

Alpha Asymmetry in Foreign Exchange Markets

An Investigation of Exploitability

Murad Farzulla^{1,2,*}

¹Dissensus, London, UK ²King's College London, London, UK

*Correspondence: murad@dissensus.ai ORCID: 0009-0002-7164-8704

July 2026

Abstract

This paper investigates whether distributional asymmetries in foreign exchange alpha signals represent exploitable market inefficiencies. Using EUR/JPY data spanning November 2015–August 2025 (504 weekly observations after rolling window warmup), we find that the asymmetry premise itself largely dissolves under dependence-robust measurement: of five alpha signal types, only the volatility-expansion (coverage) signal exhibits skewness that survives block-bootstrap inference ($\hat{\gamma}_1 = 1.75$, 95% CI [1.17, 2.15]); the signed tail signal skews *negative* rather than positive (-1.47 , CI $[-3.09, 0.54]$), and momentum, mean-reversion, and correlation signals are statistically indistinguishable from symmetry. An earlier version of this paper reported pronounced positive tail-signal skewness; that figure described the unsigned exceedance magnitude, which is right-skewed by construction, and we document the correction. The economic null is correspondingly stark. A skewness-threshold strategy earns 3.60% cumulative gross over the decade (17 trades, Sharpe 0.15), is nearly inert in walk-forward testing (3 trades in eight out-of-sample years), and carries no residual alpha against proxy carry, momentum, and dollar factors (intercept $t = 0.67$ on in-position weeks). Because the strategy trades so rarely, transaction costs are immaterial rather than decisive: net returns remain within 0.4 percentage points of gross even at wide retail spreads, with break-even near 19 pips round-trip. Data-snooping corrections complete the null: White's Reality Check ($p = 0.15$) and Hansen's SPA ($p = 0.28$) find no strategy in a 13-candidate universe that outperforms—the best realized mean return belongs to the random benchmark. The GPD shape parameter is not significantly different from zero ($\xi = -0.25$), and tail clustering is weak ($\theta = 0.83$, 95% CI [0.60, 1.00]). We conclude that alpha signal asymmetry in this setting is neither robustly detectable (with one distributional exception that does not yield a tradable edge) nor exploitable, and we caution that unsigned-magnitude constructions can manufacture the appearance of asymmetry where none exists.

Keywords: null result, alpha asymmetry, foreign exchange, skewness, market efficiency, extreme value theory

JEL Codes: G11, G14, G15, C58

1 Introduction

Financial markets exhibit persistent deviations from the efficient market hypothesis (Fama, 1970), with alpha signals—measures of risk-adjusted excess returns—displaying systematic patterns that sophisticated traders exploit (Jegadeesh and Titman, 1993; Moskowitz et al., 2012). While considerable literature examines alpha generation and decay, less attention has been paid to the *distributional properties* of alpha signals themselves. This is surprising given well-documented evidence that asset returns deviate substantially from normality, exhibiting fat tails (Mandelbrot, 1963; Cont, 2001) and asymmetric distributions (Harvey and Siddique, 2000).

The preference for skewed returns has deep roots in asset pricing theory. Kraus and Litzenberger (1976) established that investors prefer positive skewness and dislike negative skewness, implying that

assets with lottery-like payoffs command lower expected returns. Subsequent work has confirmed that skewness affects both individual asset pricing (Harvey and Siddique, 2000; Bali et al., 2011) and portfolio construction (Mitton and Vorkink, 2007; Brunnermeier et al., 2007).

This paper investigates whether alpha signals exhibit systematic asymmetries in their probability distributions, and whether such asymmetries can inform profitable trading strategies. We report negative findings that begin earlier in the chain than the usual backtest disappointment: under dependence-robust measurement, most of the asymmetry is not there. We investigate the following questions:

1. Do alpha signals in forex markets deviate significantly from normal distributions? (*Mostly—four of five types reject normality; the normalized-return signal is exactly Gaussian*)
2. Do these deviations manifest as exploitable skewness and heavy tails? (*Largely no—one signal is robustly skewed, the tail signal skews negative and fragilely, and there are no heavy tails*)
3. Do distributional asymmetries persist across market regimes? (*No—the tail signal’s skewness flips sign across regimes*)
4. Do asymmetry-aware strategies outperform benchmarks? (*No—returns are statistically indistinguishable from zero before costs and trail buy-and-hold outright, the strategy is dormant under walk-forward selection, and nothing in the candidate universe beats a random benchmark*)

We document a plausible-sounding trading idea that does not survive rigorous testing, and we document *where* it fails: earlier than expected. The asymmetry-exploitation thesis does not fall at the translation from statistical pattern to economic profit; it falls at signal detection, where an unsigned-magnitude construction had manufactured the appearance of pervasive positive skewness that corrected measurement does not support. Null results of this kind are underreported in quantitative finance (Harvey, 2017), yet they prevent wasted research effort and capital allocation to spurious patterns.

The FX factor literature has seen recent advances in factor construction methodology. Fan et al. (2025) introduce a framework that dynamically optimizes currency factor strategies—including carry, momentum, and value—via spot and forward trading, evaluating 24,336 portfolio optimization approaches and demonstrating that optimized factors significantly outperform naïve constructions after correcting for data snooping bias. Our work complements this optimization-focused agenda by examining a more fundamental question: whether the *distributional asymmetry* in factor returns (long vs. short side) itself constitutes an exploitable signal. Where Fan et al. optimize how factors are constructed, we test whether the asymmetric properties of the resulting return distributions carry economic information. Hertrich (2025) provides additional context through analysis of G10 carry, momentum, and value strategies under forward-looking Conditional Value-at-Risk conditioning, finding that the carry trade risk premium has remained unexpectedly low since the global financial crisis despite negative covariance with global FX volatility—suggesting a potential failure of currency pricing theory that our distributional analysis may help illuminate.

Our empirical contribution is threefold. First, we provide a systematic characterization of distributional properties across five distinct alpha types in EUR/JPY forex data, contributing to the growing literature on FX market microstructure (King et al., 2013; Evans and Lyons, 2002). Second, we demonstrate that asymmetry-based trading strategies fail statistically before frictions even enter—gross returns are indistinguishable from zero, walk-forward selection leaves the strategy dormant, and transaction costs prove immaterial at the realized trade frequency—extending work on technical trading rules (Brock et al., 1992; Lo et al., 2000; Neely et al., 2014) by showing where such rules fail. Third, we

test cross-market transfer and find nothing stable: applied unchanged to GBP/USD, SPY, and GLD, the strategy’s gross outcomes range from +17% to −30% cumulative with no relationship to the underlying asymmetry statistics, and the EUR/JPY tail-skew signature itself reverses sign in all three markets—in contrast to the pervasiveness of classic factor premia across asset classes (Asness et al., 2013).

2 Methodology

2.1 Data and Sample Construction

We analyze EUR/JPY forex data spanning November 2015 to August 2025, sourced from Yahoo Finance (EURJPY=X). The raw sample comprises 513 weekly observations; after discarding the first 9 observations required for 20-week rolling window initialization of skewness and volatility estimates, the analysis-ready sample contains 504 weekly observations. We note that Yahoo Finance quotes are indicative mid-rates rather than executable prices; accordingly, all transaction cost analysis (Section 5.4) uses separate spread assumptions calibrated to institutional and retail FX execution venues. The sample period encompasses multiple market regimes including the post-Brexit volatility spike (2016), COVID-19 pandemic shock (2020), and subsequent monetary policy divergence between the Federal Reserve, ECB, and Bank of Japan.

Data Frequency and Aggregation. Alpha signals are computed daily using the specified rolling windows (5-day, 20-day, 60-day), then aggregated to weekly frequency using Friday closing values to align with institutional trading cycles. The 504 analysis-ready weekly observations represent end-of-week snapshots of daily-computed signals, ensuring consistency between high-frequency alpha construction and lower-frequency backtesting. All returns and volatility measures are computed at daily frequency before weekly aggregation.

The dataset includes five pre-computed alpha signals, each capturing distinct aspects of market dynamics:

- **Tail Alpha (α_{tail}):** Captures extreme price movements using a modified z -score of returns beyond the 95th percentile. Formally, $\alpha_{\text{tail},t} = \mathbf{1}_{|r_t| > q_{0.95}} \cdot \text{sgn}(r_t) \cdot |r_t|$, where $q_{0.95}$ denotes the rolling 52-week 95th percentile of absolute returns.
- **Fast Alpha (α_{fast}):** Short-term momentum signal computed as the 5-day return normalized by 20-day realized volatility: $\alpha_{\text{fast},t} = (P_t - P_{t-5}) / (\sigma_{20,t} \sqrt{5})$.
- **Pricing Alpha (α_{price}):** Mean-reversion signal measuring deviation from fair value, computed as $\alpha_{\text{price},t} = (P_t - \text{MA}_{60,t}) / \sigma_{60,t}$, where MA_{60} and σ_{60} denote 60-day moving average and standard deviation.
- **Coverage Alpha (α_{cov}):** Volatility compression ratio measuring regime shifts in realized volatility:

$$\alpha_{\text{cov},t} = \frac{\sigma_{20,t}}{\sigma_{20,t-5}} - 1 \quad (1)$$

where $\sigma_{20,t}$ denotes 20-day realized volatility computed from daily close-to-close returns. Data source: Yahoo Finance (EURJPY=X). Values above zero indicate volatility expansion; values below zero indicate compression.

- **Hedge Alpha (α_{hedge}):** Captures the interaction between dollar correlation and interest rate carry. When EUR/JPY moves in tandem with dollar strength (high positive correlation with DXY), carry trades face amplified risk during dollar rallies. The hedge alpha is positive when: (i) correlation

is positive and Japan offers higher rates (favorable for short EUR/JPY hedges); or (ii) correlation is negative and US offers higher rates (favorable for long EUR/JPY hedges). Formally:

$$\alpha_{\text{hedge},t} = \rho_t^{(\text{EUR/JPY,DXY})} \times \Delta r_t^{(\text{JPY-USD})} \quad (2)$$

where ρ is a 100-trading-day (≈ 20 -week) rolling correlation of daily EUR/JPY and DXY returns (US Dollar Index, DX-Y.NYB from Yahoo Finance), and $\Delta r = r_{\text{JPY}} - r_{\text{USD}}$ is the annualized JPY–USD interest rate differential, approximated in the replication pipeline by its sample-average value of -2% (a constant proxy, so the correlation term carries all time variation in the signal). This alpha captures regimes where directional (correlation) and carry (rate differential) signals align, suggesting exploitable hedging opportunities.

For cross-market validation, we apply the same transformations to GBP/USD, SPY (S&P 500 ETF), and GLD (Gold ETF) daily closes from Yahoo Finance, constructing tail, fast, pricing, and coverage alphas identically to EUR/JPY. Hedge alpha is excluded from the cross-market set because its ingredients—the JPY–USD rate differential and the DXY correlation—are specific to the EUR/JPY setting.

2.2 Asymmetry Metrics

We employ four complementary measures to characterize distributional asymmetry, each capturing distinct aspects of non-normality relevant to trading strategy design.

Sample Skewness measures the degree of asymmetry around the mean. For a sample $\{x_1, \dots, x_n\}$ with sample mean $\bar{x}$ and standard deviation s , we compute the adjusted Fisher-Pearson coefficient:

$$\hat{\gamma}_1 = \frac{n}{(n-1)(n-2)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s} \right)^3 \quad (3)$$

This bias-corrected estimator is consistent under standard regularity conditions. Positive skewness ($\hat{\gamma}_1 > 0$) indicates a right-tailed distribution with more extreme positive values, while negative skewness indicates left-tail heaviness. The standard error of skewness under normality is approximately $\sqrt{6/n} \approx 0.109$ for our sample size ($n = 504$).

Excess Kurtosis captures tail heaviness relative to a Gaussian benchmark:

$$\hat{\gamma}_2 = \frac{n(n+1)}{(n-1)(n-2)(n-3)} \sum_{i=1}^n \left(\frac{x_i - \bar{x}}{s} \right)^4 - \frac{3(n-1)^2}{(n-2)(n-3)} \quad (4)$$

This estimator subtracts 3 (the kurtosis of a normal distribution) and applies finite-sample bias correction. Values exceeding zero indicate leptokurtic (fat-tailed) distributions; the standard error under normality is approximately $\sqrt{24/n} \approx 0.218$.

Asymmetry Index (AI) quantifies the ratio of upside to downside semi-variance, providing a risk-management-oriented asymmetry measure:

$$AI = \frac{\text{Var}^+(X)}{\text{Var}^-(X)} = \frac{\sum_{i: x_i > \bar{x}} (x_i - \bar{x})^2 / n^+}{\sum_{i: x_i < \bar{x}} (x_i - \bar{x})^2 / n^-} \quad (5)$$

where n^+ and n^- denote observations above and below the mean. Values above 1 indicate greater dispersion in positive deviations—relevant for assessing “lottery ticket” payoff structures (Bali et al., 2011).

Positive Observation Ratio (POR) measures the unconditional probability of positive realizations:

$$POR = \frac{1}{n} \sum_{i=1}^n \mathbf{1}_{x_i > 0} \quad (6)$$

Under normality with zero mean, $POR = 0.5$; deviations indicate location-scale asymmetry or non-zero drift.

2.3 Statistical Tests

We employ a battery of complementary tests to assess departures from normality, each with distinct power properties against different alternatives.

Shapiro-Wilk Test. The Shapiro-Wilk statistic (Shapiro and Wilk, 1965) tests the null hypothesis $H_0 : X \sim \mathcal{N}(\mu, \sigma^2)$ against general alternatives. The test statistic is:

$$W = \frac{(\sum_{i=1}^n a_i x_{(i)})^2}{\sum_{i=1}^n (x_i - \bar{x})^2} \quad (7)$$

where $x_{(i)}$ are order statistics and a_i are tabulated coefficients. This test has high power against asymmetric and heavy-tailed alternatives for moderate sample sizes.

D'Agostino-Pearson Omnibus Test. Following D'Agostino et al. (1990), we test skewness and kurtosis jointly using the K^2 statistic:

$$K^2 = Z_1(\hat{\gamma}_1)^2 + Z_2(\hat{\gamma}_2)^2 \sim \chi_2^2 \quad (8)$$

where Z_1 and Z_2 are normalizing transformations of sample skewness and kurtosis. This test is particularly powerful against asymmetric alternatives.

Jarque-Bera Test. As a robustness check, we employ the Jarque-Bera statistic (Jarque and Bera, 1980):

$$JB = \frac{n}{6} \left(\hat{\gamma}_1^2 + \frac{(\hat{\gamma}_2)^2}{4} \right) \stackrel{a}{\sim} \chi_2^2 \quad (9)$$

This Lagrange multiplier test is asymptotically equivalent to the likelihood ratio test for normality.

Skewness Significance. We test $H_0 : \gamma_1 = 0$ using the t -ratio $t = \hat{\gamma}_1 / \text{SE}(\hat{\gamma}_1)$, where $\text{SE}(\hat{\gamma}_1) \approx \sqrt{6(n-2)/[(n+1)(n+3)]}$. For $n = 504$, this yields a standard error of approximately 0.109.

Multiple Testing Correction. Given five alpha types and multiple test statistics, we apply the Bonferroni correction to control the family-wise error rate at $\alpha = 0.05$, yielding adjusted significance thresholds of $\alpha^* = 0.05/5 = 0.01$ per alpha type.

2.4 Backtesting Framework

We evaluate the asymmetry strategy against three benchmarks, with explicit signal generation and execution rules to ensure reproducibility.

Asymmetry Strategy:

Signal Generation:

- **Entry Long:** Rolling skewness of fast alpha (20-week window) exceeds 0.75 AND current fast alpha > 0
- **Entry Short:** Rolling skewness of pricing alpha (20-week window) exceeds 0.75 AND current pricing alpha $> 0.5\sigma_{\text{pricing}}$

- **Exit:** Signal reversal (opposite entry condition met) OR 4-week maximum holding period reached

Position Sizing:

$$\text{Position} = \max(0.5, \min(2.0, 1 + |AI_t - 1.0|)) \quad (10)$$

where AI_t is the contemporaneous asymmetry index. This scales positions up when asymmetry is pronounced and down when distributions approach symmetry.

Execution:

- Entry: Monday open following Friday signal generation
- Exit: Friday close or upon signal reversal
- Rebalancing: Weekly (end of Friday close)
- No leverage; positions bounded to $[0.5, 2.0]$ units

Momentum Benchmark: Weekly time-series momentum (Moskowitz et al., 2012): long one unit when the trailing 20-week return is positive, short one unit when it is negative, so the position is the sign of the trailing return and the strategy is always in the market once the lookback window fills.

Mean Reversion Benchmark: Contrarian positions on the z -score of the weekly close against its 20-week moving average: short one unit when $z_t > 2$, long one unit when $z_t < -2$, flat otherwise (positions close when $|z|$ falls back inside the two-standard-deviation band).

Buy & Hold: A constant long position of one unit, entered at the start of the sample and held throughout; its cumulative return equals the EUR/JPY spot move over the sample.

The momentum and mean-reversion benchmarks are the same 20-week momentum and 2.0σ mean-reversion rules that enter the data-snooping universe of Section 5.7, evaluated on the same weekly execution plumbing as the asymmetry strategy (positions applied to the following week’s return). Performance metrics include cumulative return, annualized volatility, Sharpe ratio, Sortino ratio, maximum drawdown, and trade count.

3 Results

3.1 Asymmetry Detection

Table 1 presents distributional statistics for each alpha type.

Table 1: Alpha Asymmetry Metrics				
Alpha Type	Skew	Kurt	AI	POR
Tail	−1.47	19.00	0.17	2.78%
Fast	0.01	−0.00	0.96	51.98%
Pricing	−0.17	−0.73	0.80	57.54%
Coverage	1.75	7.57	3.45	53.17%
Hedge	0.15	−0.52	1.40	53.57%

Note: Skew = skewness, Kurt = excess kurtosis, AI = asymmetry index (variance of positive deviations over variance of negative deviations), POR = positive observation ratio (proportion of positive values). All statistics computed on the signed weekly signal series ($n = 504$) generated by the replication pipeline (`analysis/full_pipeline.py`). An earlier version of this table reported tail-signal skewness of 5.05; that figure described the *unsigned* exceedance magnitude, which is non-negative and hence right-skewed by construction. The signed series the strategy actually trades on is reported here.

The corrected statistics tell a very different story from the one asymmetry-based trading would require. The single robust asymmetry is in *coverage* alpha—the volatility-expansion signal—which is strongly right-skewed (1.75) with heavy tails (7.57) and an asymmetry index of 3.45: volatility expands abruptly and contracts gradually, so the ratio $\sigma_{20}(t)/\sigma_{20}(t-5)$ spikes upward far more violently than it falls. This is a well-understood feature of volatility dynamics rather than a tradable anomaly in itself. Tail alpha, by contrast, skews *negative* (−1.47): conditional on exceeding the rolling 95th-percentile magnitude threshold, large EUR/JPY moves are more often large depreciations than large appreciations over this sample—the opposite of the “positive outlier” narrative, and closer to the crash-risk asymmetry familiar from the carry literature (Burnside et al., 2011). Its POR of 2.78% reflects that the series is zero except at exceedances. Fast alpha is, by construction, a normalized return z -score, and behaves like one: skewness 0.01, excess kurtosis indistinguishable from zero. Pricing and hedge alphas are likewise near-symmetric. Statistical detectability—already limited to one series—does not imply economic exploitability, a distinction the remainder of the paper examines.




Figure 1: Alpha Asymmetry Analysis. Distributions of the five signed weekly alpha signals (EUR/JPY, 2015–2025, $n = 504$) with skewness point estimates and 95% circular block bootstrap confidence intervals; the final panel summarizes the dependence-robust skewness inference. Only coverage alpha’s interval excludes zero.

3.2 Statistical Significance

Table 2 presents comprehensive hypothesis test results for each alpha type. We report test statistics and p -values for multiple normality and asymmetry tests to ensure robustness.

Table 2: Statistical Test Results

Alpha	$\hat{\gamma}_1$	t -stat	Block CI	SW	JB	K^2	LB $Q(4)$
Tail	-1.47	-13.5***	[-3.09, 0.54]	0.335***	7759***	244.4***	12.6*
Fast	0.01	0.1	[-0.21, 0.21]	0.997	0.0	0.0	6.5
Pricing	-0.17	-1.5	[-0.44, 0.09]	0.984***	13.4**	30.1***	679.1***
Coverage	1.75	16.0***	[1.17, 2.15]	0.865***	1459***	217.6***	72.6***
Hedge	0.15	1.3	[-0.25, 0.54]	0.990**	7.5*	11.7**	1716.2***

Note: t -stat tests $H_0 : \gamma_1 = 0$ with $SE = \sqrt{6/n} = 0.109$ ($n = 504$) and assumes i.i.d. observations; Block CI = 95% circular block bootstrap confidence interval for $\hat{\gamma}_1$ (block length 13 weeks, $B = 2000$), which is robust to the serial dependence the i.i.d. tests ignore; SW = Shapiro-Wilk statistic; JB = Jarque-Bera statistic; K^2 = D’Agostino-Pearson omnibus; LB $Q(4)$ = Ljung-Box statistic at four lags, reported per series. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (unadjusted); the qualitative conclusions below apply Bonferroni correction at family size $m = 5$. Because pricing, coverage, and hedge alphas are built from overlapping rolling windows, their Ljung-Box statistics are large by construction and the block-bootstrap intervals are the appropriate basis for skewness inference.

Normality. Four of five alpha types reject normality; the exception is instructive. Fast alpha—a return z -score by construction—is statistically indistinguishable from Gaussian on every test (SW $p = 0.55$, JB ≈ 0 , K^2 $p = 0.99$), a useful check that the measurement pipeline does not manufacture non-normality on its own. Tail and coverage alphas depart severely (SW 0.335 and 0.865); pricing and hedge reject mildly, driven by shape rather than asymmetry.

Skewness. Under the i.i.d. t -test, tail and coverage alphas show large-magnitude skewness of opposite signs. But the alphas are serially dependent—the three built from overlapping windows have Ljung-Box statistics ranging from 72.6 to 1716—so the block-bootstrap intervals are the honest inferential yardstick. On that yardstick only *coverage* alpha survives: its interval $[1.17, 2.15]$ excludes zero decisively, while the tail-alpha interval $[-3.09, 0.54]$ widens to include zero once dependence is respected (the point estimate is dominated by a small number of clustered exceedance weeks). Fast, pricing, and hedge alphas are consistent with symmetry throughout.

What survives. The asymmetry inventory that motivates the rest of the paper is therefore modest: one robustly right-skewed volatility-expansion signal, one negatively-skewed-but-fragile tail signal, and three symmetric signals. This is consistent with the broader stylized facts on asset-price distributions (Cont, 2001; Mandelbrot, 1963) while falling well short of the pervasive “alpha asymmetry” a trading strategy would want.

3.3 Backtest Performance

Table 3 compares the asymmetry strategy against the three benchmarks of Section 2.4 on EUR/JPY over the full sample period (2015–2025). We report standard performance measures alongside risk-adjusted metrics to enable fair comparison across strategies with different return and volatility profiles.

Table 3: Strategy Performance Comparison (EUR/JPY, 2015–2025)

Strategy	Return	Vol	Sharpe	Sortino	MDD	Trades
Asymmetry	3.60%	2.69%	0.149	0.062	−7.96%	17
Momentum (20w)	−18.92%	8.40%	−0.215	−0.309	−31.32%	27
Mean Rev. (2.0 σ)	10.74%	3.02%	0.364	0.198	−4.32%	27
Buy & Hold	33.53%	8.60%	0.390	0.584	−15.78%	1

Note: Return = cumulative gross return over the 504-week sample; Vol = annualized volatility; Sharpe = annualized Sharpe ratio (risk-free rate = 0); Sortino = annualized mean return over the standard deviation of negative-return weeks; MDD = maximum drawdown; Trades = completed round trips (position-change events divided by two, a sign flip counting as one event), with Buy & Hold’s single initial purchase counted as one trade. Benchmark definitions are given in Section 2.4 and are identical to the corresponding candidates in the Section 5.7 data-snooping universe. Transaction costs not included. Generated by `analysis/full_pipeline.py`; an earlier version of this table reported benchmark rows (momentum −15.66%, mean reversion 34.03%, buy-and-hold 2.31%) that traced to no committed code—the buy-and-hold figure was inconsistent with the +33.5% EUR/JPY spot move over the sample—and all benchmark rows have been recomputed from the replication pipeline.

Risk-Adjusted Performance. The comparison is unflattering to the asymmetry strategy. Its cumulative gross return of 3.60% is an order of magnitude below the 33.53% buy-and-hold return—which is simply the EUR/JPY spot appreciation over the decade—and its Sharpe ratio (0.149) trails both buy-and-hold (0.390) and the plain 2.0 σ mean-reversion benchmark (0.364). The strategy’s one distinction is its low volatility (2.69% annualized), which reflects inactivity rather than skill: it is flat in roughly 95% of weeks, holding a position only when the skewness threshold is exceeded. The 20-week momentum benchmark loses money outright (−18.92%), consistent with the absence of exploitable trend structure at this frequency.

Drawdown Analysis. The asymmetry strategy’s maximum drawdown (−7.96%) is about half of buy-and-hold’s (−15.78%) and a quarter of momentum’s (−31.32%), but the mean-reversion benchmark is more contained still (−4.32%). Like the low volatility, the modest drawdown is a mechanical consequence of being out of the market most of the time rather than evidence of downside management.

Trading Frequency. All four strategies trade sparsely at weekly frequency: 17 round trips for the asymmetry strategy over the decade, 27 each for the momentum and mean-reversion benchmarks, and a single purchase for buy-and-hold. The asymmetry strategy’s sparsity—a direct consequence of the conservative skewness threshold—keeps transaction cost drag negligible (Section 5.4), but it also means the return estimate rests on few independent observations, an important caveat for the statistical inference that follows.

Statistical Significance of Returns. To assess whether strategy returns are statistically distinguishable from zero, we conduct stationary-bootstrap inference (expected block length 4 weeks, $B = 2000$). The asymmetry strategy’s annualized return of 0.37% (the single annualization of the 3.60% cumulative figure used throughout this paper) has a 95% confidence interval of $[-1.14\%, 2.12\%]$, and the annualized Sharpe ratio’s interval spans $[-0.53, 0.66]$. Both include zero with room to spare: the strategy’s returns are statistically indistinguishable from no edge at all, before any consideration of costs. This—not transaction costs, which Section 5.4 shows to be immaterial at 17 trades per decade—is the paper’s economic null.

Strategy equity curves and drawdowns, EUR/JPY weekly 2015–2025 ($n = 504$)




Figure 2: Strategy equity curves and drawdowns. Cumulative gross returns (top) and draw-down profiles (bottom) for the asymmetry strategy, the 20-week momentum and 2.0σ mean-reversion benchmarks, and buy-and-hold on the 504-week EUR/JPY sample. Series generated by `analysis/make_backtest_figure.py` from the Table 3 strategy definitions.

3.4 Cross-Market Validation

To test whether the EUR/JPY findings generalize, we apply the identical alpha construction and the identical trading strategy—unchanged code, unchanged parameters—to GBP/USD, SPY, and GLD (Section 2.1). Hedge alpha is excluded as EUR/JPY-specific; the four remaining signals and the two-sided asymmetry strategy (0.75 threshold) transfer directly. Table 4 reports the signed weekly skewness of each signal alongside the strategy’s gross cumulative return and, for scale, each market’s buy-and-hold return over the same 504-week window.

Table 4: Cross-Market Check

Market	Tail Skew	Fast Skew	Pricing Skew	Coverage Skew	Strat. Ret.	B&H Ret.
EUR/JPY	−1.47	0.01	−0.17	1.75	3.60%	33.53%
GBP/USD	0.82	0.01	−0.14	1.57	17.18%	−7.63%
SPY	0.94	−0.20	−1.08	1.84	11.66%	294.30%
GLD	0.69	−0.00	−0.19	1.32	−30.26%	200.97%

Note: Skewness of the signed weekly alpha signals ($n = 504$ per market), constructed identically to EUR/JPY; Strat. Ret. = gross cumulative return of the two-sided asymmetry strategy at the 0.75 threshold (17, 22, 5, and 18 round trips respectively); B&H Ret. = buy-and-hold return over the same window. Skewness values are point estimates without the block-bootstrap intervals of Table 2; the exercise is illustrative (see Scope Limitation below). Generated by `analysis/full_pipeline.py`; an earlier version of this table reported aggregated “MR/TF/HAT” skewness categories and returns that traced to no committed code, and has been replaced.

Two observations, both unfavorable to generalization. First, the distributional signature does not travel: the tail signal’s negative skew—the closest thing EUR/JPY offers to a distinctive asymmetry—

reverses to positive in all three other markets (0.82, 0.94, 0.69). The only pattern that recurs is coverage alpha's right skew (1.57–1.84 across markets), which is the generic volatility-expansion asymmetry Section 3.1 already identified as a well-understood feature of volatility dynamics rather than a tradable signal. Second, strategy performance is erratic in exactly the way signal-free noise would produce: the identical rules earn +17.18% on GBP/USD (a market that *fell* 7.63% over the window), earn +11.66% on SPY while buy-and-hold earned 294.30%, and lose 30.26% on GLD while the metal nearly tripled. None of these figures carries the dependence-robust inference of Section 3.3, and given that the EUR/JPY return—the one we do test formally—is statistically indistinguishable from zero, the natural reading of the cross-market spread is sampling variation around no edge rather than market-specific exploitability.

Scope Limitation. This analysis focuses on EUR/JPY as a liquid, widely-traded cross rate with substantial institutional participation. Cross-market results (GBP/USD, SPY, GLD) are presented as an illustrative transfer check rather than a definitive multi-asset analysis. A comprehensive treatment would require: (i) asset-specific alpha calibration reflecting different market microstructures; (ii) market-specific transaction cost modeling (FX spreads vs. equity commissions vs. futures margins); and (iii) asset-by-asset statistical inference with proper multiple-testing corrections. We leave this extension to future work; on the evidence here, nothing suggests the EUR/JPY null would reverse elsewhere.

4 Discussion

4.1 Why Asymmetry Exploitation Fails

The empirical findings reveal a pattern more thoroughgoing than “asymmetry exists statistically but fails economically”: under corrected measurement, most of the asymmetry never existed to begin with, and what exists is not tradable. We identify four failure modes:

Little Asymmetry to Exploit. The dependence-robust inventory of Section 3 leaves one robustly skewed signal (coverage, a volatility-expansion effect) and one fragile negative skew (tail). The strategy's entry signals are built on rolling skewness of fast and pricing alphas—two series whose full-sample skewness is statistically zero. A strategy harvesting asymmetry from signals that are not asymmetric has nothing to harvest; its 17 in-sample trades are threshold artifacts of local skewness noise.

No Heavy Tails. The GPD analysis shows $\xi \approx 0$: exceedance magnitudes decay exponentially rather than by power law. Asymmetric strategies implicitly assume rare large gains compensate for frequent small losses; without genuine heavy tails—and with the signed exceedances leaning negative—this compensation does not materialize.

Inactivity Under Honest Selection. The walk-forward shows that when the threshold is chosen on training data alone, the strategy goes almost entirely dormant (3 trades in eight years). The in-sample activity that makes the backtest look evaluable is itself a product of full-sample threshold choice.

Non-Generalizability. Cross-market transfer yields erratic results—the identical rules gain 17% on GBP/USD, trail buy-and-hold by nearly 300 percentage points on SPY, and lose 30% on GLD—and the tail signal's negative skew reverses sign in every other market tested. Nothing about the EUR/JPY configuration replicates elsewhere.

Transaction costs, notably, are absent from this list: at 17 trades per decade they shave basis points (Section 5.4). The failure is statistical, not frictional.

4.2 What the Asymmetry Statistics Do Show

The corrected statistics still reveal genuine features of market structure, just not the ones an asymmetry trader would want:

Volatility-Expansion Asymmetry: Coverage alpha’s robust right skew (1.75, block CI [1.17, 2.15]) quantifies a familiar asymmetry in volatility dynamics— σ spikes abruptly and decays slowly (Bollerslev, 1986; Engle, 1982). It is real, persistent, and already priced into every options desk’s models; nothing in our results suggests it leaks into spot-return predictability.

Downside-Leaning Tails: The signed tail signal skews negative—conditional on a large weekly move, depreciations outnumber appreciations—which is the opposite of a lottery-ticket structure (Kumar, 2009; Bali et al., 2011) and closer to the crash-risk asymmetry of the carry literature (Brunnermeier et al., 2009; Burnside et al., 2011). The fragility of even this estimate under block-bootstrap inference (CI [−3.09, 0.54]) counsels against building anything on it.

A Symmetric Baseline: Fast alpha’s textbook normality is a finding in its own right: a volatility-normalized weekly return in a major pair is, to statistical precision, Gaussian. Whatever exploitable structure exists in EUR/JPY, it does not live in the low moments of normalized short-horizon returns (Moskowitz et al., 2012).

4.3 Limitations

Several limitations warrant acknowledgment. First, the 504-observation sample, while spanning a decade, may not capture all market regimes—particularly the kind of volatility clustering documented by Bollerslev (1986) and Engle (1982). Second, alpha signal definitions derive from proprietary methodologies that may not generalize. Third, the factor-attribution exercise uses single-market spot proxies rather than the published cross-sectional FX factors, and with 25 in-position weeks it has little power in either direction; the cost model, while now consistent, uses stylized spread tiers rather than realized quotes (professional practice incorporates richer friction structure; Menkhoff, 2010). Fourth, the cross-market check transfers the EUR/JPY constructions unchanged—thresholds, windows, and strategy parameters are not recalibrated per market, no market-specific cost model is applied, and no dependence-robust inference is attempted—so it can illustrate non-generalization but cannot establish market-specific results; asset-specific dynamics such as the asymmetric volatility documented by Glosten et al. (1993) are not separately modeled. Fifth, the underlying data are retrieved from a public aggregator (Yahoo Finance) whose historical series can be revised; the replication pipeline records retrieval counts and dates to make any drift visible.

5 Robustness Analysis

The preceding sections leave a thin asymmetry inventory—one robustly skewed signal among five—and an in-sample backtest whose returns are statistically indistinguishable from zero. This section stress-tests what remains, addressing four questions: (1) Do results hold out-of-sample? (2) Are patterns stable across market regimes? (3) Do trading costs erode strategy returns? (4) Is asymmetry alpha distinct from known FX risk factors?

5.1 Out-of-Sample Validation

A primary concern with any trading strategy is in-sample overfitting. Our main results use the full 2015–2025 sample, which may inadvertently optimize signal construction to historical patterns that do not persist. We address this through walk-forward analysis, a methodology widely used in quantitative finance to assess genuine predictive power (Lo et al., 2000).

We implement an expanding-window walk-forward: the entry threshold is selected from the grid $\{0.50, 0.75, 1.00, 1.25\}$ by Sharpe ratio on all data through year $t - 1$ (beginning with 2015–2017), then the strategy runs untouched on calendar year t , rolling forward annually through 2025. All signal con-

struction uses trailing windows only, so no test-year information enters signal values; the walk-forward isolates the threshold-selection step, which is the strategy’s only free parameter. Table 5 presents out-of-sample performance metrics.

Table 5: Out-of-Sample Walk-Forward Results

Test Year	Selected Threshold	OOS Return	OOS Sharpe	Hit Rate	Trades
2018	1.00	0.00%	0.00	—	0
2019	1.00	0.00%	0.00	—	0
2020	1.00	0.60%	0.52	50.0%	1
2021	1.00	1.85%	1.04	66.7%	2
2022	1.00	0.00%	0.00	—	0
2023	1.00	0.00%	0.00	—	0
2024	1.00	0.00%	0.00	—	0
2025	1.00	0.00%	0.00	—	0
Pooled OOS	—	2.46%	0.42	60.0%	3

Note: OOS = out-of-sample. Threshold selected by Sharpe ratio on the expanding training window; hit rate = proportion of profitable in-position weeks; Sharpe annualized assuming risk-free rate = 0. Pooled return is cumulative over all eight test years ($\approx 0.32\%$ annualized). 2025 covers January–August. Generated by `analysis/full_pipeline.py`; an earlier version of this table reported 141 pooled trades and could not be reproduced from the strategy specification.

The walk-forward verdict is not the graceful in-sample-to-out-of-sample degradation typical of over-fit strategies; it is near-total inactivity. In every training window the Sharpe criterion selects the 1.00 threshold (whose four in-sample trades produce the highest risk-adjusted return), and at that threshold the strategy fires three times in eight out-of-sample years, remaining flat through 2018–2019 and 2022–2025 entirely. The pooled cumulative return of 2.46% (roughly 0.32% annualized) rests on three trades; the pooled Sharpe of 0.42 is an artifact of a return series that is zero almost everywhere and should not be read as evidence of skill. The honest summary is that the asymmetry signal, disciplined by its own training data, almost never speaks—and when it does, the sample is far too small for inference. This inactivity is itself informative: the in-sample trade count (17 at the 0.75 baseline) was already sparse, and out-of-sample threshold selection reveals even that sparse activity to be threshold-fragile rather than a stable market phenomenon.

5.2 Subsample Stability

Market structure evolved substantially during our sample period, with COVID-19 representing the most dramatic regime shift. We test whether asymmetry patterns persist across distinct market environments using subsample analysis.

Table 6 partitions the sample into economically meaningful subperiods.

Table 6: Subsample Asymmetry Stability

Subsample	N	Tail Skew	Fast Skew	Price Skew	Strategy Ret.
Pre-COVID (2015–2019)	208	-3.37	-0.05	0.12	1.06%
COVID Shock (2020)	52	4.87	0.09	-0.07	-0.19%
Post-COVID (2021–2025)	244	-0.73	0.05	-0.38	2.71%
Low VIX ($VIX < 20$)	342	-3.24	0.04	-0.20	2.38%
High VIX ($VIX \geq 20$)	162	0.62	-0.06	-0.10	2.67%
Rate Hike Regime (2022–2025)	191	-0.70	-0.13	-0.52	2.86%

Note: VIX thresholds based on monthly average. Skewness coefficients computed within each subsample. Strategy returns are cumulative within-period returns.

Temporal Stability. Subsample skewness is *not* stable for tail alpha: it is strongly left-skewed in the pre-COVID period (-3.37) and again post-COVID (-0.73), but flips to pronounced right-skewness during the 2020 COVID shock ($+4.87$). This sign reversal is exactly the signature of a regime-dependent statistic—the extreme positive returns that dominate the COVID window are absent in calmer periods, where the tail-exceedance series is dominated by negative outliers. Fast and pricing alphas, by contrast, are close to symmetric in every subsample ($|\text{skew}| \leq 0.13$), so the asymmetry that motivates the strategy is concentrated in the tail signal and is itself regime-contingent rather than a stable structural feature.

Regime Conditioning. VIX conditioning tells a similar story. The low-VIX subsample shows strong negative tail skewness (-3.24), while the high-VIX subsample is weakly positive ($+0.62$); fast and pricing skewness are near zero in both. Strategy returns are comparable across VIX regimes (2.38% low-VIX vs. 2.67% high-VIX), offering no evidence that volatility conditioning sharpens the exploitable edge.

5.3 Entry Threshold Sensitivity

The baseline skewness threshold (0.75) deserves scrutiny: is it arbitrary or robustly selected? Table 7 reports strategy performance across alternative thresholds, holding other parameters constant.

Table 7: Sensitivity to Entry Threshold

Threshold	Return	Sharpe	MDD	Trades	Hit Rate
0.50	-6.91%	-0.15	-13.9%	37	40.8%
0.75	3.60%	0.15	-8.0%	17	48.0%
1.00	2.96%	0.44	-1.0%	4	66.7%
1.25	0.00%	0.00	0.0%	0	—

Note: Returns use the full asymmetry strategy specification (long and short legs with position sizing as described in Section 2.4). Performance varies non-monotonically with threshold. Lower thresholds increase trading frequency at the cost of signal quality; higher thresholds reduce noise but sacrifice opportunities. Simplified long-only implementations yield different (typically lower) returns due to the absence of the short-side component.

Threshold Selection. The 0.75 baseline is a full-sample choice, and we make no pretense otherwise: it corresponds approximately to the 75th percentile of historical rolling skewness magnitudes, and its selection benefits from in-sample knowledge. The out-of-sample discipline lives in the walk-forward of Section 5, where the threshold is re-selected each year on training data only—and there the Sharpe criterion consistently prefers 1.00, at which the strategy barely trades. The sensitivity analysis shows why the choice matters so much: the 0.50 threshold generates excessive trading (37 trades) with diluted

signal quality, producing a negative return (-6.91%); the 1.25 threshold never fires. Only the 0.75 and 1.00 thresholds produce positive returns, and the 1.00 threshold’s higher Sharpe (0.44) rests on 4 trades. This fragility—where the strategy’s viability depends on precise threshold calibration within a narrow band—is itself part of the paper’s null finding: the asymmetry signal is not robust enough to constitute a reliable trading strategy.

5.4 Transaction Cost Sensitivity

Backtests without transaction costs overstate implementable returns. We analyze cost sensitivity across scenarios ranging from zero-cost (theoretical benchmark) to retail-level spreads.

The EUR/JPY market is highly liquid, with typical spreads of 0.5–1.0 pips during active trading hours (King et al., 2013). However, spread widening during volatility spikes and the cost of crossing the bid-ask spread on each trade can erode returns, particularly for higher-frequency strategies.

Table 8: Transaction Cost Analysis

Cost Scenario	Round-Trip Cost (pips)	Net Return	Net Sharpe
Zero Cost (Baseline)	0.0	3.60%	0.149
Prime Brokerage	0.3	3.54%	0.147
Institutional	0.7	3.47%	0.144
Retail (Tight)	1.3	3.35%	0.140
Retail (Wide)	2.0	3.22%	0.135
Break-Even	19.2	0.00%	0.000

Note: Round-trip cost = bid-ask spread plus market impact (pips); each position-change leg is charged half the round-trip cost in proportion to the position change, at the sample-average EUR/JPY price. Carry/financing is excluded uniformly from all scenarios (the strategy is in position roughly 5% of weeks, so the carry differential contributes at most a few basis points either way). Net returns cumulative over 2015–2025. Generated by `analysis/full_pipeline.py`; an earlier version of this table showed net return *rising* with costs at low tiers, an inconsistency in the carry treatment across rows.

Cost Components. We model execution costs as bid-ask spread plus market impact, charged per position change. Financing/carry is excluded from all scenarios symmetrically: with 17 round trips and roughly 5% of weeks in position over a decade, the EUR–JPY carry differential moves net returns by basis points, not percentage points, and including it asymmetrically is how cost tables produce impossible orderings.

The Real Cost Story: Immateriality. The table’s message is not that costs erode the edge—it is that costs barely register. A strategy that trades 17 times in a decade pays the spread 17 times; even the wide-retail scenario (2.0 pips round-trip) costs only 38 basis points cumulatively, leaving net return within 0.4 percentage points of gross. Break-even requires an implausible 19.2 pips round-trip. Transaction costs therefore cannot carry the paper’s null result, and we do not ask them to: the null rests on the fact that the gross return itself—3.60% over a decade, Sharpe 0.15, confidence interval spanning zero—is statistically indistinguishable from no edge at all. Costs are immaterial rather than decisive, which is exactly what trade sparsity implies (Kyle, 1985).

5.5 Factor Attribution

Does asymmetry alpha merely proxy for known FX risk factors? Prior work documents systematic risk premia in currency markets associated with carry (Lustig et al., 2011), momentum (Menkhoff et al., 2012), and the dollar factor (Verdelhan, 2018), with Lettau et al. (2014) showing that downside risk conditioning explains a substantial portion of currency risk premia across asset classes. A parallel factor

decomposition approach applied to cryptocurrency markets (Farzulla, 2026a) similarly finds that structural factors dominate idiosyncratic signal variation. If asymmetry signals simply load on these factors, our contribution reduces to factor timing rather than genuine alpha discovery.

We regress asymmetry strategy returns on proxy FX factors:

$$R_{\text{asym},t} = \alpha + \beta_1 \cdot \text{Carry}_t + \beta_2 \cdot \text{Mom}_t + \beta_3 \cdot \text{Dollar}_t + \varepsilon_t \quad (11)$$

where the factors are honest single-market proxies constructible from our data rather than the published cross-sectional factors: Carry is the average weekly spot return of the classic carry crosses AUD/JPY and NZD/JPY (a spot-only proxy for the Lustig et al., 2011 carry factor, omitting interest accrual), Mom is the return of a 12-week time-series momentum rule on EUR/JPY itself (in the spirit of Menkhoff et al., 2012), and Dollar is the DXY weekly return (proxying Verdelhan, 2018). Constructing the published factors requires a G10 forward-rate cross-section we do not use; the proxies capture the same directional exposures at the single-pair level, and we label them as proxies throughout. Because the strategy is out of the market most weeks, we estimate the regression twice: on the full weekly sample (where the dependent variable is structurally zero in out-of-position weeks) and on in-position weeks only, whose count is the honest effective sample size.

Table 9: Factor Attribution Regression (HAC-Corrected)

Variable	Full sample ($n = 491$)		In-position weeks ($n = 25$)	
	Coefficient	t -stat	Coefficient	t -stat
Intercept (α)	0.00008	0.48	0.0030	0.67
Carry (β_1)	−0.0007	−0.05	0.067	0.14
Momentum (β_2)	−0.0070	−0.37	−0.247	−0.51
Dollar (β_3)	−0.0123	−0.83	−0.502	−0.84
R^2	0.001		0.067	
F -statistic	0.48		0.62	

Note: Newey-West HAC standard errors with 4 lags (Newey and West, 1987). Dependent variable is weekly asymmetry strategy returns; factors are the single-market proxies described in the text. The full-sample column includes the out-of-position weeks in which the dependent variable is structurally zero; the in-position column is the honest effective sample. No coefficient in either specification is significant at the 10% level. Generated by `analysis/full_pipeline.py`; an earlier version of this table reported a marginally significant intercept of 21 bps weekly (10.9% annualized) with $F = 20.8$, which could not be reproduced and was internally inconsistent with its own R^2 and with the 3.60% cumulative backtest return.

No Residual Alpha. The intercept is indistinguishable from zero in both specifications ($t = 0.48$ full sample; $t = 0.67$ on in-position weeks). There is no factor-adjusted alpha to explain, and consequently no tension between regression alpha and realized returns to reconcile: a strategy earning 3.60% cumulative over a decade has a weekly mean of well under one basis point, and that is what the regression finds.

Factor Loadings and Effective Sample. No factor loading approaches significance, and R^2 is 0.001 on the full sample—but both facts are dominated by a structural feature rather than by orthogonality: the strategy is out of the market in roughly 95% of weeks, so the full-sample dependent variable is mostly exact zeros and the informative sample is the 25 in-position weeks. With an effective n of 25, the factor regression has essentially no power, and we present it as a specification-honesty exhibit rather than as evidence that asymmetry returns are orthogonal to factor risk. The one defensible conclusion is negative: nothing in this exercise supports a residual-alpha claim.

5.6 Tail Distribution Analysis

A concern with any exceedance-based signal is whether the extreme observations that drive it reflect a stable tail distribution or sparse outliers that may not recur out-of-sample. We characterize the magnitude distribution of tail events using Extreme Value Theory (McNeil and Frey, 2000; Coles, 2001), bearing in mind the corrected result of Table 1: the *signed* tail signal skews negative (large moves are more often depreciations), so the question here is the shape of the exceedance-magnitude tail, not an upside-asymmetry story.

Exceedance Preprocessing. Financial tail events exhibit temporal clustering (volatility clustering), violating the GPD independence assumption. We apply runs declustering with minimum inter-exceedance separation of 5 weeks, following Coles (2001). Of the 26 raw exceedances above the 95th percentile threshold, declustering yields 20 cluster maxima used for GPD estimation. We estimate the extremal index $\theta = 0.83$ via the Ferro–Segers intervals estimator (Ferro and Segers, 2003), with a 95% confidence interval of $[0.60, 1.00]$ obtained by bootstrap resampling of the inter-exceedance times themselves ($B = 2000$)—a scheme that preserves the gap distribution the estimator is built on and, unlike a naive index bootstrap, produces an interval containing its own point estimate.

Methodology. We fit a Generalized Pareto Distribution (GPD) to cluster maxima above the 95th percentile threshold. The GPD is characterized by shape parameter ξ (determining tail heaviness) and scale parameter σ . Positive ξ indicates Pareto-type heavy tails; $\xi = 0$ corresponds to exponential decay; negative ξ implies bounded tails.

Table 10: GPD Parameter Estimates—Tail Alpha (Declustered)

Parameter	Estimate	95% CI
Shape (ξ)	−0.25	[−1.62, 0.25]
Scale (σ)	0.0089	—
Threshold (95th pctl)	0.0244	—
Extremal index (θ)	0.83	[0.60, 1.00]
Raw exceedances	26	—
Cluster maxima	20	—
KS test p -value	0.971	Fail to reject

Note: GPD fitted via maximum likelihood to cluster maxima above the 95th percentile after runs declustering (5-week separation). ξ CI from a nonparametric bootstrap of cluster maxima ($B = 2000$); θ CI from bootstrap resampling of inter-exceedance times ($B = 2000$). KS = Kolmogorov-Smirnov goodness-of-fit test. Generated by `analysis/full_pipeline.py`; an earlier version of this table reported a θ confidence interval of $[0.42, 0.63]$ that excluded its own point estimate—an invalid interval produced by a biased index-resampling scheme, replaced here.

Tail Characterization. The estimated shape parameter $\hat{\xi} = -0.25$ (95% CI: $[-1.62, 0.25]$) is not significantly different from zero, suggesting tails that are neither heavy (Pareto-type) nor strictly bounded (Weibull-type), but rather consistent with exponential decay. The wide confidence interval reflects the limited number of cluster maxima (20) available for estimation—a common challenge in EVT applications to financial data (McNeil and Frey, 2000). The Kolmogorov-Smirnov test fails to reject the GPD null hypothesis ($p = 0.97$), supporting the appropriateness of the EVT framework despite parameter uncertainty.

Clustering Dynamics. The extremal index $\theta = 0.83$ indicates weak temporal clustering in tail events: the mean cluster size is $1/\theta \approx 1.2$, and the confidence interval reaches 1.0 (no clustering). Exceedances arrive close to independently, with mild GARCH-type persistence (Bollerslev, 1986) at most. Clustering therefore does little to reduce the effective sample size for tail inference; the binding

constraint is simply the small number of exceedances.

Implications. The EVT analysis characterizes the *magnitudes* of extreme weekly moves and finds them well-behaved: exponential-type decay, no evidence of power-law tails, weak clustering. Combined with the corrected skewness results—the signed exceedances lean negative, and not robustly so—there is no distributional foundation here for an upside-capture trading story. Risk management can rely on finite-variance assumptions; traders hoping for lottery-like right tails in EUR/JPY weekly returns will not find them in this sample.

5.7 Multiple Testing Corrections

With five alpha types examined across multiple statistical tests, data-snooping concerns arise: apparent significance may reflect chance discoveries from multiple hypothesis testing. We address this through both formal corrections and bootstrap-based reality checks.

Bonferroni Corrections. Table 2 applies plain Bonferroni correction at family size $m = 5$ (one skewness test per alpha type; $\alpha^* = 0.01$ for the 5% family level). At the corrected threshold, four of five alpha types reject *normality*, but the asymmetry inventory is far thinner: only coverage alpha’s skewness survives both the correction and—more demanding—the block-bootstrap dependence adjustment. Fast alpha, indistinguishable from Gaussian on every test, serves as the natural control: the measurement pipeline does not manufacture non-normality where none exists.

Bootstrap Implementation. White’s Reality Check and Hansen’s SPA employ the stationary block bootstrap of Politis and Romano (1994) with expected block length $\ell = 4$ weeks. The candidate strategy universe comprises 13 strategies: simplified long-only asymmetry strategies using each of the 5 alpha types’ own rolling-skewness signal ($n = 5$), the full two-sided asymmetry strategy of Section 2.4 ($n = 1$), time-series momentum with 10-, 20-, and 40-week lookback windows ($n = 3$), mean-reversion at 1.5σ and 2.0σ thresholds ($n = 2$), an always-long (carry-proxy) strategy ($n = 1$), and a seeded random benchmark ($n = 1$). Bootstrap replications: 1,000. Loss function: negative weekly return, benchmark zero.

White’s Reality Check. Following White (2000), we test whether the best-performing candidate significantly outperforms across the universe. The Reality Check does not reject the null of no superior strategy ($RC = 0.02$, $p = 0.15$). More telling than the p -value is the identity of the best performer: the highest realized mean return in the universe belongs to the *random benchmark*—a coin-flip position sequence—which is precisely what one expects when no candidate carries genuine signal.

Hansen’s SPA Test. The Superior Predictive Ability test (Hansen, 2005) agrees ($SPA = 1.90$, $p = 0.28$): no candidate in the universe demonstrates superior predictive ability over the zero benchmark at any conventional level.

Table 11: Data-Snooping Correction Results

Test	Statistic	p -value	Interpretation
White’s Reality Check	0.02	0.15	Not significant
Hansen’s SPA	1.90	0.28	Not significant
# Candidate Strategies	13	—	—

Note: Candidate strategies: simplified asymmetry (5 alpha types), the full two-sided asymmetry strategy, time-series momentum (3 windows), mean-reversion (2 thresholds), always-long, and a seeded random benchmark. Stationary bootstrap with 1,000 replications, expected block length 4 weeks. Generated by `analysis/full_pipeline.py`; an earlier version of this table reported $RC = 2.14$ ($p = 0.042$), which could not be reproduced from any specification of the stated universe.

Interpretation. The data-snooping corrections complete the null result rather than qualifying it.

Nothing in a 13-strategy universe spanning asymmetry, momentum, mean-reversion, and carry beats a zero benchmark once dependence-aware resampling is applied—and the best raw performer is the random sequence. Whatever in-sample structure the asymmetry signals appeared to offer, it does not survive honest multiple-testing discipline.

6 Conclusions

This paper investigated whether distributional asymmetries in FX alpha signals represent exploitable inefficiencies. Our findings are negative at every layer:

1. Under dependence-robust inference, only one of five alpha signals (coverage, a volatility-expansion effect) exhibits robust skewness; the tail signal skews negative and fragilely so, and the remaining signals—including the two the trading strategy is built on—are statistically symmetric
2. EVT analysis reveals no heavy tails ($\xi \approx 0$) and weak clustering ($\theta = 0.83$); exceedance magnitudes are exponential-type, with no lottery-like right tail
3. Asymmetry-based trading generates a small positive gross return in-sample (3.60% cumulative, 17 trades), but bootstrap confidence intervals for both annualized return and Sharpe ratio comfortably include zero, and honest walk-forward selection leaves the strategy dormant (3 trades in eight out-of-sample years)
4. Transaction costs are immaterial rather than decisive at this trade frequency; the null is statistical, not frictional
5. Data-snooping corrections find no strategy in a 13-candidate universe with superior predictive ability—the best raw performer is the random benchmark—and the cross-market check finds neither the distributional signature nor any consistent strategy performance beyond EUR/JPY

These null findings contribute to the literature in three ways. First, they demonstrate the gap between statistical and economic significance in quantitative trading research—a pattern likely more common than publication bias reveals (Harvey, 2017). Second, they document, from the inside, how easily higher-moment “stylized facts” can be manufactured by measurement choices: this paper’s own earlier version reported pronounced positive tail skewness that turned out to describe an unsigned magnitude series, right-skewed by construction. The corrected analysis is a case study in why skewness claims should be accompanied by the sign conventions, sampling frequency, and dependence-robust intervals that make them checkable. Third, they provide a documented negative result that may prevent other researchers from pursuing this particular dead end.

The methodological lesson is sharper than “backtests flatter”: most of the asymmetry was never there. Distributional anomalies in financial data are easily overstated at the measurement stage, before a single trade is simulated—and the honest sequence of checks (signed variables, dependence-robust intervals, training-data-only selection, data-snooping correction) dismantled this one layer by layer.

Future research might examine whether asymmetry exploitation is viable at higher frequencies (intraday) where both signal counts and cost structures scale differently, or in markets with different microstructure. Parallel findings in cryptocurrency markets—where an apparent infrastructure-versus-regulatory volatility asymmetry dissolves once cross-asset dependence and heavy tails are handled correctly (Farzulla, 2026b)—and in cross-market event studies more broadly (Farzulla, 2025) reinforce that distributional asymmetries are easily overstated and rarely translate into exploitable strategies. Our results echo that the burden of proof for asymmetry-based strategies should be high: an apparently sharp

effect is insufficient evidence of economic value, and is often not robust evidence of a statistical effect either.

Acknowledgments

Portions of this manuscript were drafted and revised with assistance from Claude (Anthropic). The author retains full responsibility for all intellectual content, analytical decisions, and interpretive claims. This work was conducted as part of the Adversarial Systems Research program at Dissensus. Working papers and replication materials are available through the [Adversarial Systems & Complexity Research Initiative \(ASCRI\)](#). Comments and correspondence are welcome at murad@dissensus.ai.

Declarations

Conflict of Interest. The author declares no competing interests.

Funding. This research received no external funding.

Data Availability. EUR/JPY exchange rate data obtained from publicly available sources. Processed datasets and analysis outputs available at <https://github.com/studiofarzulla/alpha-asymmetry>.

Code Availability. Full replication code available at <https://github.com/studiofarzulla/alpha-asymmetry> under MIT License.

AI Assistance. Claude (Anthropic) assisted with manuscript drafting, code review, and statistical exposition. All research design, data analysis, and interpretation are solely the author's.

Author Contributions. Sole author.

References

- Clifford S. Asness, Tobias J. Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013. doi: 10.1111/jofi.12021.
- Turan G. Bali, Nusret Cakici, and Robert F. Whitelaw. Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics*, 99(2):427–446, 2011. doi: 10.1016/j.jfineco.2010.08.014.
- Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986. doi: 10.1016/0304-4076(86)90063-1.
- William Brock, Josef Lakonishok, and Blake LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance*, 47(5):1731–1764, 1992. doi: 10.1111/j.1540-6261.1992.tb04681.x.
- Markus K. Brunnermeier, Christian Gollier, and Jonathan A. Parker. Optimal beliefs, asset prices, and the preference for skewed returns. *American Economic Review*, 97(2):159–165, 2007. doi: 10.1257/aer.97.2.159.
- Markus K. Brunnermeier, Stefan Nagel, and Lasse Heje Pedersen. Carry trades and currency crashes. *NBER Macroeconomics Annual*, 23:313–348, 2009. doi: 10.1086/593088.
- Craig Burnside, Martin Eichenbaum, Isaac Kleshchelski, and Sergio Rebelo. Do peso problems explain the returns to the carry trade? *Review of Financial Studies*, 24(3):853–891, 2011. doi: 10.1093/rfs/hhq138.
- Stuart Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer, London, 2001. doi: 10.1007/978-1-4471-3675-0.
- Rama Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001. doi: 10.1080/713665670.
- Ralph B. D’Agostino, Albert Belanger, and Ralph B. D’Agostino Jr. A suggestion for using powerful and informative tests of normality. *The American Statistician*, 44(4):316–321, 1990. doi: 10.1080/00031305.1990.10475751.
- Robert F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50(4):987–1007, 1982. doi: 10.2307/1912773.
- Martin D.D. Evans and Richard K. Lyons. Order flow and exchange rate dynamics. *Journal of Political Economy*, 110(1):170–180, 2002. doi: 10.1086/324391.
- Eugene F. Fama. Efficient capital markets: A review of theory and empirical work. *The Journal of Finance*, 25(2):383–417, 1970. doi: 10.2307/2325486.
- Minyou Fan, Fearghal Kearney, Youwei Li, and Jiadong Liu. Optimizing currency factors. *Financial Review*, 2025. doi: 10.1111/fire.70000. Dynamic optimization of carry, momentum, and value currency factor strategies via spot and forward trading.

- Murad Farzulla. Infrastructure vs regulatory shocks: Asymmetric volatility response in cryptocurrency markets. *Research Square*, 2025. doi: 10.21203/rs.3.rs-8323026/v1. Under review at Digital Finance (Springer).
- Murad Farzulla. Do whitepaper claims predict market behavior? evidence from cryptocurrency factor analysis. *arXiv preprint arXiv:2601.20336*, 2026a. doi: 10.48550/arXiv.2601.20336. Preprint.
- Murad Farzulla. Do cryptocurrency markets differentiate infrastructure from regulatory shocks? a multi-moment event study with dependence-robust inference, 2026b. Zenodo concept DOI (v4.0.0), resolves to the current version. An arXiv posting (2602.07046) exists under the predecessor title pending replacement.
- Christopher A. T. Ferro and Johan Segers. Inference for clusters of extreme values. *Journal of the Royal Statistical Society: Series B*, 65(2):545–556, 2003. doi: 10.1111/1467-9868.00401.
- Lawrence R. Glosten, Ravi Jagannathan, and David E. Runkle. On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801, 1993. doi: 10.1111/j.1540-6261.1993.tb05128.x.
- Peter Reinhard Hansen. A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4):365–380, 2005. doi: 10.1198/073500105000000063.
- Campbell R. Harvey. Presidential address: The scientific outlook in financial economics. *The Journal of Finance*, 72(4):1399–1440, 2017. doi: 10.1111/jofi.12530.
- Campbell R. Harvey and Akhtar Siddique. Conditional skewness in asset pricing tests. *The Journal of Finance*, 55(3):1263–1295, 2000. doi: 10.1111/0022-1082.00247.
- Daniel Hertrich. *Implications of Changes in Conditional Value at Risk for the Currency Carry Trade and Equilibrium Currency Pricing: Empirical Evidence from the G10 Currencies*. PhD thesis, University of Bayreuth, 2025. urn:nbn:de:bvb:703-epub-8176-9. Forward-looking risk conditioning on G10 carry, momentum, and value strategies.
- Carlos M. Jarque and Anil K. Bera. Efficient tests for normality, homoscedasticity and serial independence of regression residuals. *Economics Letters*, 6(3):255–259, 1980. doi: 10.1016/0165-1765(80)90024-5.
- Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993. doi: 10.1111/j.1540-6261.1993.tb04702.x.
- Michael R. King, Carol L. Osler, and Dagfinn Rime. The market microstructure approach to foreign exchange: Looking back and looking forward. *Journal of International Money and Finance*, 38: 52–74, 2013. doi: 10.1016/j.jimonfin.2013.05.001.
- Alan Kraus and Robert H. Litzenberger. Skewness preference and the valuation of risk assets. *The Journal of Finance*, 31(4):1085–1100, 1976. doi: 10.1111/j.1540-6261.1976.tb01961.x.
- Alok Kumar. Who gambles in the stock market? *The Journal of Finance*, 64(4):1889–1933, 2009. doi: 10.1111/j.1540-6261.2009.01483.x.

- Albert S. Kyle. Continuous auctions and insider trading. *Econometrica*, 53(6):1315–1335, 1985. doi: 10.2307/1913210.
- Martin Lettau, Matteo Maggiori, and Michael Weber. Conditional risk premia in currency markets and other asset classes. *Journal of Financial Economics*, 114(2):197–225, 2014. doi: 10.1016/j.jfineco.2014.07.001.
- Andrew W. Lo, Harry Mamaysky, and Jiang Wang. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. *The Journal of Finance*, 55(4):1705–1765, 2000. doi: 10.1111/0022-1082.00265.
- Hanno Lustig, Nikolai Roussanov, and Adrien Verdelhan. Common risk factors in currency markets. *Review of Financial Studies*, 24(11):3731–3777, 2011. doi: 10.1093/rfs/hhr068.
- Benoit Mandelbrot. The variation of certain speculative prices. *The Journal of Business*, 36(4):394–419, 1963.
- Alexander J. McNeil and Rüdiger Frey. Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach. *Journal of Empirical Finance*, 7(3-4):271–300, 2000. doi: 10.1016/S0927-5398(00)00012-8.
- Lukas Menkhoff. The use of technical analysis by fund managers: International evidence. *Journal of Banking & Finance*, 34(11):2573–2586, 2010. doi: 10.1016/j.jbankfin.2010.04.014.
- Lukas Menkhoff, Lucio Sarno, Maik Schmeling, and Andreas Schrimpf. Currency momentum strategies. *Journal of Financial Economics*, 106(3):660–684, 2012. doi: 10.1016/j.jfineco.2012.06.009.
- Todd Mitton and Keith Vorkink. Equilibrium underdiversification and the preference for skewness. *Review of Financial Studies*, 20(4):1255–1288, 2007. doi: 10.1093/revfin/hhm011.
- Tobias J. Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. Time series momentum. *Journal of Financial Economics*, 104(2):228–250, 2012. doi: 10.1016/j.jfineco.2011.11.003.
- Christopher J. Neely, David E. Rapach, Jun Tu, and Guofu Zhou. Forecasting the equity risk premium: The role of technical indicators. *Management Science*, 60(7):1772–1791, 2014. doi: 10.1287/mnsc.2013.1838.
- Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.
- Dimitris N. Politis and Joseph P. Romano. The stationary bootstrap. *Journal of the American Statistical Association*, 89(428):1303–1313, 1994. doi: 10.1080/01621459.1994.10476870.
- S. S. Shapiro and M. B. Wilk. An analysis of variance test for normality (complete samples). *Biometrika*, 52(3/4):591–611, 1965. doi: 10.2307/2333709.
- Adrien Verdelhan. The share of systematic variation in bilateral exchange rates. *The Journal of Finance*, 73(1):375–418, 2018. doi: 10.1111/jofi.12587.
- Halbert White. A reality check for data snooping. *Econometrica*, 68(5):1097–1126, 2000. doi: 10.1111/1468-0262.00152.

A Enlarged Figures

For improved readability, this appendix reproduces all figures at enlarged scale.

Distributional properties of the five signed weekly alpha signals, EUR/JPY 2015–2025 (n = 504)




Figure 3: Alpha Asymmetry Analysis (enlarged). Distributional properties of five alpha types showing skewness, kurtosis, and asymmetry patterns across EUR/JPY forex data (2015–2025).

Strategy equity curves and drawdowns, EUR/JPY weekly 2015–2025 (n = 504)




Figure 4: Strategy equity curves and drawdowns (enlarged). Cumulative gross returns and drawdown profiles for the asymmetry strategy, the 20-week momentum and 2.0σ mean-reversion benchmarks, and buy-and-hold on the 504-week EUR/JPY sample.