BNY | INVESTMENTS

COST, CONVEXITY AND THE CASE FOR SYSTEMATIC TAIL HEDGING

August 2026

Tail-risk hedging occupies a contentious position in institutional portfolio construction. Proponents see convex protection as a structural tool for improving compounded outcomes and preserving strategic risk-taking capacity. Critics see it as a persistent cost, difficult to justify through long quiet periods.

This paper sets out both sides of that debate before arguing that the perceived drawbacks are largely a function of how tail-risk hedging is implemented. A systematic, cost-aware approach, delivered through quantitative investment strategies (QIS), can manage those costs without eliminating the protection itself.

THE CASE FOR TAIL-RISK HEDGING

Path dependency and compounding. A 50% drawdown requires a 100% gain to recover. Convex protection changes the terminal wealth path, not just the average return. This matters for institutions with liabilities, or client bases prone to redeeming near the bottom of a drawdown.

Behavioural discipline. Hedged portfolios allow allocators to remain invested through a drawdown rather than de-risking at the worst possible moment. In this sense the hedge is as much a governance tool as a return driver.

Rebalancing optionality. Convex payoffs generate liquidity precisely when risk assets are cheapest, funding opportunistic rebalancing into the recovery. This is often the more valuable benefit of the strategy than the payoff itself.

Correlation regime shift. Traditional diversifiers such as bonds, gold and trend-following can fall simultaneously with equities during genuine liquidity or inflation shocks, as seen in 2022. Explicit convexity does not depend on historical correlations holding.

Sizing efficiency. A modest, permanent allocation to convex protection can support a larger strategic allocation to growth assets than an unhedged portfolio at the same overall risk budget. This is a risk transformation argument, not simply an insurance one.

THE CASE AGAINST TAIL-RISK HEDGING

Persistent negative carry. Options-based protection bleeds theta continuously. Over long non-crisis periods this is a real, visible cost that shows up in performance attribution.

Timing and path risk. Tail hedges can expire worthless repeatedly before the one event that pays off. This is the classic long-volatility fatigue that leads to programs being cancelled shortly before they are needed.

Basis and implementation risk. Hedges rarely match a portfolio's actual tail exposure precisely, whether that is index versus specific holdings, duration mismatch, or gap risk on entry and exit. Realised protection can therefore disappoint relative to back-tested expectations.

Opportunity cost. Critics argue that the same budget deployed in liquid alternatives, cash, or dynamic de-risking rules can achieve similar downside mitigation without the structural cost drag.

Governance and optics risk. A multi-year losing hedge is hard to defend to a board or client even when it is working exactly as designed. It is insurance that draws complaints until the fire arrives.

THE MYTH OF COSTLESS PROTECTION

The price of an option is, broadly speaking, the market's estimate of the expected value of its payoff. Strip out the cost and convexity goes with it, since there is no mechanism by which protection becomes genuinely free without also becoming less effective. This point cuts against a common but flawed argument: that a smarter structure ought to deliver the same tail protection at near-zero ongoing cost. That claim usually points to a structure that has quietly reduced the payoff, the trigger sensitivity, or the strike level, rather than a genuine improvement in efficiency.

A few mechanisms illustrate the trade-off. Moving further out of the money or converting a put into a put spread with a capped payoff reduces cost. However, by construction it also reduces the protection delivered in a genuine tail event, or caps it below the depth of the actual drawdown. Dynamically harvesting the volatility risk premium to fund cheaper convexity can reduce net cost too, but only up to a point, since the short-volatility leg funding the strategy can itself become a source of tail risk and reintroduce the correlation-breakdown vulnerability that the hedge was meant to solve. Systematic frameworks that reduce cost by only activating protection once a drawdown or volatility threshold is crossed leave a gap in which the first leg of a shock is unhedged by design. The zero-cost collar is perhaps the cleanest illustration: achieving zero net premium by capping upside participation does not eliminate cost, it simply relocates it from cash paid to upside forgone.

A related but distinct route to low-cost protection is to substitute explicit convexity with statistical diversifiers such as government bonds, trend following, gold, or low-volatility equity factors, on the reasoning that these have historically shown negative or low correlation to equities in drawdowns, at a fraction of the running cost of options-based protection. This is a legitimate diversification tool, but it is not the same thing as a hedge, and conflating the two is where the fragility hides. A long options position has a payoff defined by contract terms and does not depend on any historical relationship holding. A statistical hedge's protection lasts only as long as the observed correlation persists, and that correlation tends to break down exactly during the shock the hedge exists for. Genuine liquidity crises, inflation shocks, and rate-driven selloffs are precisely the environments in which multiple asset classes get sold simultaneously for the same reason. 2022 remains the clearest recent case: equities and bonds fell together, defeating the classic 60/40 diversification assumption at the moment it was most needed.

In both cases, an apparent cost saving is not a free lunch. It is an unpriced or mispriced risk that only becomes visible once an assumption fails. Statistical diversifiers carry little

or no explicit running cost, so they look unambiguously superior to options-based protection on a cost basis, but the comparison is incomplete: it sets the visible, continuous cost of one instrument against the invisible, contingent tail risk of another. Any defensive approach that appears to have solved the cost problem should prompt the question of what has been given up, or what assumption is being relied upon, to achieve that. Sometimes the answer is a legitimate, disclosed trade-off. Sometimes it is a correlation assumption doing the job an explicit contract should be doing, and that assumption tends to fail exactly when its failure is most costly.

A further variant, common in more complex tail-hedging constructions, is to fund the cost of protection by combining it with an uncorrelated source of positive carry drawn from other QIS. These sources of carry could include:

- Equity volatility carry, which systematically sells implied volatility to harvest the variance risk premium
- Cross-asset carry, exploiting roll yield and forward-rate differentials across foreign exchange, rates or commodities
- Skew or correlation risk premium harvesting
- A covered-call or put-write overlay running alongside the hedge book.

The logic is appealing on its face. If the carry source is genuinely uncorrelated to the hedge and to the tail event itself, it can offset the running cost of protection without touching the convexity being purchased.

The difficulty is that uncorrelated is a normal-markets property, and by definition most carry QIS are carry strategies precisely because they are short some form of risk or liquidity premium. Equity volatility carry is structurally short volatility. Cross-asset carry strategies are typically short liquidity or short a currency or rate differential. Skew harvesting is short the very tail risk the hedge exists to address. All these strategies tend to perform well in calm or moderately stressed conditions and worst in exactly the dislocation the tail hedge is designed to protect against, which is close to the definition of a risk premium rather than a coincidence. The correlation the strategy was funded on can compress or invert at the moment it matters most, so the carry sleeve loses money at the same time the hedge is paying off, quietly eroding the net benefit the investor was relying on the hedge to deliver. In more adverse cases, where the carry sleeve runs its own leverage or margin requirements, as short-volatility and skew-harvesting strategies commonly do, a loss there can force deleveraging or unwinds elsewhere in the portfolio precisely during the crisis window, working directly against the hedge's purpose rather than merely offsetting its cost.

This is a variant of the same theme running through this section: an apparently low-cost or self-funding structure is rarely getting something for nothing. Here the trade-off is not a capped payoff or a reliance on historical correlation holding, but a structural tension between the QIS carry source funding the programme and the event being hedged against. A genuinely disciplined programme can still use carry-generating QIS components, but only with clear-eyed recognition of when that carry source is likely to fail. The component should be sized so that its potential drawdown in a crisis does not offset the very protection it was meant to subsidise and never levered in a way that could compound the shock the hedge exists to address.

None of this weakens the paper's central argument; if anything, it sharpens it. The task for active, systematic management is not to engineer cost out of the equation, but to manage it efficiently. This means timing it against valuation, monetising it via fixed spreads and sizing it dynamically, while remaining explicit about what is being paid for and why. Cost, honestly priced and actively managed, is the signature of genuine convexity, not a flaw a well-run programme should be trying to eliminate.

RULES BEAT REFLEXES IN TAIL HEDGING

The tactical approach to tail-risk hedging, in which a portfolio manager, allocator or committee decides when to add or remove protection based on macro views or volatility levels, has an obvious appeal: judgment can react to information that a systematic, rules-based approach cannot see. But it imports the very behavioural fragility that tail hedging is meant to solve. Tactical hedging depends on someone correctly timing both entry and exit, and the evidence on discretionary market timing generally is not favourable. Once protection becomes discretionary, it becomes vulnerable to the same psychology that causes investors to de-risk at the bottom of a cycle: hedges get pared back after a long quiet period, just before they are needed, and added back after a shock, once they are expensive.

A SYSTEMATIC, QIS-DELIVERED APPROACH

A systematic, QIS-delivered approach addresses this fragility, in our view, through several concrete mechanisms.

Removal of timing discretion.
Rules-based triggers, such as volatility level, skew, realised-implied spread, or drawdown thresholds, govern sizing and rebalancing. This does not eliminate the risk of a wrong call. It relocates that risk upstream into strategy design and back-testing, where it can be stress-tested and governed, rather than being dependent on the day-to-day judgement of a portfolio manager under pressure.

Cost monetisation rather than pure cost bearing.
This is arguably the strongest technical argument for QIS delivery. A systematic framework can dynamically size the hedge ratio or shift structure, whether that is outright puts, put spreads, or skew-harvesting overlays. It can be based on the pricing of volatility itself, reducing exposure when the volatility risk premium is expensive and re-loading it when protection is cheap. This addresses the persistent negative carry critique directly: the carry is not eliminated, but it is actively managed rather than paid blindly and continuously.

Liquidity guarantees precisely when it is most valuable.
This is an advantage that is easy to overlook. In genuine stress events, the underlying options market is often exactly where liquidity disappears first: bid-ask spreads widen sharply, market makers withdraw quotes, and unwinding a physical options position, or sourcing fresh protection, can become difficult or punitively expensive at the moment it matters most. A QIS structure, typically delivered as a swap or index-linked note through a bank counterparty, carries a contractual liquidity guarantee, so the investor can monetise or unwind the exposure on the rule book's terms, independent of the live tradability of the underlying options market. The bank, rather than the end investor, absorbs the execution risk of sourcing or unwinding the options in the market.

Fixed spreads as a cost-effective monetisation mechanism.
Compounding the point above, QIS rule books typically price monetisation off a fixed, pre-agreed spread rather than a live, stressed market spread. At the exact moment a discretionary options seller would face the widest, most punitive bid-ask spreads, and therefore the worst cost of monetising protection, the QIS investor instead transacts against a spread fixed well in advance of the crisis. The strategy's convexity is realised more efficiently, and more cheaply, than an equivalent physical options position would allow, precisely when that efficiency matters most.

Capital and operational efficiency.
Delivered via a swap or index-linked structure through a bank counterparty, a QIS approach avoids the balance-sheet and operational drag of running a physical options book internally.

Governance and explainability.
A rules-based framework is fully transparent and replicable to a risk committee or client: it is possible to show precisely why a hedge ratio changed, rather than relying on a portfolio manager's shift in view. This counters the governance and optics drawback set out above. A multi-year losing hedge can be considerably easier to defend when it is operating exactly as a disclosed, tested rule set intended, rather than reflecting an individual's discretionary call.

Consistency of risk budget.
Because sizing is rule-driven rather than sentiment-driven, the strategy maintains a more stable risk budget through the cycle. This avoids the classic failure mode where tactical hedges are undersized precisely when correlations break down, which tends to be exactly when a discretionary overlay is most likely to have been scaled back on the assumption that diversification is working fine.

SYSTEMATIC EXECUTION, ACTIVE OVERSIGHT

A genuinely novel regime – a structural break the model was not designed around – is the one scenario in which tactical judgement can add value that a rules engine cannot. A well-constructed programme does not treat this as an unresolved weakness of systematic delivery. Instead, it addresses the issue directly, by combining a systematic core with active portfolio management oversight.

In this hybrid design, the systematic framework continues to govern day-to-day sizing, rebalancing and cost monetisation, removing the timing discretion and behavioural fragility that undermine purely tactical hedging. The portfolio management team's role is narrower and more precisely defined: to identify, recognise and react to genuine structural breaks – that is regime shifts the model's assumptions were not built around – where continuing to apply historical rules mechanically would reduce effectiveness. This is fiduciary oversight of the systematic engine, not a replacement for it.

This design avoids reintroducing the fragility set out above. The portfolio management team is not making forecasting calls on the future direction of volatility, nor timing entries and exits. That discretion, and its associated behavioural risk, remains removed. Oversight is instead reserved for a narrower and rarer judgement: recognising when the environment itself has changed in a way the systematic rules did not anticipate, and adjusting the framework's parameters or constraints accordingly. Any such adjustment would be subject to the same governance and documentation standards as the systematic rules themselves.

The resulting programme captures the discipline, cost-efficiency and governance benefits of systematic delivery, while retaining a safeguard against the one failure mode a rules-based approach cannot resolve unassisted. Systematic execution combined with active fiduciary oversight is a more complete answer to the tail hedging debate than either pure discretion or pure systematisation on its own.

A MORE COMPLETE APPROACH TO TAIL-RISK HEDGING

The drawbacks most commonly levelled at tail-risk hedging, notably persistent cost, poor timing and governance discomfort, are largely a function of how protection is implemented, not an inherent property of hedging itself. A systematic approach, delivered through QIS, converts tail-risk hedging from a discretionary insurance bet into a disciplined, cost-aware risk-transformation tool. With this strategy, costs can be actively managed rather than borne blindly, liquidity is contractually guaranteed at the moment it is most valuable, and the framework remains transparent and defensible to governance stakeholders.

The most complete version of this approach is not purely systematic but systematic plus oversight: a rules-based engine handles the day-to-day discipline that removes behavioural fragility, while a portfolio management team retains a narrowly defined fiduciary mandate to recognise and respond to genuine structural breaks the model was not built to anticipate. We believe that the benefits of tail-risk hedging are not diminished by its costs. They are best realised by managing those costs intelligently, within a framework that combines systematic rigour with active oversight, rather than attempting to remove the costs altogether or ceding the entire decision either to pure discretion or pure automation.

Terminology

Tail-risk hedging. A strategy designed to protect a portfolio against rare but severe market declines or extreme events. It often involves instruments like put options or other positions that gain value during market stress.

Convexity. Convexity is the non-linear, curved relationship between an option's price and the underlying asset's price. Unlike stocks, where profits and losses move in a straight 1:1 line, an option's profit accelerates when the market moves in your favour, representing positive convexity.

Theta. An options Greek that measures how much an option's value declines as time passes, assuming other factors stay constant. For option buyers, theta is usually negative because time decay erodes option value.

Volatility risk premium. The difference between the implied volatility priced into options and the actual volatility that ultimately occurs. Investors who sell options often seek to capture this premium as compensation for bearing volatility risk.

Diversification. The practice of spreading investments across assets, sectors, or regions to reduce concentration risk. The goal is to lower overall portfolio volatility without necessarily reducing expected returns proportionally.

Options. Financial contracts that give the buyer the right, but not the obligation, to buy or sell an asset at a specified price before or at expiration. Calls provide the right to buy. Puts provide the right to sell.

Correlation. A statistical measure of how two assets move in relation to each other. A correlation of +1 means they move together perfectly, 0 means no consistent relationship, and -1 means they move in opposite directions.

Skew. In options markets, skew refers to the difference in implied volatility across strike prices or maturities. It often reflects greater demand for downside protection, causing put options to trade with higher implied volatility than calls.

Negative carry. The cost of holding an investment when its financing, hedging, or income characteristics produce a net loss over time. For example, buying protection or holding a low-yield asset funded at a higher rate can create negative carry.

Governance. The framework of rules, oversight, and decision-making processes used to manage an investment strategy or portfolio. It includes roles, accountability, risk controls, and policies that guide how decisions are made.

Roll yield. The return generated when a futures position is rolled from a near-dated contract into a later-dated one. It can be positive or negative depending on whether the market is in backwardation or contango.

Covered call. An options strategy where an investor holds an asset and sells a call option on that same asset. It generates income from the option premium but limits upside if the asset rises above the strike price.

Put-write overlay. A strategy that involves selling put options, usually on an equity index or portfolio, to generate premium income. It can enhance yield, but it also creates downside exposure if the underlying asset falls significantly.

Quantitative investment strategy. A strategy that uses mathematical models, statistical analysis, and data to identify investment opportunities and manage risk. It relies on systematic rules rather than subjective judgment to make portfolio decisions.

Important information

MARK-989205-2026-08-19

GU-929-31 August 2027

BNY | INVESTMENTS