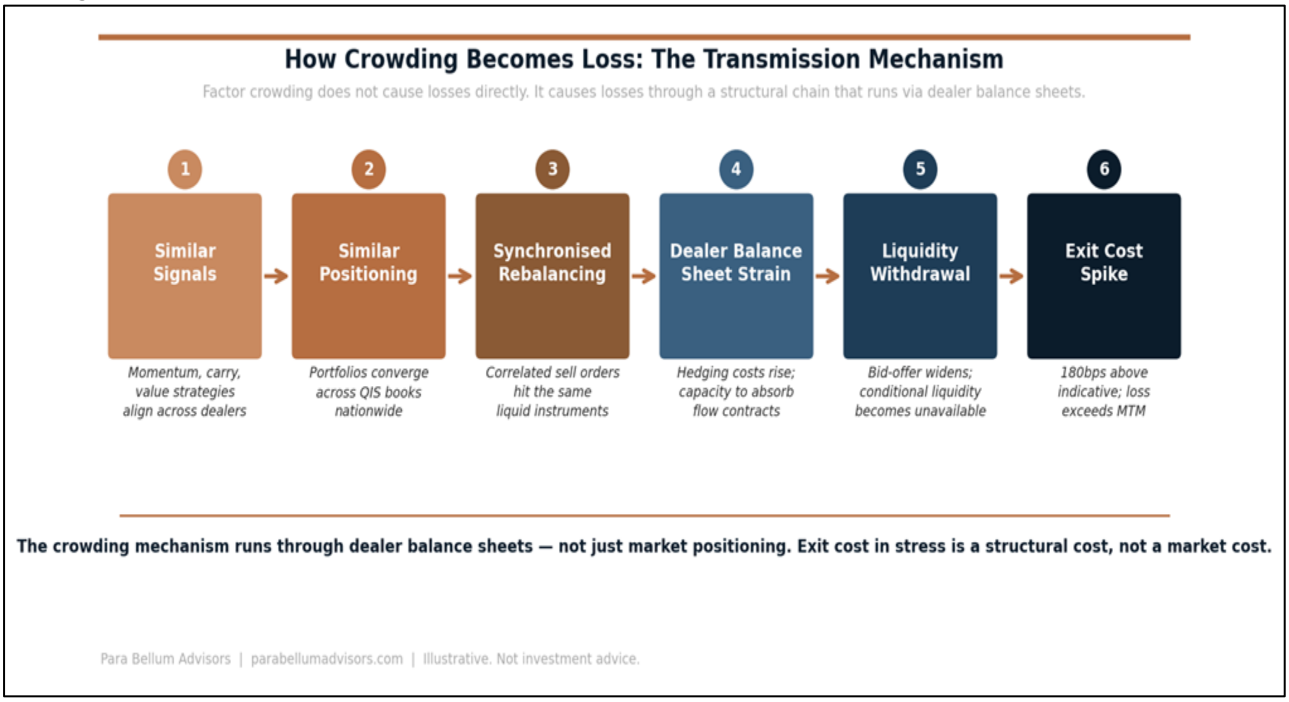

Figure 2: How crowding becomes loss – the transmission mechanism through dealer balance sheets.

7.2 Regime shifts

Factor premia are regime-dependent. Momentum requires persistent trends. Carry requires stable yield differentials. Value requires mean reversion. None of these conditions hold permanently and some fail dramatically. The 2009 momentum crash, the 2022 carry unwind in rates, and episodic vol spikes that reverse short volatility strategies are all examples of factor-specific regime failures. Vol-targeting mechanics make this worse: at the top of a good run, realised volatility is low, positions are large, and the strategy is maximally exposed.

7.3 Path dependency and outcome dispersion

Two investors in the same QIS strategy can experience materially different outcomes. Path dependency operates through three mechanisms: sequence of returns, vol control mechanics, and rebalancing timing. An investor who enters before a strong run has a larger capital base when the drawdown occurs. Vol-targeting forces buying into low volatility environments and selling into high volatility environments. Small differences in when a portfolio rebalances relative to a market event can produce meaningfully different results.

7.4 Correlation breakdown under stress

The diversification case for QIS typically rests on historically low correlations between QIS returns and traditional asset classes. These correlations are real in calm conditions. They are not reliable in stress. When investors need to raise cash or reduce risk quickly, they sell what they can sell, not what they want to sell. QIS strategies that normally exhibit low correlation to equities can see that correlation spike when both need to be liquidated to meet margin calls. The stress correlation is the relevant figure for portfolio construction, not the calm-market historical number.

8. Use Cases and Failure Mapping

QIS strategies are used across a range of institutional applications. Each use case has a coherent logic. Each also has a characteristic failure mode that follows from the structural properties described above.

Use Case	Intended Outcome	Structural Reality	Characteristic Failure Mode
Overlay management	Systematic factor exposure at low cost alongside core holdings	OTC counterparty dependency: collateral drag interacts with core portfolio liquidity	Margin calls during market stress force liquidation of core holdings to fund overlay collateral
Portable alpha	Uncorrelated return stream added to passive market exposure	Most QIS premia are liquidity-dependent and become correlated under stress; financing cost erodes the alpha	Correlation rises precisely when alpha is needed most; net alpha after financing often negative
Yield enhancement	Higher income without meaningful credit or duration risk	Embedded short optionality in carry or volatility strategies; tail risk disguised as steady income	Catastrophic loss during vol spike or carry reversal offsets years of income accumulation
Liquidity sleeve	Capital deployed in daily-liquidity strategies accessible at short notice	Conditional liquidity dependent on dealer balance sheet; exit costs elevated in stress	Cannot exit at required speed or size when the liquidity event that triggered the need occurs
Tail hedging	Convex payoff during drawdown events that funds recovery	Long vol QIS strategies bleed premium over time; monetisation requires active governance	Premium bleed abandoned strategy before the event it was designed to protect against

The failure modes share a structural characteristic: in most cases the strategy does not fail at the signal level. It fails because the structural properties of the product interact badly with the investor's governance, balance sheet, or liquidity framework under stress.

The portable alpha failure mode deserves specific attention. Most marketed portable alpha strategies are not portable alpha in the meaningful sense. They are liquidity-dependent factor premia that exhibit low correlation in historical data drawn from calm-market periods. Genuine uncorrelated alpha requires convexity: payoff that is positive when everything else is negative. Most QIS strategies do not clear that bar.

9. Performance: Backtest Versus Reality

The gap between backtest and live performance is one of the most consistent empirical regularities in institutional finance. In QIS, the gap is structural rather than incidental.

9.1 Why backtests look better

- **Execution cost assumption:** backtests use mid-market prices. Live execution crosses the spread, incurs market impact, and rolls at real basis costs. For high-turnover strategies, this alone can account for 100 to 200bps annually.
- **Funding cost omission:** backtests frequently do not reflect the actual cost of funding the strategy. The implied financing embedded in an OTC swap structure is a real cost that reduces net return.
- **Survivorship bias:** the strategies presented to investors are the ones that worked during the backtest period. The strategies that looked poor in backtest were not brought to market.
- **Regime selection:** strategy design choices are often optimised against more recent periods. The backtest looks better than live performance because the design was calibrated to the history it is tested against.
- **Rebalancing idealisation:** backtests assume instantaneous rebalancing at the close price. Live rebalancing takes time, trades at market impact, and interacts with the dealer's own inventory management.

9.2 Total cost of ownership

The total cost of a QIS investment has four components that should be assessed separately: explicit fees (management fee, performance fee); execution costs (trading costs, market impact, spread costs); funding costs (swap spread, collateral funding, financing of leverage); and structuring margin (dealer margin embedded in the product structure).

In practice, it is the last three that dominate. In many institutionally distributed QIS allocations, 40 to 60 percent of gross return is consumed by these implicit costs rather than the explicit fee.

For a typical institutionally distributed QIS product, the management fee may represent 30 to 40 percent of the all-in cost. The rest is hidden in implicit implementation drag – execution slippage, wrapper spreads, financing, and hedging costs. That means a product with a 50 bps fee and 150 bps of implicit costs is more expensive than one with a 100 bps fee and 50 bps of implicit costs. Looking at headline fees alone is misleading.

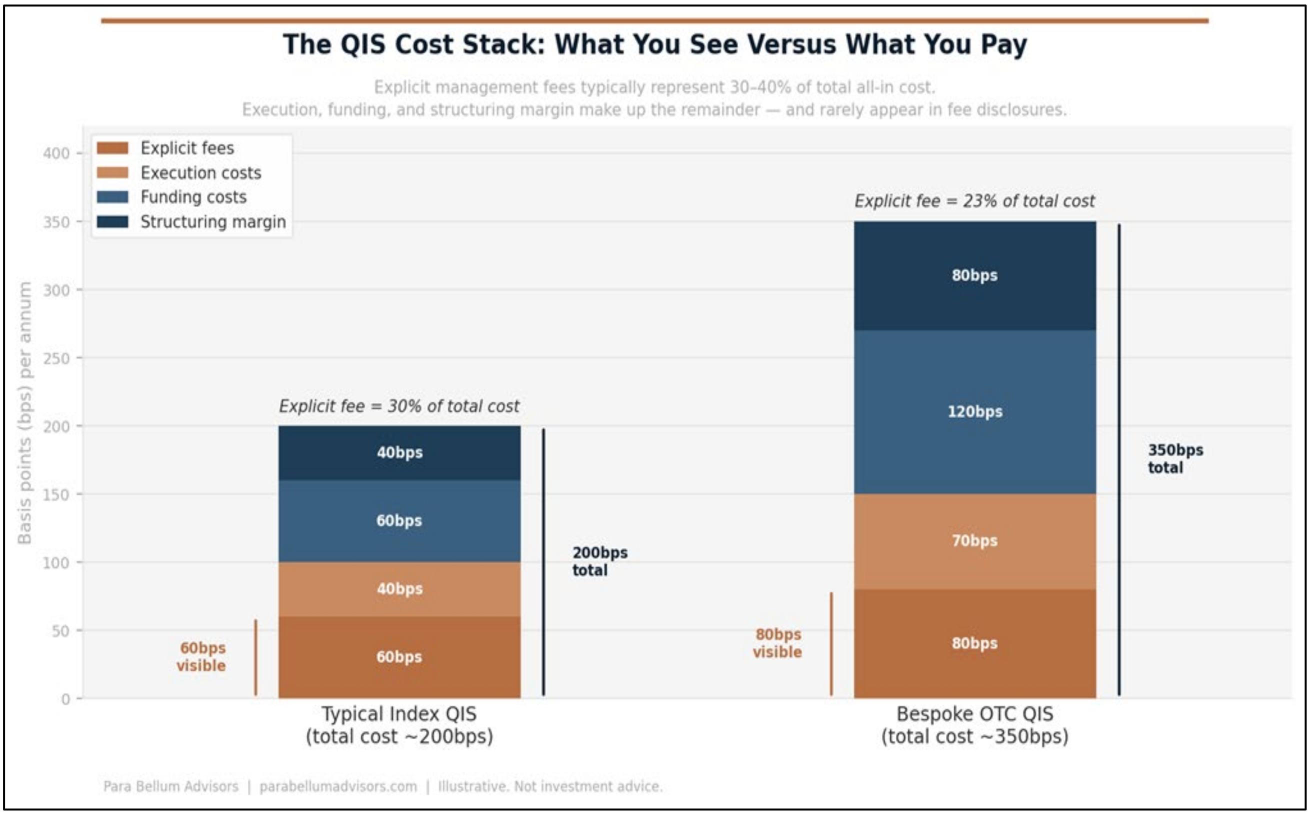

Figure 3: The QIS cost stack – explicit fees typically represent 30–40% of total all-in cost.

10. Case Studies

The following case studies are drawn from real institutional outcomes in QIS-adjacent strategies. They have been simplified and made generic to remove identifying information, but the structural dynamics in each are real and representative. The purpose is not to illustrate failure. It is to show how the structural properties described above behave in practice.

10.1 Momentum reversal and crowded exit

Setup

A multi-strategy asset manager held a large position in a cross-asset momentum QIS strategy via an index swap with a major dealer. The strategy had performed well for 18 months and had attracted additional capital from the manager's overlay programme. By the time the reversal occurred, the position was three times its initial size.

What happened

A sharp reversal in global equity momentum, triggered by a rapid rotation from growth to value factors, produced a 15 percent drawdown over eight trading days. The manager sought to reduce exposure by approximately 40 percent to bring the position back within risk limits. The exit request coincided with identical exit pressure from multiple other allocators in the same strategy. The dealer widened the bid-offer substantially. The actual exit cost was 180bps higher than the pre-request indicative.

The lesson

The loss wasn't 15 percent. It was 15 percent plus the cost of getting out. Position sizing in QIS should account for stressed exit cost, not just strategy volatility.

10.2 Short vol in a yield enhancement mandate

Setup

An insurance-adjacent investment manager ran a yield enhancement programme using a QIS strategy that sold variance in equity indices and interest rate markets. It had delivered consistent monthly income for two and a half years and had been approved for a larger allocation.

What happened

A rapid rise in rates volatility, unrelated to equity markets, triggered simultaneously in multiple markets in which the strategy was short. The strategy moved from 2 percent positive to 18 percent negative in three weeks. The income accumulated over the prior two and a half years was eliminated, and a further 8 percent of NAV was lost beyond that. The governance review had not modelled a stress scenario for the short vol strategy explicitly. The investment committee had approved the allocation based on historical Sharpe ratio.

The lesson

Strategies that derive returns from option selling are not low-risk strategies with favourable risk statistics. They are strategies that transfer risk from the near term to the tail.

The risk was never the volatility. It was the scenario the strategy was implicitly short.

10.3 Carry unwind in rising rates

Setup

A sovereign wealth fund had a meaningful allocation to a fixed income carry QIS strategy, accessed via bespoke OTC swaps with two dealers. The 2022 rates shock produced a rapid and synchronised rise in yields across all developed markets. The carry differential between long and short positions compressed sharply as central banks raised rates simultaneously. The strategy declined 22 percent over the calendar year.

What happened

The carry differential compressed by over 120bps in the span of six weeks as central bank rate decisions landed across the US, EU, UK, and Australia in rapid succession. The long legs lost on capital as yields rose; the short legs did not offset because yields were rising there too. Both dealers simultaneously widened their unwind bids as the macro environment deteriorated. The fund's risk framework had modelled rate rises as idiosyncratic – a sensible assumption in any prior decade. In 2022, the assumption was wrong. The strategy declined 22 percent over the calendar year, with the bulk of the loss concentrated in a 45-day window in the first half of the year.

The lesson

The diversification within a QIS strategy does not protect against regime changes that affect all the strategy's positions in the same direction. Macro correlation risk is a real risk in carry strategies. When the regime shifts, design matters less than your ability to exit.

10.4 Implementation done well

Setup

A family office allocated to an equity factor QIS strategy via a managed account structure, retaining full position transparency and independent exit capability. The allocation was sized conservatively at 8 percent of the portfolio, with a pre-agreed drawdown tolerance of 15 percent before a mandatory governance review.

What happened

The strategy underwent a sharp drawdown over six weeks. The drawdown reached 11 percent, short of the governance review trigger. The CIO, having full visibility into the underlying positions, was able to assess independently that the drawdown was regime-driven rather than structural. The family office held the position through the drawdown. The strategy recovered within four months. Three structural features made the difference: the managed account gave independent position visibility, governance had pre-defined drawdown tolerances with delegated authority, and conservative sizing meant the drawdown did not create liquidity pressure elsewhere.

The lesson

QIS works when structure is deliberate, not assumed. Managed account where possible, conservative sizing, pre-defined governance, independent monitoring capability.

11. The Economics of QIS

QIS is often described as an efficient product. The efficiency depends on the total cost, not the visible cost. Total cost includes the explicit fee, execution cost embedded in turnover, financing cost of the derivative structure, and the structuring margin retained by the dealer. For most investors in most QIS products, the total cost is materially higher than the explicit fee suggests.

11.1 The dealer's economics

Dealers distribute QIS products because they are profitable. The profitability comes from multiple sources. The primary source is the bid-offer embedded in ongoing product operation: execution of the underlying trades, roll costs, and rebalancing mechanics all pass through the dealer at a spread. For a strategy that turns over 300 to 400 percent of its gross notional annually, these spreads aggregate into a substantial annual cost that appears nowhere in the fee disclosure.

The secondary source is the funding spread. The dealer funds the strategy at the risk-free rate and charges the investor the risk-free rate plus a spread. In a normal interest rate environment, this is typically 20 to 50bps on gross notional. On a levered strategy, this multiplies. The tertiary source is the structuring margin in the initial transaction.

11.2 Fee stacking

In strategies distributed through intermediaries, fees can stack in ways that make the total cost very difficult to assess. An asset manager that accesses QIS via a fund vehicle may be paying on top of the embedded dealer economics in the underlying strategies. The most useful question to ask of any QIS product: what is the all-in cost as a percentage of expected gross return? If the explicit fee is 50bps and the expected gross return is 300bps, the visible cost is 17 percent of expected return. If implicit costs add another 100bps, the total cost rises to 50 percent of expected return.

> ***QIS often looks efficient at the fee level and expensive at the total cost level. The relevant comparison is all-in cost against expected gross return, not management fee against benchmark.***

12. Governance and Due Diligence

Governance of QIS positions is an area where standard institutional frameworks are frequently inadequate. The usual approach treats QIS as an investment with a mark-to-market performance series and applies the same monitoring tools used for traditional fund mandates. This misses most of the structural risks described in this paper.

12.1 The internal governance gap

One of the least discussed risks in institutional QIS is the internal governance gap: the absence of a person or committee within the institution with the mandate and capability to challenge dealer valuation marks. In most institutions, the investment committee reviews QIS performance using numbers provided by the dealer. There is often no internal capability to replicate or stress-test the valuation. The governance review is, in practical terms, a review of the dealer's representation of performance.

Closing the internal governance gap requires either building independent valuation capability or structuring the investment in a way that provides independent pricing from the outset. The managed account structure is the most effective way to close this gap. Where managed accounts are not available, commissioning independent model replication of the strategy as a governance tool is worth the cost.

12.2 Due diligence checklist

The following questions should be standard in any QIS due diligence process:

- What is the true exposure? Can you describe the risk in terms of what market conditions produce gains and losses, not in terms of factor labels?
- What drives turnover, and what is the actual annual cost at live execution prices?
- What is the all-in cost: explicit fees plus execution cost plus funding cost plus structuring margin?
- Who prices this position on an ongoing basis, and is independent verification possible?
- How do you exit? What is the exit mechanism and the estimated cost under stress conditions?
- What happens to this strategy in a regime change? Which specific market conditions produce the worst outcomes?
- What is the counterparty concentration across the entire QIS book, not just this individual position?
- Can this strategy be replicated independently, even partially, using listed instruments?
- What governance rights are retained post-investment? Who has authority to exit or reduce position size?
- Has the strategy been reviewed for path dependency? What is the sensitivity of expected return to entry timing and vol environment?

13. QIS Versus Alternative Implementations

QIS is frequently selected as the implementation vehicle for factor exposure without explicit comparison to alternative methods of accessing the same exposure. The selection is often driven by convenience: a single ISDA agreement, familiar documentation, straightforward operational setup. Convenience is a legitimate consideration. It should not be the only one.

Dimension	Direct Implementation	QIS via OTC Swap	Active Mandate
Transparency	Full position visibility at all times	methodology controlled by dealer; bespoke strategies have no independent benchmark	Depends on mandate terms; typically, high for separately managed accounts
Liquidity	True secondary market for listed; bilateral for OTC	Conditional on dealer balance sheet; not market liquidity	Generally good for liquid mandates; subject to fund redemption terms
Counterparty risk	Minimal for listed instruments; bilateral dealer exposure for OTC	Concentrated bilateral dependency; often 1–2 dealers across the whole book	Fund-level exposure to manager; risk is more transparent and separable
Cost	Execution costs only; no structuring margin or funding spread	Explicit fee plus significant implicit costs (execution, funding, structuring margin)	Explicit management fee; minimal implicit costs
Governance burden	High – requires internal pricing, monitoring, and execution capability	Low operationally; high structurally – dealer marks accepted without independent challenge	Medium – standard fund reporting and mandated review
Exit flexibility	Listed: full market liquidity. OTC: negotiated exit.	Bilateral; exit price set by dealer; cost unknown until executed; elevated in stress	Predictable redemption terms; generally straightforward

The comparison is not uniformly in favour of direct implementation. For smaller investors without the operational infrastructure for direct implementation, QIS provides access to factor exposures that would otherwise be unavailable.

The legitimate use case for QIS is the access it provides, not its theoretical cost efficiency. The problem arises when QIS is selected primarily because it is convenient and the structural trade-offs relative to alternatives are not assessed.

14. When QIS Works and When It Fails

14.1 When it works

QIS works consistently when four conditions are simultaneously met:

- **Clear role:** the strategy has a defined purpose in the portfolio with a specific return target, specific diversification objective, or specific risk function. The mandate is not 'add alpha' in an unspecified way.
- **Controlled costs:** total cost has been assessed including implicit costs and is proportionate to expected gross return. The all-in cost is known before commitment.
- **Proper sizing:** the position is sized to survive a stressed exit without creating liquidity pressure elsewhere in the portfolio. Stressed exit cost is modelled explicitly.
- **Strong governance:** the institution has pre-defined drawdown tolerances, delegated exit authority, independent monitoring capability, and a mandatory review process triggered by conditions rather than a calendar cycle.

14.2 When it fails

QIS fails consistently when one or more of the following conditions apply:

- **Bought for Sharpe:** the selection was driven by historical Sharpe ratio rather than analysis of what the strategy is implicitly short and under what conditions that short position loses.
- **Overlapping exposures:** multiple QIS strategies provide apparent diversification at the strategy level while creating concentrated counterparty exposure and correlated behaviour in stress.
- **Ignored structure:** the investment decision was made at the signal level without assessing counterparty, pricing, liquidity, and exit mechanics.
- **Blind trust in models:** ongoing governance relied entirely on dealer-provided performance numbers with no independent validation or stress testing.
- **Governance mismatch:** the institution's decision cycle was incompatible with the speed at which QIS positions need to be managed in stress.

15. Practitioner Framework

The following framework is designed as a practical decision tool for institutions evaluating, approving, or reviewing QIS allocations. It organises the analysis into five dimensions covering both the strategy and structural properties of any QIS product.

15.1 The five P framework

Purpose

Define the portfolio role precisely – the specific return target, the protection objective, or the factor premium being accessed – and specify the conditions under which the strategy will not deliver. See Appendix A for the full framework.

Purity

Assess whether the return can be attributed to the claimed factor, or whether construction choices, leverage, and execution are doing most of the work. See Appendix A for the full framework.

Price

Quantify the all-in cost – explicit fees plus execution, funding, and structuring margin – and express it as a proportion of expected gross return. See Appendix A for the full framework.

Path

Model the stressed exit cost, assess sensitivity to entry timing and vol environment, and confirm the portfolio can absorb an adverse path without forced selling elsewhere. See Appendix A for the full framework.

Practicality

Confirm the institution has independent monitoring capability, delegated exit authority, and governance triggers that operate at market speed rather than committee speed – and be candid about where gaps exist. See Appendix A for the full framework.

15.2 Allocator decision tree

Decision Point	If Yes	If No
Do I understand what market conditions cause this strategy to lose significantly?	Proceed to next check	Do not allocate until this is answered
Can I replicate or independently proxy the exposure outside the dealer wrapper?	Proceed to next check	Ensure independent monitoring capability is in place before proceeding
Is the structure necessary, or could direct implementation serve the same purpose?	Evaluate direct implementation alongside QIS	If QIS is the only viable access vehicle, acknowledge the structural trade-offs explicitly
Is the all-in cost proportionate to expected gross return?	Proceed to next check	Renegotiate or consider alternatives
Have I modelled stressed exit cost and confirmed portfolio can absorb it?	Proceed to next check	Re-size to ensure stressed exit is survivable
Does the institution have pre-defined drawdown tolerances and delegated exit authority?	Proceed to allocation	Establish governance framework before committing capital

16. The Future of QIS

The QIS market is evolving in response to some of the pressures described in this paper. The direction of travel is toward customisation, transparency, and managed accounts. But the pace of change is slow relative to the scale of the structural issues.

16.1 Customisation over commoditisation

Standardised index-based QIS products are under fee pressure. Margins on vanilla factor strategies have compressed as the market has matured. The dealer response is to move up the complexity curve into bespoke, customised solutions. For investors, this is not unambiguously good. The structural issues described in this paper are more acute in bespoke structures. More customisation typically means more counterparty dependency, less replicability, and more reliance on dealer marks.

16.2 Pressure on fees and total cost transparency

Institutional investors are becoming more sophisticated in their cost analysis. The shift toward total cost of ownership frameworks, rather than explicit fee benchmarking, is slow but real. The strategies that will persist are those that deliver net return after total cost that justifies the structural trade-offs.

16.3 Crowding risk remains

The structural crowding risk described in Section 7 has not diminished – if anything, the market has grown more concentrated. Assets under management in systematic factor strategies have continued to expand. The universe of signals across competing QIS products has converged as providers benchmark against each other. The result is that a larger pool of capital is positioned in more similar ways than at any prior point in the product category's history.

A significant, sustained factor reversal has not yet tested exit dynamics at this scale. When it does, the mechanism described in Section 7 – signals to positioning to synchronised rebalancing to dealer balance sheet strain to liquidity withdrawal – will run faster and harder than any historical episode suggests. Institutions that have not pre-positioned their exit options, sized for stressed unwind costs, and established governance that can act in hours rather than weeks will find conditional liquidity has become no liquidity.

16.4 Integration into total portfolio frameworks

The most constructive development in institutional QIS usage is the shift toward treating QIS as one component within an integrated portfolio framework rather than a standalone allocation. When QIS is sized, governed, and monitored within a framework that considers total portfolio liquidity, counterparty concentration, and stress dynamics, the structural risks are real but manageable. The product category is not broken. It is frequently implemented without the discipline its structural characteristics require. That gap between what QIS demands of its owners and what most governance frameworks actually deliver is the gap this paper has tried to close.

Conclusion

QIS is not misunderstood because it is complex. It is misunderstood because it is framed incorrectly.

It is sold as a strategy. It behaves as a structure.

The factor premia are real. The outcomes are determined by everything wrapped around them.

Two investors can access the same signal and experience completely different results. The difference is not skill. It is structure.

Allocators who understand this:

- Select structures deliberately, not for convenience
- Price total cost, not headline fees
- Size for stressed exit, not backtest volatility
- Govern at the speed the market requires, not the speed the committee cycle allows

> ***In QIS, implementation is the strategy. Structure determines the outcome. The signal is where the story starts – not where it ends.***

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/

Appendix A. The Five Ps for QIS Governance

Apply before committing capital to any quantitative investment strategy.

	Question	✓ Yes	✗ No	Key Test
PURPOSE	***Does this strategy have a specific, defined role in the portfolio?***	✓ YES Clear objective – a drawdown threshold, liquidity source, or factor premium with defined conditions for when it will succeed and when it will not.	✗ NO "Add alpha" or "diversify" without specifics. If you cannot state when it will fail, the purpose is undefined.	**Key test:** *Can you name the exact market conditions under which this strategy loses significantly?*
PURITY	***Is the factor exposure clean and independently observable?***	✓ YES Signal is replicable outside the dealer wrapper. Construction choices are disclosed, stable, and have not been retroactively optimised against backtest history.	✗ NO Black-box methodology. Construction parameters can change without disclosure. The strategy cannot be independently proxied or stress-tested.	**Key test:** *Can you replicate or independently proxy the exposure? If not, you cannot fully risk-manage it.*
PRICE	***Have you assessed the all-in cost – not just the explicit fee?***	✓ YES Total cost (fees + execution + funding + structuring margin) is known and proportionate to expected gross return.	✗ NO Fee schedule only. Execution drag, funding cost, and structuring margin have not been disaggregated. Fee-only comparison is meaningless.	**Key test:** *What is the all-in cost as a percentage of expected gross return? Above 40% warrants hard scrutiny.*
PATH	***Have you stress-tested the sequence of returns, not just the end-point?***	✓ YES Stressed exit cost modelled explicitly. Vol targeting mechanics understood. Entry timing and rebalancing sequence sensitivity assessed.	✗ NO Risk analysis based on historical Sharpe ratio and average return only. No stressed exit scenario has been constructed or reviewed.	**Key test:** *What does a stressed exit cost, and can the portfolio absorb it without forced selling elsewhere?*
PRACTICALITY	***Can your institution actually govern this investment at the right speed?***	✓ YES Managed account or equivalent transparency. Pre-defined drawdown triggers. Delegated exit authority that can act at market speed – not committee speed.	✗ NO Monthly committee decision cycle. Dealer marks accepted without independent challenge. No internal monitoring or replication capability.	**Key test:** *Who inside the institution can challenge a dealer valuation mark? If no one, governance is incomplete.*

✓ All five answered satisfactorily: proceed, with position sizing limits and governance triggers documented.

✗ Any gap identified: resolve before committing capital. A gap in Practicality is not a governance detail – it is a portfolio risk.