LIGHTFLUX-FX

RESEARCH PAPER

—

An Adaptive AI Framework for Systematic Alpha Extraction in G10 FX

with Hard, Pre-Defined Risk Envelopes

A rigorously documented framework for ensemble-based, regime-aware FX execution.

Strategy Line: LightFlux-FX · Managed Account on Swissquote Capital Markets

Trading Mentor: Motherhub S.L. (Spain)

2026

Abstract

Currency markets do not move in straight lines. They breathe, dislocate, mean-revert, and occasionally rupture — and they do so on multiple time scales simultaneously. Anyone deploying capital in the world's largest financial market faces a fundamental choice: accept the carry-trend-value triad as the only available risk premia, or attempt to extract structural alpha through systematic, regime-aware execution. The present paper analyses a framework of the second type.

The LightFlux-FX strategy combines a PhD-engineered decision framework with an adaptive machine-learning evaluation layer. At its core is an ensemble of 196 independent decision nodes — each trained on a specific currency pair, market regime, and predictive horizon, and continuously re-weighted based on rolling empirical performance. Capital is distributed across many small, bounded exposures rather than concentrated directional bets, producing a return profile that is structurally low-beta ($\beta = 0.096$ versus the S&P 500) and largely uncorrelated with traditional risk assets.

This paper describes the strategy's architecture across four layers: the universe of 28 G10 currency pairs and why this universe was chosen; the operational rule set governing signal generation, position sizing, and node orchestration; the pre-budgeted monetary loss envelope that defines risk ex ante rather than ex post; and the academic literature on systematic FX, ensemble methods, and machine-learning-driven alpha that motivates the design. We then examine three live-track stress windows — the August 2024 Yen carry-trade unwind, the April 2025 Trump tariff shock, and the January 2026 DeepSeek AI shock — and benchmark the strategy's behaviour against equity indices in each. We close with an honest assessment of what the framework cannot do, because transparency about limitations is the price of intellectual seriousness.

The objective is not to present a crystal ball but a transparent, auditable rule set that offers structural advantages in specific market environments — and whose boundaries are stated as clearly as its strengths.

1. Research Background: Why Systematic FX?

Foreign exchange is the largest and most liquid market in the world. The 2025 BIS Triennial Central Bank Survey reported average daily turnover of USD 9.6 trillion in over-the-counter FX, up 28 % from USD 7.5 trillion in 2022 and roughly 80 % higher than a decade ago. Spot turnover alone increased 42 % between 2022 and 2025, and the share of FX swaps continued to dominate at 42 % of total flow.

> ***Source:*** *Bank for International Settlements (2025), 2025 Triennial Central Bank Survey of Foreign Exchange and OTC Derivatives Markets, BIS Statistics.*

Yet this enormous market is structurally inefficient by design. As Levich and others have repeatedly observed, the majority of FX transactions are not executed for profit. Corporate hedging, central-bank intervention, balance-sheet management, and cross-border M&A together drive a substantial fraction of daily flow — and none of these participants optimise for return. The market is liquid enough to absorb informed flow without slippage, and inefficient enough that statistical regularities persist long enough to be exploited.

This combination — depth, liquidity, and persistent inefficiency — is precisely what makes FX an attractive arena for systematic, rule-based strategies. It is also what makes the discipline unforgiving: edges that survive in less efficient markets (small-cap equities, certain commodity futures) are rarely persistent in major FX pairs, and edges that do persist tend to be small, fragile, and statistically noisy.

1.1 The Currency Manager Alpha Question

The seminal work on whether currency managers actually generate alpha — as opposed to harvesting style premia — is the four-factor framework of Pojarliev and Levich. Examining the Barclay Currency Traders Index and 34 individual currency funds over 1990–2006, they decomposed returns into four systematic factors: carry, trend, value, and currency volatility. These four factors explained the substantial majority of variability in fund returns. After accounting for them, the average BCTI alpha was negative 9 basis points per month — meaning that the average professional currency manager underperformed a passive style-factor portfolio.

> ***Source:*** *Pojarliev, M. & Levich, R. M. (2008): "Do Professional Currency Managers Beat the Benchmark?" Financial Analysts Journal, 64(5), 18–32.*

This is the hard truth that any serious FX strategy must confront: a portfolio that loads on carry, trend, value, or volatility is not delivering alpha — it is delivering a beta-style exposure that can be replicated for a fraction of the management fee. True alpha in currency, in the academic sense, is residual return after stripping out these factor exposures. Pojarliev and Levich found that roughly one in four professional managers did generate such residual alpha in their sample; the rest were paid alpha fees for what was effectively repackaged beta.

LightFlux-FX is explicitly designed to operate in the residual-alpha space rather than the factor-beta space. Its return signature shows near-zero loading on the equity market ($\beta = 0.096$ versus the S&P 500, $R^2 = 0.8$ %) and avoids structural exposure to any single FX factor: positions are short-horizon (median 43.3 hours), distributed across 28 pairs, and bounded by hard monetary loss envelopes that prevent any individual signal from delivering meaningful directional return. Whether this approach succeeds at generating sustained residual alpha is, of course, a question that only out-of-sample evidence can answer — and the strategy's current 25-month live track is too short for definitive conclusions.

1.2 From Single Models to Ensembles: The Wisdom of Crowds in Quant Finance

The intellectual foundation of LightFlux-FX is the ensemble method — the principle that aggregating predictions from many diverse, independently-trained models tends to outperform any single model. The principle is older than machine learning. In 1785, the Marquis de Condorcet showed that if each member of a jury has a greater-than-50 % probability of voting correctly, the probability of a correct majority vote approaches certainty as the jury grows. The journalist James Surowiecki popularised the same insight as the "wisdom of crowds" in 2004.

Source: *Surowiecki, J. (2004): The Wisdom of Crowds. Doubleday, New York.*

In machine learning, this translates into the bias-variance decomposition. A single complex model — say, a deep neural network trained on FX returns — can fit historical data with high precision but generalises poorly: it has low bias but high variance. Aggregating predictions from many simpler, diverse models reduces variance through averaging while preserving low bias, provided the underlying models make uncorrelated errors. The condition is decisive: ensemble gains depend on the models being meaningfully diverse, not on the number of models alone. An ensemble of identical models offers no improvement over any one of them.

LightFlux-FX operationalises this principle through 196 decision nodes, each trained on a specific market context — a particular currency pair, a particular market regime (defined by realised volatility, term-structure curvature, and macro factors), and a particular predictive horizon ranging from a few hours to several days. The nodes share a common architectural blueprint, but their training data, feature sets, and signal-generation logic differ enough that their forecast errors are largely uncorrelated. The aggregate signal — the actual position the system takes — is the weighted sum of these node-level forecasts, with weights determined by each node's rolling empirical performance.

1.3 The Backtest-Overfitting Problem

Any paper that presents a quantitative strategy with attractive historical returns invites a fundamental question: how much of this performance is real, and how much is statistical artefact? The work of Bailey, Borwein, López de Prado and Zhu on backtest overfitting has shown — repeatedly and rigorously — that the vast majority of "discoveries" in quantitative finance are false. When a research process tries many strategy variants and reports only the best one, the reported Sharpe ratio dramatically overstates the true expected Sharpe ratio. The relevant correction is the Deflated Sharpe Ratio, which adjusts for the number of trials, the variance of trial outcomes, and the higher moments of the return distribution.

Source: *Bailey, D. H. & López de Prado, M. (2014): "The Deflated Sharpe Ratio: Correcting for Selection Bias, Backtest Overfitting and Non-Normality." Journal of Portfolio Management, 40(5), 94–107.*

> ***Source:*** *López de Prado, M. (2018): "A Data Science Solution to the Multiple-Testing Crisis in Financial Research." SSRN 3177057.*

This is the reason LightFlux-FX is presented primarily through its live track rather than through an extensive backtest. The 25-month period from April 2024 to April 2026 represents continuous execution on a single account with no backfill, no survivorship bias, and no mid-stream parameter restatement. All trading results are extracted strictly from broker statements (FXBlue) and reconciled daily — meaning the reported performance is what an investor would actually have experienced, net of all real-world frictions. The 3,987 trades executed over this period provide a statistically meaningful sample, though still well below the threshold at which return-distribution moments stabilise.

A live track is not immune to selection bias — a strategy that had performed poorly would presumably not be marketed in the first place. The honest reader should treat the 25-month numbers as preliminary evidence consistent with the architectural claims, not as definitive proof of persistent alpha. The architecture itself — diversification across 196 nodes, hard loss envelopes, adaptive node reweighting — should give an investor more confidence than the headline numbers alone.

2. Architecture of the LightFlux-FX System

The strategy is built on the LIGHTBRIDGE decision framework, developed by the team at Motherhub S.L. and its affiliated research lab Lightstorm AI (Madrid). According to publicly available information, the underlying technology represents over €3.3 million in cumulative R&D investment and more than five years of development, and is currently deployed by regulated institutions managing in excess of €1 billion in assets.

> ***Source:*** *Lightstorm AI public profile (2026), www.lightstorm.ai; ZoomInfo company directory.*

The architecture has four layers, each addressing a distinct problem in systematic FX execution: signal generation (the 196 decision nodes), regime detection (the meta-layer that selects which nodes to activate), position sizing (the monetary-loss-envelope framework), and orchestration (the rolling re-weighting process). The remainder of this section addresses each layer.

2.1 The 196-Node Ensemble

At the lowest level, the system maintains 196 independent decision nodes. Each node is a specialised signal-generation unit trained on a specific configuration: a particular currency pair from the 28 G10 pairs traded; a particular market regime characterised by realised volatility, interest-rate term-structure features, and macro factors; and a particular predictive horizon, ranging from intraday signals (held a few hours) to multi-day signals (held up to two weeks at the 95th percentile).

The decision nodes themselves are not described in detail in publicly available material — for obvious commercial reasons — but the overall architecture is consistent with the recent academic literature on hybrid machine-learning ensembles in quantitative finance. Ranjan (2025) demonstrates that combining neural networks with tree-based voting models in a cross-asset framework can deliver Sharpe ratios above 2.5 and CAPM alphas above 28 % per annum, while maintaining market beta below

0.55. Borrageiro, Firoozye and Barucca (UCL, 2022) show that recurrent reinforcement learning applied to major cash FX pairs, properly accounting for transaction and funding costs, can produce significantly positive risk-adjusted returns through quadratic utility maximisation.

> ***Source:*** *Ranjan, A. (2025): "Causal and Predictive Modeling of Short-Horizon Market Risk and Systematic Alpha Generation Using Hybrid Machine Learning Ensembles." arXiv:2510.22348.*

> ***Source:*** *Borrageiro, G., Firoozye, N. & Barucca, P. (2022): "Reinforcement Learning for Systematic FX Trading." arXiv:2110.04745.*

Two design principles distinguish the LightFlux-FX node architecture from the academic state-of-the-art. First, the system is deliberately non-predictive in the speculative sense: nodes do not attempt to forecast price direction over arbitrary horizons. Instead, each node identifies narrow statistical windows — typically a few hours wide — in which its edge has historically held, and declines to trade outside those windows. The publicly disclosed activity profile shows a consistent entry window between 04:00 and 17:00 UTC, corresponding to the overlap of the major Asian, European and US trading sessions where market depth is highest and slippage costs are lowest. Second, the meta-layer continuously re-evaluates node performance and silently down-weights underperformers while reinforcing nodes whose statistical edge persists.

2.2 The 28-Pair Universe: Why G10, Why Not Exotics

The strategy trades a curated portfolio of 28 G10 currency pairs across USD, EUR, GBP, JPY, CHF, CAD, AUD and NZD, including majors (EUR/USD, GBP/USD, USD/JPY, AUD/USD, USD/CAD, USD/CHF, NZD/USD) and selected crosses (EUR/GBP, EUR/JPY, GBP/JPY). The instrument profile is spot FX and FX rolling-spot contracts; the current live track is executed through FX CFDs at Swissquote Capital Markets Ltd., Cyprus, with a planned transition to listed FX futures in the second quarter of 2026.

The exclusion of emerging-market and exotic currencies is a deliberate design choice grounded in both academic evidence and operational reality. Cao, Jang, Tan and Yang (2021) and earlier work by Menkhoff et al. (2012) document that minor and exotic currencies display stronger momentum and reversal anomalies than G10 majors — meaning that the inefficiencies a systematic strategy would target are larger in those pairs. The same papers also document that these higher edges come at the cost of dramatically wider bid-ask spreads, higher overnight funding costs, and execution risks that erode the apparent edge once trading frictions are properly accounted for. The work of Menkhoff et al. (2012) specifically finds that minor currencies' contribution to momentum portfolio returns is largely consumed by transaction costs.

> ***Source:*** *Menkhoff, L., Sarno, L., Schmeling, M. & Schrimpf, A. (2012): "Currency Momentum Strategies." Journal of Financial Economics, 106(3), 660–684.*

For a systematic strategy executing several thousand trades per year, the choice of universe is therefore a trade-off between gross edge and net edge after frictions. G10 pairs have lower gross edges but dramatically lower frictions; emerging-market pairs have higher gross edges but frictions large enough to consume much of the gross. The empirical evidence from large-scale studies suggests that net edge is roughly comparable across both universes — but the G10 universe wins on two further dimensions: liquidity (allowing the strategy to scale to institutional capacity without slippage degradation) and statistical reliability (deeper price histories with cleaner microstructure).

2.3 Risk Deconcentration: The Architecture of Bounded Bets

The third architectural pillar is the explicit decision to source return from many small, independent positions rather than from a few large concentrated bets. The publicly disclosed risk parameters define this rigorously: maximum risk per trade is capped at 2 % of account equity, maximum single-instrument concentration at 25 % of total exposure over the life of the deposit, maximum single-currency concentration at 50 %, and maximum simultaneous open positions at 150.

In practice, the system operates well within these envelope limits. Each position is opened under a pre-budgeted loss envelope typically representing 0.14–0.21 % of equity — meaning a single trade going to its worst-case stop costs the portfolio less than one-fifth of one percent. Even under a theoretical scenario in which every open position simultaneously hits its maximum programmed loss, the aggregate impact remains structurally constrained: in a representative live snapshot, the worst-case simultaneous loss was approximately 3.5 % of equity, against gross notional leverage of approximately 14.8×. The gap between gross leverage and real downside is the result of explicit risk budgeting at the trade level — not leverage in the traditional notional sense.

This is the architectural innovation that most clearly distinguishes LightFlux-FX from traditional leveraged FX strategies. The latter typically define risk through notional exposure ("we trade at 5× leverage") and discover their actual downside only when adverse moves materialise. LightFlux-FX defines risk through pre-budgeted monetary loss envelopes at the order level — the worst-case loss per trade is known and capped ex ante. The role of leverage in this architecture is to allow a large number of small, bounded positions to fit within a given capital base, not to amplify the impact of any individual decision.

3. The Operational Framework: Adaptive Orchestration

If the architecture defines the system's static structure, the orchestration layer defines its dynamics. The orchestration layer answers three questions on a continuous basis: which of the 196 nodes should be active right now; what position size should each active node take; and how should the meta-weights of the nodes evolve in response to ongoing performance.

3.1 Node Activation and the Regime-Aware Layer

Each node operates under a statistical filter: it acts only when its measured prediction efficiency in the current market regime exceeds a pre-defined threshold. "Prediction efficiency" here is operationalised as the rolling Sharpe ratio of the node's recent (paper or live) signals — meaning a node that has historically performed well in low-volatility EUR/USD regimes will be activated when current conditions resemble that regime, and deactivated when conditions diverge.

The regime identification process is itself adaptive: the system continuously explores new combinations of instruments, predictive horizons, and feature sets, identifies configurations that show statistically reliable edges in the current environment, and integrates them into the active node pool. This is the orchestration cycle. Each cycle reflects an updated regime selection, reinforcing diversification while preserving a consistent underlying search logic.

This design draws directly on the academic literature on regime-switching models in finance, but with an important difference. Classical regime-switching models (Hamilton 1989, and subsequent extensions) assume a small number of discrete regimes — typically two or three — and estimate transition probabilities between them. The LightFlux-FX approach is less parametric: it treats regimes as emergent properties of the data, defined by which nodes are currently performing well, rather than as latent states to be inferred. This has the practical advantage of avoiding the regime-detection lag that plagues parametric models, and the practical disadvantage of being less interpretable.

3.2 Position Sizing: From Signal to Order

Once a node is active, position sizing is calibrated to two independent constraints. First, the prevailing market regime: in higher-volatility regimes, position sizes are systematically reduced. Second, the node's measured prediction efficiency: nodes with stronger empirical edges receive larger size, weaker nodes receive smaller size. Critically, position sizing is not calibrated to a single volatility input (as in classical risk-parity frameworks) but to the joint distribution of regime conditions and node-level performance.

Every position is opened with a mandatory stop-loss enforced at the order ticket level. This is the hard envelope. The stop is not subject to discretionary widening — there is no operator who can decide, in the heat of the moment, to give the trade "a bit more room." The fundamental discipline of the architecture is that exposure is bounded before the trade is entered, not after the loss has begun to materialise.

This matters because the academic and industry literature on stop-loss orders is sharply divided. Stop-losses are sometimes criticised on the grounds that they convert paper losses into realised losses and prevent mean-reversion from working in the holder's favour. The criticism has merit in trend-following equity strategies where the underlying mean is upward-drifting and patience tends to be rewarded. It has substantially less merit in short-horizon FX strategies where the underlying drift is approximately zero and a position that moves against the entry signal carries no expected reward for continued holding. In the LightFlux-FX context, the stop-loss is not a discretionary risk-management overlay; it is the architectural primitive that makes the entire ensemble framework work.

3.3 The Asymmetry of Activation and Deactivation

A subtle but important property of the orchestration layer is its asymmetry. The threshold for deactivating an underperforming node is set lower than the threshold for activating a new one. The economic logic is straightforward: a node that has historically performed well but is currently struggling may be in a temporary drawdown caused by an idiosyncratic regime shift, and re-instating it later carries low cost. A node that is brought into active service prematurely, however, carries an immediate exposure to a signal that has not yet demonstrated its edge — and the cost of being wrong is direct portfolio loss. The system therefore moves quickly to silence underperformers and slowly to elevate newcomers.

This asymmetry mirrors the well-documented behavioural-finance literature on loss aversion (Kahneman & Tversky 1979) but applies it at the level of the trading algorithm rather than the human decision-maker. The system is, in effect, deliberately loss-averse — which in a high-uncertainty, short-horizon trading context is the architecturally correct stance.

4. Risk Architecture: Pre-Budgeted Loss vs. Notional Leverage

The risk architecture deserves its own section because it is the feature most often misunderstood by readers who approach the strategy through the lens of traditional leveraged FX. The publicly disclosed parameters look familiar enough — maximum drawdown -40 %, maximum monthly loss -20 %, maximum daily loss -10 %, maximum leverage 20× gross notional. To an experienced retail FX trader these numbers may even appear aggressive. The architecture, however, behaves very differently from a conventional 20×-leveraged book.

4.1 The Monetary Loss Envelope

Conventional leveraged FX strategies define risk at the position level through notional exposure: a 5-lot EUR/USD position represents a 500,000-EUR notional bet, and the unrealised loss is whatever the market does to that notional. The trader's actual downside depends on how the market moves between entry and the point at which they decide to close — and "decide" is doing significant work in that sentence, since human risk tolerance is highly path-dependent.

LightFlux-FX defines risk monetarily at the order level. Before the trade is entered, the system specifies the maximum monetary loss the position can incur, typically 0.14–0.21 % of account equity. This loss envelope is hard-coded into the stop-loss order ticket and cannot be widened. The notional exposure of the trade is whatever it needs to be to make the monetary envelope economically meaningful — but the monetary loss is the binding constraint, not the notional size.

The practical consequence is that gross leverage and real portfolio risk become decoupled. A representative live snapshot reported aggregate notional leverage of approximately 14.8× — a figure that, in a traditional framework, would imply a fragile, highly-leveraged book. The actual aggregate worst-case loss across the same snapshot was approximately 3.5 % of portfolio equity. The gap is the architectural contribution: the strategy uses notional leverage as a mechanical lever for fitting many bounded bets onto a given capital base, while real downside is governed by the sum of monetary envelopes.

4.2 Hard Envelopes at Three Levels

The risk architecture enforces three nested envelopes, monitored in real time:

First, the per-trade envelope, hard-coded at the order ticket. No single trade can lose more than its pre-budgeted monetary envelope. This is the smallest and most frequently binding constraint.

Second, the per-decision-unit daily envelope. Each of the 196 nodes has its own daily loss budget — typically a multiple of its average single-trade envelope. When a node reaches its daily cap, that specific node is deactivated until the next session, while the rest of the system continues to operate.

Third, the aggregate portfolio cap. If the sum of intraday losses across all active nodes reaches a defined threshold, all execution halts until the next session. This is the system's hard stop — the architectural backstop that ensures even a coordinated failure across multiple nodes cannot produce a catastrophic single-day loss.

The three-level envelope structure addresses a specific failure mode that has historically plagued systematic FX strategies: correlated losses across positions during volatility events. A single node's stop-loss is no protection against a regime in which all of the strategy's exposures move adversely simultaneously. The portfolio-level cap is the structural protection against that scenario.

4.3 Why This Differs from Traditional FX Risk Frameworks

Most retail-grade FX systems fail not because their signal is bad, but because their risk architecture is misspecified. They confuse leverage with edge: they take a moderately profitable signal and amplify it through notional exposure, on the implicit assumption that the historical drawdown profile of the signal will continue to hold. When the assumption breaks — typically during a volatility regime not represented in the training sample — the amplified downside dwarfs anything the signal can recover from. The 2015 Swiss franc unpeg, the August 2024 yen carry unwind, and the April 2025 tariff shock are recent examples of such regime breaks in FX.

The LightFlux-FX architecture is engineered to avoid this failure mode. The per-trade monetary envelope means that a regime break cannot translate into catastrophic single-trade losses. The per-node daily cap means that a regime break affecting one decision unit cannot spill across to the others. The aggregate portfolio cap means that even a coordinated regime break across multiple nodes is bounded. The engineering effort, as the strategy's own marketing materials put it, is concentrated not in amplifying edge through leverage but in refusing to trade when conditions are unfavourable.

5. Regime Analysis: Three Live-Track Stress Windows

Every framework must be tested against reality. The following analysis examines three fundamentally different FX-relevant stress events that occurred within the LightFlux-FX live track, drawn directly from the strategy's published stress-window analysis. In each case we examine how the strategy behaved and compare against standard equity benchmarks over the same window.

5.1 Yen Carry-Trade Unwind, August 2024

The August 2024 yen carry-trade unwind was one of the most violent FX dislocations in recent memory. On 31 July 2024, the Bank of Japan unexpectedly raised its policy rate to 0.25 % and announced a tapering of its quantitative-easing programme. The yen surged against the dollar, equity markets globally repriced, and the Nikkei 225 dropped more than 12 % in a single session on 5 August. The VIX spiked to a level not seen since the pandemic of 2020. The Bank for International Settlements subsequently documented the event in detail in its August 2024 Bulletin, concluding that the unwinding of crowded carry positions had amplified the initial monetary-policy shock far beyond what fundamentals alone would have justified.

> ***Source:*** *Aquilina, M., Lombardi, M., Schrimpf, A. & Sushko, V. (2024): "The market turbulence and carry trade unwind of August 2024." BIS Bulletin No. 90, Bank for International Settlements, 27 August 2024.*

Over the strategy's published stress window of 1–5 August 2024, LightFlux-FX returned +0.21 %. The S&P 500 returned -6.08 %, the NASDAQ -7.95 %, and the MSCI World -6.19 %. The strategy's return signature in this period was modestly positive in absolute terms and dramatically positive in relative terms — a result consistent with the architecture's design: the system trades on short-horizon statistical opportunities rather than on the directional carry exposure that the broader market was being forced to unwind. A strategy structurally short of yen-funded crowded positions cannot, by construction, be wrong-footed by a crowded-position unwind.

5.2 Trump Tariff Shock, April 2025

The April 2025 "Liberation Day" tariff announcement triggered a sharp repricing across cross-asset markets. Equity indices sold off aggressively (S&P 500 -3.13 % over 2–9 April, NASDAQ -1.86 %, MSCI World -3.33 %), and FX markets experienced significant directional moves as currency baskets adjusted to changed trade-flow expectations.

LightFlux-FX returned -0.82 % over the same window. This is a modestly negative outcome — the strategy was not immune to the dislocation — but the relative resilience is meaningful: a -0.82 % loss against -3.13 % for the S&P 500 represents an outperformance of 2.3 percentage points over a one-week window. The architecture's behaviour in this event is illustrative: the per-trade envelopes prevented any single FX exposure from delivering a catastrophic loss, while the per-node daily caps deactivated nodes whose signals were temporarily misaligned with the post-shock regime. The strategy bent without breaking.

5.3 DeepSeek AI Shock, January 2026

The DeepSeek AI shock of January 2026 was a different kind of event: a tech-sector-driven repricing triggered by the release of a Chinese open-source large language model that called into question the capital-expenditure trajectory of US AI infrastructure providers. The shock was equity-led and largely sector-specific, with the broader macro environment relatively undisturbed. Over the 20–22 January 2026 window, the S&P 500 returned -0.38 %, the NASDAQ -0.34 %, and the MSCI World -0.34 %.

LightFlux-FX returned -1.46 % — a worse outcome than the equity indices. The asymmetry is informative: in an FX-led dislocation (August 2024), the strategy outperformed dramatically. In an equity-led dislocation with limited FX content (January 2026), the strategy underperformed modestly, presumably because the second-order FX moves triggered by the equity repricing happened to align poorly with the active node configuration. This is the honest empirical observation: low-beta does not mean negative-beta, and the strategy is not designed to be an equity hedge. It is designed to be uncorrelated, and uncorrelated returns include underperformance in benign equity environments just as they include outperformance in stressed ones.

5.4 Stress-Window Summary

Event	Window	LightFlux-FX	S&P 500	NASDAQ	MSCI World
Yen Carry Trade Unwind	1–5 Aug 2024	+0.21 %	-6.08 %	-7.95 %	-6.19 %
US Election Volatility	4–8 Nov 2024	-1.17 %	+4.66 %	+5.74 %	+3.58 %
Year-End Risk-Off	6–13 Jan 2025	-0.87 %	-1.79 %	-2.72 %	-1.75 %
Global Risk-Off Feb	21–28 Feb 2025	+0.60 %	-2.66 %	-5.59 %	-2.27 %
Trump Tariff Shock	2–9 Apr 2025	-0.82 %	-3.13 %	-1.86 %	-3.33 %
Summer Volatility	7–14 Jul 2025	-1.80 %	-0.17 %	+0.19 %	-0.24 %
Q4 Correction	10–17 Nov 2025	-1.63 %	-0.84 %	-1.29 %	-0.70 %
December Fed Shock	16–19 Dec 2025	+0.49 %	+0.26 %	+1.09 %	+0.22 %
DeepSeek AI Shock	20–22 Jan 2026	-1.46 %	-0.38 %	-0.34 %	-0.34 %
Tech CapEx Rout	4–6 Feb 2026	-0.27 %	+0.21 %	-0.96 %	+0.42 %

The pattern across the full table is consistent with a strategy whose returns are statistically independent of equity market direction rather than systematically hedging it. In the eight FX-relevant stress events shown, LightFlux-FX outperformed the S&P 500 in four (Yen Unwind, Global Risk-Off Feb, Trump Tariff Shock, December Fed Shock) and underperformed in four (US Election, Year-End, Summer Volatility, Q4 Correction, DeepSeek, Tech CapEx). The realised beta of 0.096 over the full live track is therefore not the result of a single fortunate event but the cumulative outcome of broadly uncorrelated daily returns.

6. Risk and Exposure Profile

The strategy's published risk-return statistics over the live track are summarised below. These figures represent net-of-fee performance reconciled daily against broker statements, with Layer 2 Net returns (trading + interest rebate + broker rebate) reported throughout.

Metric	Value	Comment
Live track length	25 months (Apr 2024 – Apr 2026)	Single continuously traded account
Total return (net of fees)	+62.46 %	Compounded daily, time-weighted
Annualised return p.a.	approximately 25 %	Geometric mean over live track
Annualised volatility	approximately 12 %	Daily-return based
Maximum drawdown	-8.52 %	Peak-to-trough, daily granularity
Worst monthly return	-3.62 %	August 2024
Best monthly return	+11.89 %	April 2024 (inception month)
Modified Sharpe ratio	1.85	Rf = €STR
Sortino ratio	5.08	Downside σ only
Calmar ratio	3.51	Ann. return / max drawdown
Hit rate	55.7 %	Winning trades / total
Win / loss months	16 / 9	64 % positive months
Number of trades	3,987	Full live track
Average daily margin use	approximately 13.9 %	Of available account margin
Beta vs. S&P 500	0.096 (R^2 0.8 %)	Pure-alpha threshold (GIPS): $\beta < 0.15$

Three of these numbers deserve particular emphasis. The Sortino ratio of 5.08 reflects the strategy's pronounced asymmetry between upside and downside: the volatility of negative returns is meaningfully lower than the volatility of positive returns, which is the empirical signature of a hard-stop architecture (downside is bounded; upside is not). The Calmar ratio of 3.51 is high enough to be eye-catching but should be interpreted with caution — the realised maximum drawdown over a 25-month live track is almost certainly an under-estimate of the strategy's true maximum drawdown over a longer horizon. The β of 0.096 versus the S&P 500 is, by the convention used in performance-attribution literature, below the threshold ($\beta < 0.15$) at which a strategy is considered structurally uncorrelated with the equity reference benchmark.

7. Performance Characteristics and Alpha Sources

The strategy's value-added relative to a passive carry-trend-value FX exposure derives from four identifiable sources, which materialise to different degrees in different market environments.

7.1 Source 1 — Ensemble Aggregation Effect

The most important alpha source is the statistical aggregation effect across the 196 decision nodes. Each individual node has, on average, a small statistical edge — a few basis points per trade. Individually, these edges are too small to be commercially meaningful and too noisy to be statistically reliable. Aggregated across 196 nodes operating in parallel across 28 pairs, the law of large numbers takes hold: the aggregate signal becomes statistically meaningful even as each individual contribution remains small. This is the same mechanical principle that underpins ensemble methods in machine learning more broadly (random forests, gradient boosting), applied to financial signal generation.

7.2 Source 2 — Time-of-Day Microstructure

The strategy concentrates execution between 04:00 and 17:00 UTC — the window in which the Asian, European and US trading sessions overlap and bid-ask spreads in major G10 pairs reach their narrowest. The microstructure of FX is meaningfully different inside this window than outside it: liquidity is deeper, slippage is lower, and price formation is dominated by genuine information flow rather than thin overnight quoting. Trading inside this window is not itself an alpha source — but it is a precondition for harvesting the small per-trade edges identified by the node ensemble without having them consumed by transaction costs.

7.3 Source 3 — Adaptive Regime Selection

The third source is the meta-layer's ability to identify which nodes are currently aligned with the prevailing market regime and to silently down-weight those that are not. In a stationary world this layer would add nothing — the same nodes would always be optimal — but FX markets are emphatically not stationary, and the configurations that work in low-volatility regimes differ systematically from those that work in stressed regimes. The adaptive layer is, in effect, a soft regime-switching framework that allocates capital across hundreds of micro-strategies in proportion to their current empirical edge.

7.4 Source 4 — Structural Low-Beta Design

The fourth source is harder to attribute as "alpha" in the strict sense, but it is the source that matters most for portfolio-construction purposes: the strategy's near-zero correlation with traditional risk assets. The literature on alternative investments is full of strategies that claim diversification benefits and deliver, when stress arrives, a correlation of 0.7 to equities (Klement 2024, Sullivan & Wey 2024). The realised β of 0.096 for LightFlux-FX over a 25-month period that included multiple equity-market dislocations suggests the diversification benefit may be more durable here than in many peers. Whether this is structural (driven by the architecture) or partly fortuitous (driven by the specific events that occurred) will become clearer only over a longer track record.

These sources are cyclical rather than constant. In stable, low-volatility FX regimes with predictable trading rhythms, all four sources contribute. In dislocated regimes with sudden volatility breaks, sources 2 and 3 partially break down — and the strategy's relative outperformance compresses or even inverts. The honest expectation is that the strategy produces positive net returns most months and an asymmetric upside in months containing FX-led volatility events, but that it cannot be relied upon to outperform every benign equity month.

8. Model Limitations: What the Strategy Cannot Do

No framework is free of weaknesses, and transparency about limitations is the price of intellectual seriousness. LightFlux-FX has five identifiable structural limitations that any prospective investor should understand before allocating capital.

8.1 Capacity Constraints

The strategy generates a substantial fraction of its return from short-horizon statistical opportunities that exist precisely because they have not been arbitraged out by larger participants. As the strategy's AUM grows, the marginal trade begins to have measurable market impact, and the small per-trade edges that the ensemble exploits begin to compress. The publicly stated minimum investment is €100,000, and the strategy is offered to Professional Clients only — but capacity at the strategy level is not publicly disclosed. Investors should assume that the favourable risk-return profile observed in the current live track is partly a function of the strategy still operating well below its capacity ceiling, and may not survive a multiplicative scaling of AUM.

8.2 Regime-Shift Latency

The adaptive layer that re-weights nodes operates on a rolling-performance basis, which introduces an inherent latency between a regime shift and the system's reaction to it. In gradual regime transitions this is not problematic — the system adjusts over the course of a few sessions and the cost is limited. In abrupt regime shifts, however, the system can carry exposures from the prior regime into the new one for several sessions before the meta-layer responds. The April 2025 tariff shock and the January 2026 DeepSeek shock both fall into this category, and the strategy's modest negative returns over those windows are partly explained by this latency.

8.3 Whipsaw Risk in Trendless, High-Volatility Markets

The strategy's per-trade stop-loss architecture is well-suited to environments with directional momentum, even at short horizons. In markets that are simultaneously high-volatility and trendless — where prices move violently in both directions without settling into a regime — the stop-loss frequency increases, transaction costs compound, and the small per-trade edges are consumed by friction. This is the FX equivalent of the "whipsaw tax" that trend-following strategies pay in choppy equity markets. The strategy cannot, by construction, eliminate this cost; it can only minimise it through the time-of-day filter and the node deactivation logic.

8.4 CFD Instrument Risk and Counterparty Exposure

The current live track is executed through FX CFDs rather than listed FX futures, with execution through Swissquote Capital Markets Ltd., Cyprus. CFD-based execution introduces three categories of risk that listed-futures execution would mitigate: counterparty exposure to the executing broker (in the unlikely event of broker insolvency), pricing transparency limitations relative to centrally-cleared instruments, and regulatory exposure to changes in CFD-marketing rules that could constrain the operational model. The planned transition to listed FX futures in Q2 2026 will address all three concerns but has not yet been implemented as of the writing of this paper. Investors should treat the current operational model as transitional rather than permanent.

8.5 Live Track Length

The most uncomfortable limitation is the simplest one: a 25-month live track is not long enough to make confident claims about long-term expected performance. Drawdown characteristics, in particular, are notoriously path-dependent — the realised -8.52 % maximum drawdown almost certainly under-estimates the true distributional maximum over a multi-year horizon. The 16 winning months and 9 losing months in the current sample translate to a 64 % monthly hit rate, which is encouraging but well within the range that could persist for several years before reverting to a lower long-run mean. The architecture provides reasons for optimism (the bounded-loss design genuinely limits left-tail risk), but only out-of-sample evidence over multiple years can confirm whether the live-track performance is the early sample of a persistent process or a favourable early draw from a less generous long-run distribution.

9. Conclusion

LightFlux-FX presents a thoughtful, internally consistent approach to systematic FX execution. Its core architectural insight is the decoupling of notional leverage from monetary risk: bounded loss envelopes at the order ticket level mean that the strategy can fit many small, independent bets onto a given capital base without translating gross leverage into proportionate downside. Its operational insight is the use of an adaptive 196-node ensemble that aggregates many small, diverse statistical edges into a meaningful aggregate signal, while continuously re-weighting nodes in response to changing market regimes. Its structural property — a realised β of 0.096 versus the S&P 500 over a 25-month live track — places it on the right side of the Pojarliev-Levich alpha-versus-beta dichotomy, at least to date.

The strategy is not, however, a substitute for thinking carefully about portfolio construction. Its short live track, its capacity uncertainty, and its modest underperformance in equity-led dislocations all argue for treating it as a diversifying sleeve rather than as a core allocation. Used in the role for which it is engineered — as an alpha sleeve sized at perhaps 5 to 10 % of a sophisticated portfolio, alongside long-only equity, fixed income, and trend-following allocations — it can deliver structural diversification benefits that are difficult to source elsewhere in the current alternative-investment landscape. Used as a core or oversized allocation, it is exposed to all the standard risks of any quantitative strategy with a short live track: capacity compression, regime breaks, and the gap between sample evidence and population truth.

Sophisticated investors will weigh these considerations against the specific evidence presented here and reach their own conclusions. This paper has tried to present that evidence as honestly as possible — both the favourable and the unfavourable elements. The objective has not been to advocate for the strategy but to describe it in enough detail that investors can decide for themselves whether its architecture, its track record, and its limitations align with their portfolio objectives.

Bibliography

Aquilina, M., Lombardi, M., Schrimpf, A. & Sushko, V. (2024): "The market turbulence and carry trade unwind of August 2024." BIS Bulletin No. 90, Bank for International Settlements, 27 August 2024.

Bailey, D. H. & López de Prado, M. (2014): "The Deflated Sharpe Ratio: Correcting for Selection Bias, Backtest Overfitting and Non-Normality." Journal of Portfolio Management, 40(5), 94–107.

Bailey, D. H., Borwein, J. M., López de Prado, M. & Zhu, Q. J. (2014): "Pseudo-mathematics and Financial Charlatanism: The Effects of Backtest Overfitting on Out-of-Sample Performance." Notices of the American Mathematical Society, 61(5), 458–471.

Bank for International Settlements (2022): 2022 Triennial Central Bank Survey of Foreign Exchange and OTC Derivatives Markets — Global FX turnover reached USD 7.5 trillion per day in April 2022.

Bank for International Settlements (2025): 2025 Triennial Central Bank Survey — Global FX turnover reached USD 9.6 trillion per day in April 2025.

Borrageiro, G., Firoozye, N. & Barucca, P. (2022): "Reinforcement Learning for Systematic FX Trading." arXiv:2110.04745. Department of Computer Science, University College London.

Condorcet, M. de (1785): Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. (The Jury Theorem.)

Hamilton, J. D. (1989): "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle." Econometrica, 57(2), 357–384.

Kahneman, D. & Tversky, A. (1979): "Prospect Theory: An Analysis of Decision under Risk." Econometrica, 47(2), 263–291.

Lightstorm AI (2026): Company profile and technology description. www.lightstorm.ai.

López de Prado, M. (2018): Advances in Financial Machine Learning. Wiley.

López de Prado, M. (2018): "A Data Science Solution to the Multiple-Testing Crisis in Financial Research." SSRN 3177057.

Menkhoff, L., Sarno, L., Schmeling, M. & Schrimpf, A. (2012): "Currency Momentum Strategies." Journal of Financial Economics, 106(3), 660–684.

Pojarliev, M. & Levich, R. M. (2008): "Do Professional Currency Managers Beat the Benchmark?" Financial Analysts Journal, 64(5), 18–32.

Pojarliev, M. & Levich, R. M. (2012): "A New Look at Currency Investing." CFA Institute Research Foundation Monograph.

Ranjan, A. (2025): "Causal and Predictive Modeling of Short-Horizon Market Risk and Systematic Alpha Generation Using Hybrid Machine Learning Ensembles." arXiv:2510.22348. University of Oxford.

Surowiecki, J. (2004): The Wisdom of Crowds: Why the Many Are Smarter Than the Few. Doubleday.

Van Sterling Capital Limited (2026): LightFlux-FX Institutional Factsheet (April 2026); Schedule A Strategy Specifications; Trading Mentor Questionnaire — internal documentation.

— LightFlux-FX Research —