RETHINKING TAIL RISK

Why Convexity Is Not an Insurance Product, but a Strategy

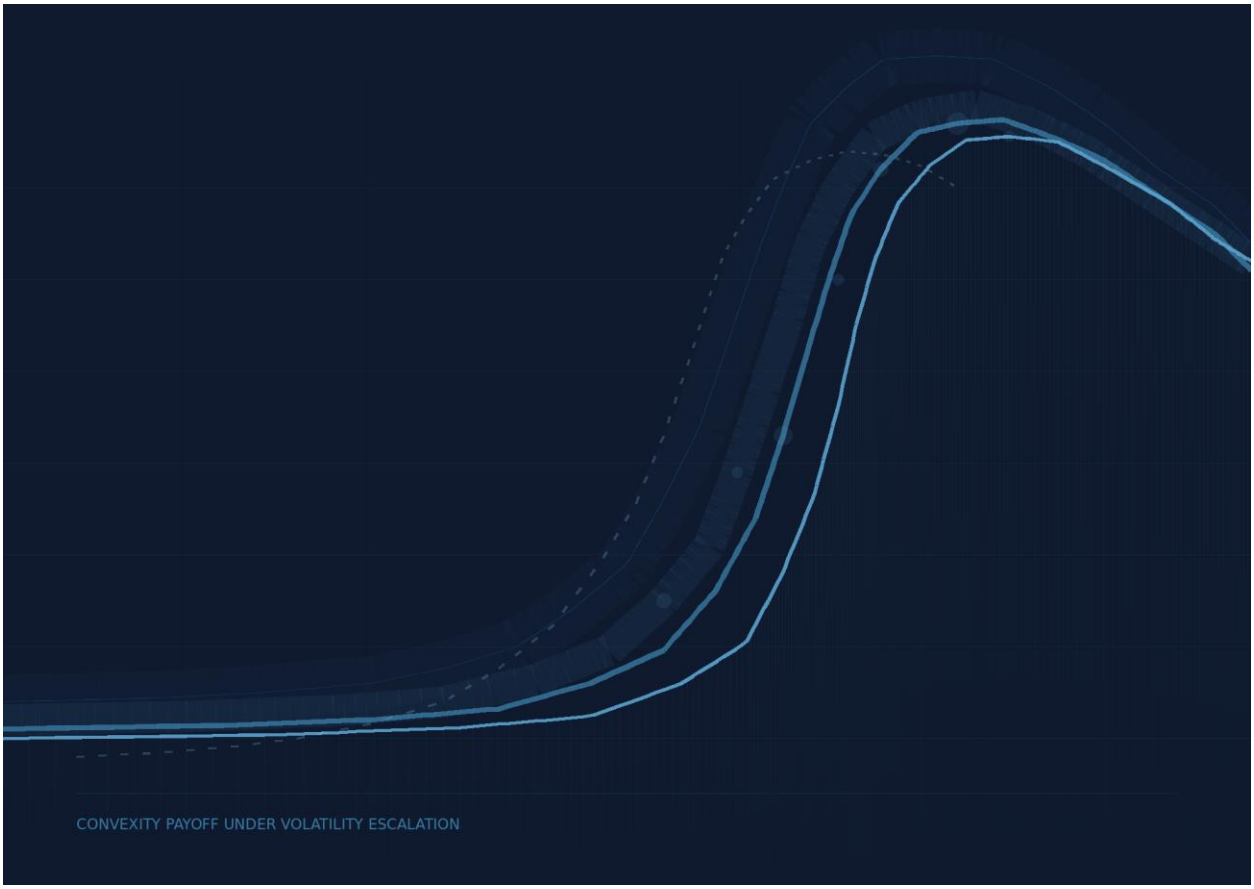

Wolfgang Landl

Dubai – March 2026

FOREWORD

Why This Article Was Written

Three distinct groups of institutional investors encounter tail risk strategies. All three, in the end, decline to invest – each for entirely different reasons.

The first group lacks familiarity with volatility as an asset class. Options, volatility surfaces, convexity – these concepts sit outside their daily toolkit. The discomfort is genuine, though seldom stated openly. Instead, the response tends to be: “The timing isn’t right” or “We’re monitoring the space.” What they really mean is that they do not yet understand the subject well enough to commit capital. That is not a weakness – it is honest caution, and it deserves respect. But it also deserves an answer.

The second group has experience – sometimes too much of it. They have invested in volatility strategies before, lost money, and concluded that the entire concept is flawed. What went unexamined was whether it was the concept that failed or the implementation. An investor who buys a return-seeking strategy marketed as a hedge should not be surprised when protection is absent during a crisis. The bad experience was real. But it was not evidence against convexity – it was evidence against mislabelling.

The third group is today the largest. These are investors who have watched, over years and in some cases decades, as markets reached new highs again and again despite every disruption. Fear has largely disappeared from the language of investment committees. The prevailing concern is FOMO: the fear of missing the next rally, of appearing too defensive, of leaving returns on the table. In this environment, hedging feels like a drag – psychologically, from an accounting perspective, and politically.

Underlying all three groups is a deeper paradox. Consider the decisions investors make routinely and without hesitation: buy a car, purchase car insurance; buy a house, purchase home insurance. Nobody considers these conversations difficult. They are simply responsible asset protection. Yet an investor who would never dream of driving uninsured or leaving a home unprotected will hold a multi-million-dollar portfolio with no downside protection whatsoever. The same individual who pays two thousand dollars annually to insure a fifty-thousand-dollar vehicle will balk at “spending” a percentage to protect a multi-million-dollar investment portfolio.

This inconsistency reveals not a rational cost-benefit calculation, but a cognitive bias. *Car and home insurance protect against visible, tangible risks. Portfolio losses, by contrast, feel abstract until they materialise* – at which point protection becomes prohibitively expensive or simply unavailable. The difficulty in placing tail hedges is not inherent to the product. It stems from recency bias and the false comfort of low-volatility regimes. The very conditions that make

protection most affordable are precisely the conditions that create the greatest psychological resistance to purchasing it.

This is precisely the moment to pause.

Not because a crash is imminent. But because the quietest periods in capital markets have historically represented the most attractive hedging windows – and simultaneously the moments when hedging is easiest to implement: cheap, unobtrusive, and free from pressure. Investors who think about portfolio structure during these periods are not acting defensively. They are acting intelligently.

The real question is not *whether* to hedge. The question is: *with what – and at what cost.*

This article attempts to answer that question. Not with a product pitch, but with a framework for thinking. Because investors who understand how convexity truly arises – and why most approaches fail structurally before the first crisis ever arrives – make better decisions. Regardless of which group they belong to.

Wolfgang Landl

Dubai, March 2026

RETHINKING TAIL RISK

Why Convexity Is Not an Insurance Product, but a Strategy

Extreme market events are still widely treated as anomalies. Yet the record of recent decades paints a different picture: financial crises, pandemic-driven shocks, geopolitical escalations, and monetary policy regime shifts occur far more frequently than classical risk models would suggest.

Institutional investors have responded by adopting so-called tail risk strategies in growing numbers. And yet, experience shows that many of these approaches disappoint over the long run. They are either too costly to sustain – or they fail to perform precisely when performance matters most.

The reason lies not in the instruments, but in the conceptual framework: tail risk is not an insurance problem. It is a convexity problem.

1. Diversification Reaches Its Limits in Extreme Markets

Diversification is the foundation of professional asset management. In stress periods, however, it loses much of its power. Correlations spike, liquidity evaporates, and the dominant risk driver – equity exposure – overwhelms everything else.

Empirical analysis consistently shows that in standard balanced portfolios, up to ninety percent of total risk is attributable to equities. In precisely those phases where protection is most urgently needed, diversification delivers only limited stability.

The implication is straightforward: supplementary mechanisms to cushion extreme losses are warranted – **provided they are economically viable.**

2. Why Conventional Tail Hedges Rarely Deliver

The most common hedging approach involves the systematic purchase of deep out-of-the-money put options. In theory, these instruments offer strongly convex payoffs and positive skew. In practice, they carry significant structural disadvantages.

Options are expensive over the long run. Equity markets rise on average, and volatility is typically priced at a premium. Buying protection continuously means paying that premium continuously. Factor in timing decisions, maturity selection, and the psychological burden of years of unrewarded costs, and the picture darkens further.

The outcome is predictable: many investors abandon their hedges shortly before they are needed. The strategy fails not because of the instrument, but because of its **lack of sustainability.**

3. Volatility Is Not a Static Risk Factor

A central conceptual error lies in treating volatility as a passive, stable input. Volatility is neither stable nor linear. It is highly path-dependent, responding to market structure, liquidity conditions, and shifting expectations – and it often moves in abrupt, discontinuous jumps.

Effective tail risk implementation therefore resembles not an insurance policy but a **professional volatility trading desk:**

- Dynamic adjustment of strikes and maturities
- Exploitation of mispricings along the volatility surface
- Consideration of liquidity, dealer positioning, and market stress
- Consistent price and timing discipline

The decisive difference lies not in the instrument but in the **manner of execution.** Or to put it more bluntly: ***structures are not strategies.***

4. The Gradual Nature of Volatility Escalation

A persistent misconception in financial markets is that volatility spikes appear suddenly, without warning. The empirical record tells a fundamentally different story: volatility typically builds gradually over weeks or months before reaching crisis levels. This observation has profound implications for how portfolio protection should be designed, timed, and governed.

The Empirical Evidence: Volatility Does Not Appear Overnight

The CBOE Volatility Index (VIX) – derived from S&P 500 index options and widely regarded as the market's forward-looking fear gauge – reveals a consistent signature across major crises: a progressive escalation from baseline levels to extreme readings, rather than instantaneous jumps.

The 2008 Global Financial Crisis: A Fifteen-Month Escalation

The 2008 financial crisis exemplifies the gradual buildup pattern most clearly. Market volatility began its ascent well before the September 2008 Lehman Brothers collapse that is commonly identified as the crisis trigger:

Period	VIX Behaviour
2007 baseline	VIX averaged 12–15, consistent with the low-volatility regime of 2004–2007
Mid-2007 initial stress	VIX moved into the 20s as subprime mortgage concerns emerged
September 2008 acceleration	Following Lehman's bankruptcy, VIX surged into the 30–40 range
October–November 2008 peak	VIX reached its all-time high of 80.86 on November 20, 2008

This progression spans approximately fifteen to eighteen months from initial elevation to maximum panic. Research by Manda (2010) demonstrates that volatility levels rose to 43.6 percent during the crisis period, compared to pre-crisis levels of 13.4 percent – a transformation that occurred gradually rather than instantaneously.

The COVID-19 Crisis: A Compressed but Observable Timeline

The COVID-19 market disruption in early 2020, while more compressed, still exhibited discernible warning signals:

Period	VIX Behaviour
January 2020	VIX at 13–15 (normal baseline levels)
Late February 2020	VIX climbing into the 20s as pandemic concerns intensified
March 16, 2020	VIX peaked at 82.69 – the highest reading in the index's history

Even in this accelerated timeline, the progression from baseline to peak volatility took six to eight weeks. The peak was preceded by weeks of observable escalation that provided a window for action.

March 2026: A Real-Time Illustration

The geopolitical shock following US–Israeli strikes on Iran beginning February 28, 2026, provides evidence of the same pattern unfolding in real time. Market tensions had been building through February, with the VIX closing at 19.86 on February 27. Following the weekend strikes, the VIX surged substantially, reaching highs above 35 in early March – a significant increase from the 16–17 range seen just weeks earlier.

The trajectory remains consistent with the familiar pattern of gradual escalation rather than instantaneous panic. Whether this episode develops into a full crisis or mean-reverts from current levels, the structure of the buildup follows the same logic observed in every major dislocation of the past two decades.

Why Volatility Clusters and Mean-Reverts

The gradual nature of volatility buildup is rooted in two fundamental properties of financial markets.

Volatility clustering, first documented by Mandelbrot (1963) and formalised through ARCH/GARCH models by Engle (1982) and Bollerslev (1986), describes the empirical observation that large changes tend to be followed by large changes, and small changes by small changes. While returns themselves exhibit minimal autocorrelation, absolute returns display positive and slowly decaying autocorrelation over horizons ranging from minutes to several weeks. This persistence explains why crisis episodes extend over weeks and months rather than resolving immediately. During volatile periods, market makers widen spreads, reduce position sizes, and may withdraw entirely. Algorithmic traders often reduce activity or shut down, creating liquidity vacuums that amplify price swings and extend the duration of volatility clusters.

Mean reversion operates simultaneously: volatility tends to return to historical average levels over extended horizons. This property creates a natural lifecycle for volatility episodes – gradual buildup, peak intensity, and eventual decay back toward baseline levels. Extreme volatility does not persist indefinitely, but neither does it resolve instantly. The reversion process operates over weeks or months, mirroring the pace of the initial escalation.

Underlying both dynamics is the progressive nature of information processing in financial markets. Major crises rarely emerge from single events. They develop as accumulating evidence forces a step-by-step reassessment of systemic risk. The subprime problems behind the 2008 crisis were visible to astute observers as early as 2006. The COVID-19 pandemic evolved from a regional health concern in January 2020 to a recognised global threat by late February. Each stage of reassessment corresponded to measurable increases in volatility.

Implications for Portfolio Protection

The “frog in boiling water” problem. Because volatility rises gradually, portfolio managers often fail to recognise developing crisis conditions until substantial damage has occurred. The progression from a VIX of 15 to 25 feels manageable and temporary. By the time it reaches 40 or higher, the cost of protection has multiplied and portfolio losses may already be severe.

The cost-timing paradox. Tail hedge instruments – out-of-the-money puts, volatility products – become progressively more expensive as volatility rises. The optimal time to establish protection, when volatility is low and insurance is cheap, is precisely when it appears least necessary. By the time crisis conditions are obvious to all participants, protective instruments have already appreciated substantially.

The visibility window. The six-to-eight-week minimum to fifteen-month-plus timeline for volatility escalation provides a window for informed action. Investors who actively monitor volatility

indicators can observe developing stress conditions and adjust risk exposure or establish protective positions before peak crisis levels are reached. But only those with structures already in place can act on this window. Those who wait for committee approval to begin the process of hedging will find that by the time the decision is made, the opportunity has passed.

The time to buy insurance is before the fire, not during it. The gradual buildup of volatility provides a visibility window – but only for those already positioned to act.

5. Why Conventional Tail Hedge Products Fail Structurally

Most institutional tail risk solutions follow a familiar pattern: periodic purchase of deep out-of-the-money options, fixed maturities and strikes, rules-based rebalancing, and historical optimisation calibrated to past crises.

These approaches appear sensible at first glance, but they are undermined by several **structural headwinds:**

Equity Risk Premium (ERP): Equity markets rise over the long term. Passive hedges are therefore systematically positioned against the dominant return driver.

Volatility Risk Premium (VRP): Options are, on average, overpriced relative to realised volatility. Continuously buying protection means continuously paying a premium to the market.

The No-Alpha Dilemma: Without a genuine trading edge, the expected value of the hedge is negative. Structure alone does not substitute for strategy.

Variance Drag: Years of capital erosion substantially reduce the economic value of even large one-off gains. A strategy that compounds negatively for long stretches must generate extraordinary returns in crises simply to break even.

The result: superficially reassuring but economically inefficient protection that costs more than it delivers over the long run.

6. Which Strategies Actually Deliver Convexity?

Not every defensive strategy generates convexity, and not every form of convexity is sustainable. A review of the main categories reveals critical distinctions.

Category 1: Return Strategies Disguised as Hedges

These strategies implicitly sell risk and collect premiums. They perform well in benign environments – but in genuine crises they fail, or worse, amplify losses.

Risk Premia and Alternative Risk Premia Funds systematically harvest established premia such as the volatility risk premium, carry, and momentum. The problem: in stress periods, these premia collapse simultaneously. They function only in calm markets and offer no meaningful protection when it counts.

Volatility Arbitrage and Dispersion Trading aim to exploit relative mispricings between volatility instruments. The approach sounds compelling, but it is fundamentally unsuited as a hedge. There is no defined positive payoff in a crisis – returns depend on spread convergence, not on the severity of the drawdown.

Managed Volatility and Low-Volatility Equity reduce portfolio volatility by rotating into lower-beta equities. In systemic events, however, all equities decline – including low-volatility names, as March 2020 demonstrated. The approach dampens day-to-day swings but does not protect against genuine tail events.

Risk Parity balances risk contributions across asset classes. When correlations surge – as they invariably do in crises – the diversification premise breaks down entirely. Passive risk balancing offers no asymmetric payoff.

VIX Roll-Yield Strategies (Long VIX Futures) appear to be the most intuitive hedge, but they are among the most expensive. The VIX futures curve is structurally in contango, imposing daily roll losses of five to ten percent per month in calm periods. The strategy destroys capital long before any crisis arrives.

Category 2: Structurally Asymmetric and Convex

These strategies possess a mathematically defined, non-linear payoff. They lose little in normal markets and gain disproportionately in crises.

Long OTM Index Puts (e.g. S&P 500) represent the purest form of convexity: defined maximum loss equal to the premium paid, and unlimited upside. Purchased passively and mechanically, they bleed steadily. Managed actively and with discipline, they become highly effective. The payoff rises non-linearly with the severity of the crisis.

Long OTM VIX Calls exploit the tendency of the VIX to spike from levels of fifteen to forty or even eighty during dislocations, generating massive gains on call positions. Strike selection is critical: a one-month VIX 30 call can return fifty percent in one scenario or three hundred percent in another, depending on positioning and market microstructure. The payoff is directly linked to rising volatility.

Long Variance Swaps profit when realised volatility exceeds implied volatility – a condition that typically obtains during crises. The instrument is inherently convex to the stress event, with gains that accelerate as disruption deepens.

Tail Risk Funds with Active Management combine long OTM puts with active timing and strike selection, often operating a two-sleeve structure in which flow trading offsets option premiums. This is the approach that delivers maximum convexity with minimal sustained bleed.

Between the Categories: CTAs and Trend-Following

CTAs react more slowly to market breaks but deliver robust positive skew over extended stress periods, as demonstrated in 2008. In sudden crashes – a flash crash, or the rapid COVID selloff of March 2020 – they often underperform or incur losses. Trend-following serves well as a structural diversification component but does not substitute for short-duration hedging.

Convexity is not a product feature. It is the result of active, price-disciplined trading.

---

7. From Tail Hedging to Strategic Convexity

For institutional investors, a shift in perspective is overdue. Rather than deploying isolated hedging products, convexity should be understood as a strategic resource – governed through an active convexity budget.

The objective is not the elimination of risk. It is the deliberate introduction of asymmetric payoff profiles that provide liquidity in stress periods, preserve the capacity to act, and limit losses – without imposing excessive costs over the full cycle.

In practice, this typically corresponds to an allocation of two to four percent of total assets, maintained on a long-term basis and clearly embedded within the portfolio's governance framework.

The empirical evidence on volatility escalation reinforces this point. Because crises build gradually, a standing allocation to convexity – one that is already in place when the first signals appear – is fundamentally different from a reactive decision to hedge after the fact. The visibility

window documented in Section 4 is only useful to investors who already have the infrastructure, the mandate, and the positions to act on it.

8. A Practical Example: What Separates a Long-Vol Position from an Expensive Lottery Ticket

How does an active approach systematically address the structural problems described above – and exploit the visibility window that gradual volatility escalation provides? Using the Ambrus process as an illustration, four decisive filters can be identified that govern every positioning decision.

Filter 1: Value Assessment – Buy Only at Dislocation

Most solutions-based funds purchase options irrespective of the prevailing price level – on the first Monday of the month, regardless of whether implied volatility stands at twelve or thirty-five. The Ambrus approach instead employs a five- to ten-year lookback on floating-strike volatility. A position is initiated only when a strike appears historically cheap – that is, when a genuine dislocation is present.

Practical significance: the purchase price of protection directly determines the bleed. A VIX call bought at an implied volatility of eighty following a crash is typically prohibitively expensive. The same call purchased at an implied volatility of fifteen during a calm market is structurally cheap. The gradual buildup of volatility documented in this paper means that the cheapest entry points occur precisely during the low-volatility regimes that precede crises – not during the crises themselves.

Filter 2: Sensitivity to Spot Declines – Buy Only What Responds

The ***"dead vol" problem:*** an option can be technically positioned on a cheap strike and still fail to respond when the market falls.

Consider a concrete example: a one-month VIX 30 call can deliver a gain of fifty percent in one scenario – and three hundred percent in another, given the same underlying move. The sixfold difference arises entirely from market microstructure: asymmetric demand driven by corporate hedging programmes, or dealer positioning that is offside, can transform the same option into a radically different payoff.

Ambrus analyses market microstructure before committing capital. Sensitivity must be verified – and this is proprietary knowledge that develops only through daily, active market participation.

Filter 3: Carry Profile – Actively Manage the Bleed

A poorly chosen structure can lose half its vega exposure in a single day – without any movement in the underlying market. Through targeted analysis of the term structure and skew curve, Ambrus positions to occasionally generate positive carry along the curve, producing returns even absent a significant move in spot prices.

This is the foundation of the ***bleed-equals-zero philosophy:*** the market's time structure pays for the protection, rather than the investor continuously burning through premiums.

Filter 4: Execution Efficiency – Convexity Is Defined at the Point of Purchase

The convexity of an option is not a fixed property – it is defined at the purchase price. Institutional options are frequently traded through brokers with wide spreads and no limit-order capability. A cost of five to fifteen basis points per transaction, compounded across hundreds of trades per year, creates meaningful drag.

The Ambrus approach – informed by the culture of a Chicago proprietary trading desk – relies on precise price targeting through limit orders, low-latency infrastructure, adaptive execution across multiple exchanges, and the posting of liquidity to capture exchange rebates. The objective is not merely to control costs but to neutralise them entirely and generate a genuine alpha contribution over time.

Value → Sensitivity → Carry Profile → Execution. A solutions-based fund applies none of these filters. An active fund applies all four – and refrains from buying altogether when the criteria are not met.

---

9. Governance Over Product Labels

For pension funds and other institutional allocators, the category to which a strategy belongs matters far less than its portfolio-level economic effect. Three questions should be central to any evaluation:

1. How does the strategy behave in stress periods?
2. What are the running costs over the full market cycle?
3. Does it deliver liquidity precisely when it is needed most?

Convexity is not a by-product. It is a deliberate **design decision.**

Conclusion

The widespread claim that tail risk hedging is inevitably expensive and inefficient is less a market law than a consequence of poor implementation.

The empirical evidence is equally clear on a second point: volatility does not spike without warning. Major market disruptions across different decades and crisis types – the 2008 financial crisis, the 2020 COVID-19 shock, and the developing geopolitical tensions of early 2026 – demonstrate measurable buildup periods ranging from weeks to months. This gradual escalation reflects fundamental properties of financial markets: volatility clustering, mean reversion, and the progressive nature of information processing.

For portfolio managers and risk officers, these two realities – the structural failure of passive hedging and the observable buildup of volatility – converge on a single conclusion: tail hedge strategy must shift from reactive crisis management to proactive risk monitoring, supported by a standing allocation to actively managed convexity.

Convexity exists – but it cannot be passively consumed. It must be actively traded, continuously adjusted, and disciplined in its monetisation.

Tail risk is not an insurance problem. It is a trading problem.

The warning signals exist. The question is whether market participants will recognise them before the progression from manageable volatility to crisis-level stress completes its inevitable cycle.

SOURCES

Hedging Tail Event Risks – a "Must" in a Diversified Portfolio? – Wolfgang Landl – December 2024

Marketing vs. Math: Solutions-Based Tail Risk Hedging vs. Alpha-Focused Tail Risk Hedging – Ambrus Capital – June 2025

Gradual Nature of Volatility Escalation: Understanding Pre-Crisis Market Dynamics – Wolfgang Landl – March 2026

Academic References

Albulescu, C. T. (2020). COVID-19 and the United States financial markets' volatility. Finance Research Letters, 38, 101699.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics, 31(3), 307–327.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. Econometrica, 50(4), 987–1007.

Manda, K. (2010). Stock market volatility during the 2008 financial crisis. Glucksman Institute for Research in Securities Markets, NYU Stern.

Mandelbrot, B. (1963). The variation of certain speculative prices. The Journal of Business, 36(4), 394–419.

DISCLAIMER

The information and opinions contained in this paper have been researched, compiled, and written by the author with the utmost care. This document is provided for informational purposes only and is intended exclusively for Professional Clients and Market Counterparties. It must not be relied upon by any other person.

Any opinions expressed may change without prior notice and do not constitute investment advice. The author does not warrant as to the accuracy or completeness of any market prices, data, or information contained herein, all of which may be subject to change at any time without notice. Neither the author nor any associated firm shall have any liability for any direct, indirect, consequential, or special losses or damages arising from or in connection with the information contained herein or any reliance on any of the content.

Investors should conduct their own due diligence and not rely on the financial assumptions, estimates, or regulatory representations presented in this document. Nothing herein constitutes investment, legal, accounting, or tax advice, nor a representation that any investment or strategy is suitable or appropriate for individual circumstances. The author recommends that investors independently assess, with a professional advisor, the specific financial risks as well as legal, regulatory, credit, tax, and accounting consequences of any investment decision.

Past and simulated performance are not a reliable guide to future performance. Investments can go down as well as up and investors may receive back less than their original investment.