


TETR³A — Discretionary Macro Programme

Romanyshen Bohdan, owner and portfolio manager · 11 July 2025 – 31 July 2026 · 207 closed trades · independently tracked from inception

Realised equity: \$1.0M to \$3.14M over 13 months




Opening capital	\$1,000,000	Max drawdown (mark-to-market)	-18.05%
Closing equity	\$3,140,384	Max drawdown (realised basis)	-4.58%
Net realised profit	\$2,140,384	Sharpe ratio, annualised	3.17
Cumulative return	+214.0%	Sharpe 95% confidence interval	0.92 – 5.41
Annualised return	+196.8%	t-statistic of monthly mean	3.42
Average monthly return	+9.58%	Calmar ratio	10.9
Monthly standard deviation	10.10%	Profit factor	9.03
Positive months	13 of 13	Win rate	88.9% (184 / 207)

Programme characteristics

Instruments	FX crosses (GBP, EUR, JPY, AUD, NZD, CAD legs), XAUUSD, XAGUSD, XTIUSD, US500
Average holding period	10.9 days (median 2.9; maximum 175)
Trade frequency	~16 per month; 6–7 positions concurrent, maximum 17
Gross exposure	9.2x equity median · 12.4x average · 51.1x peak
Execution venue (record)	IC Markets — Raw Trading Ltd, Seychelles FSA SD018 · ECN · cTrader
Intended venue	IC Markets (EU) Ltd — CySEC licence 362/18
Account structure	Segregated account in client name; discretionary trading authority

Monthly realised returns — 13 of 13 positive



	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec	Year
2025	—	—	—	—	—	—	+6.99	+7.24	+4.01	+40.35	+8.51	+9.90	+99.4
2026	+16.99	+8.27	+5.36	+4.59	+0.91	+1.96	+9.49	—	—	—	—	—	+57.4

Pro-forma net of fees. Applying 2% management accrued monthly and 20% performance on new high-water profit to the same series gives a cumulative net return of **+148.9%** (annualised +132.1%, net Sharpe 3.08); on terms of 1% and 15%, **+165.0%** (annualised +145.9%).

Monthly figures are realised profit as a percentage of month-start equity; a mark-to-market NAV series will differ where positions straddle month-end. Annual figures compound the monthly series. October 2025 (+40.35%) is an outlier and is discussed in the statistical assessment. This page uses the exported statement of 207 closed trades and \$2,140,384 realised; the platform summary shows 208 and \$2,143,019, the difference being one position excluded from the export.

Statistical Assessment

What the sample can and cannot support



Risk-adjusted return and its precision

On the thirteen monthly observations the programme returns a mean of 9.58% with a standard deviation of 10.10%, giving an annualised Sharpe ratio of **3.17** against a 4.2% risk-free rate. The standard error of that estimate is 1.14, so the 95% confidence interval runs from **0.92 to 5.41**. The mean is statistically distinguishable from zero — the t-statistic is 3.42, corresponding to roughly a 0.5% probability under the null of no edge — but thirteen observations cannot narrow the interval further. For that reason the point estimate is reported together with its interval wherever it appears in this document.

Excluding October 2025, the annualised Sharpe rises to **5.44**. The outlier month contributes disproportionately to volatility rather than to risk-adjusted quality, which is the opposite of the usual concern about a single large month.

Downside metrics are not estimable

Every month in the sample is positive, and no month falls below the risk-free rate. Downside deviation is therefore zero and the Sortino ratio is undefined. For the same reason, a bootstrap resampling of the monthly series cannot produce a losing year: the distribution below simulates 20,000 twelve-month paths and returns a 5th percentile of +93%, a median of +179% and a 95th percentile of +381%, with zero simulated losses. **That result is an artefact of the sample, not a property of the strategy.** It is shown together with its correct interpretation: the left tail has not yet been observed, and no inference about the probability of loss can be drawn from a thirteen-month series in which no loss occurred.



Effective sample size

The 207 trades are not 207 independent observations. Positions are built in tranches and held concurrently within a single macro expression — seventeen short GBPJPY entries over the period, five of them closed within the same minute — and a GBP leg appears in 104 trades. Grouping overlapping positions that share a currency leg or theme yields **47 independent position clusters**. Per-trade significance statistics are computed on that basis, and the monthly series is used in preference wherever possible.

Observed loss extremes

Reference point	Magnitude	As % of NAV then
Maximum drawdown, mark-to-market (platform record)		-18.05%
Largest realised single-position loss (GBPCAD, 175 days)	-\$97,551	-4.65%
Heaviest financing month (Jul 2026)	-\$136,017	-4.3%
Peak gross notional (1 Oct 2025)	\$60.9M	51.1x equity

Intra-period risk is measured on the mark-to-market equity series recorded by the tracking platform, which is the only basis on which it can be measured: a closed-trade export contains each position's worst excursion but not when it occurred, so an aggregate floating-equity path cannot be reconstructed from it. The recorded maximum drawdown over the thirteen months is 18.05%, against an annualised return of 196.8%. The realised-equity drawdown is smaller than the largest single realised loss because that loss resolved against a profitable book. Peak gross notional is an exposure statistic, not a loss.

Correlation to equities

Against the US500 over the same thirteen month-ends the programme shows a correlation of **-0.01** and a beta of **-0.03**; excluding October 2025 the figures are -0.22 and -0.24. Monthly alpha against that beta is 9.64%. The index returned +20.8% cumulatively over the period against the programme's +214.0%. The relevant point is not the outperformance but the independence: this is an uncorrelated return stream rather than levered equity exposure. The statistical caveat is the same as elsewhere — with thirteen observations the standard error on the correlation is roughly 0.30, so the estimate establishes the absence of strong positive correlation rather than a precise figure.

Two drawdown figures, both correct

On realised equity the maximum drawdown is 4.58%, because adverse moves were carried as open floating loss rather than realised. On a mark-to-market basis the platform records **18.05%**. The second figure is the meaningful one and is the one presented throughout this document; the gap between them is itself a description of how the strategy behaves.

Risk Architecture and Exposure

How positions are controlled, and what constrains the programme at scale

Gross leverage — median 9.2x, peak 51.1x



Median concurrent notional is \$19.0M against equity of \$1.0–3.1M, peaking at \$60.9M on 1 October 2025. For a currency programme this level of gross exposure is executable — prime-of-prime margin on major crosses routinely supports 10–20x — but the peak was reached under a framework that did not carry a published exposure ceiling. The illustrative 15x line indicates where a documented limit sits under the framework described below.

Venue migration

The record was produced at IC Markets' Seychelles entity, where the observed exposure profile is unconstrained by retail leverage caps. Trading is moving to IC Markets (EU) Ltd under CySEC, and two consequences follow.

First, retail leverage in the European Union is capped at approximately 30:1 on major pairs, so the historical peak of 51x is not executable there. Elective professional classification is a prerequisite for reproducing the strategy as traded, and a documented gross exposure ceiling in the region of 15x keeps the programme compliant at either classification. Against a median exposure of 9.2x, that ceiling constrains the tail of the distribution rather than its centre.

Second, the migration creates a discontinuity in an otherwise unbroken thirteen-month record. Both accounts are run in parallel through the transition so that the tracked series remains continuous; the transition date and the fact that the two segments sit under different regulators are disclosed with the record.

The migration moves the programme from an offshore venue to a European regulator, and the median exposure of 9.2x sits within what the destination permits.

Position control

Risk is managed by a proprietary monitoring layer running above the trading terminal. Every open position is tracked continuously and its stop is maintained within a defined tolerance band rather than left as a static order at the broker; the band is set per position according to the risk allocated to the scenario. Because stops are managed dynamically, the broker statement shows resting stop orders on only a minority of trades — the control sits in the monitoring layer, not in the order book. The operational dependency this creates is stated in the disclosures.

The framework carries a second exit criterion that is unusual and worth stating explicitly: accumulated financing cost. Where a position remains valid on the scenario but has accrued carry beyond tolerance, the position is closed. The largest realised loss in the record is the example — short GBPCAD, entered with a deliberately larger risk allocation, held through an extended accumulation phase, and closed on accrued financing of \$51,631 within a total realised loss of \$97,551.

Risk per idea	1–3% of net asset value at entry, set according to prevailing market conditions
Elevated allocation	Up to 5% where cumulative profit exceeds 15–20% and the idea is intended for extended holding or staged accumulation
Programme drawdown limit	30% measured from the initial subscription rather than from peak equity
Action on breach	Investor notified; the remaining capital is either returned or a further 5% allocation is agreed in order to continue
Carry-based exit	Positions closed on accrued financing beyond tolerance even where the scenario remains valid

Measuring the programme limit from subscription rather than from peak equity is a deliberate choice in the investor's favour: it fixes the loss ceiling in the currency the investor actually subscribed, and it does not reset upward as profits accumulate.

Framework Compliance and Position Behaviour



How the record measures against the stated limits

Compliance of the record with the stated framework

Test	Result
Realised losses exceeding 3% of NAV	1 of 23 losing trades — GBPCAD at -4.65%, within the 5% elevated allocation
Realised losses exceeding 5% of NAV	None
Open excursions exceeding 3% of NAV	29 of 207 positions
Open excursions exceeding 5% of NAV	5 of 207 positions; deepest -8.57% (XAGUSD)
Minimum equity against the 30% floor	\$999,777 — the floor at \$700,000 was never approached
Aggregate risk at maximum concurrency	5 concurrent ideas x 3% = 15% against a 30% limit

The risk budget governs planned exposure at entry, not unrealised excursion. Twenty-nine positions carried open heat beyond 3% of NAV while the scenario remained intact, and none beyond 10%; these are not breaches of the budget, which is measured at entry. The risk unit is the idea rather than the ticket: a single macro expression built in tranches shares one allocation, which is what makes a 3% budget consistent with as many as seventeen simultaneous positions.

Positions are given room by design




Winning positions were held 8.3 days on average, losing positions 31.1 days, and the scatter shows how much adverse excursion positions were permitted before resolving. This asymmetry is a property of the approach rather than a lapse in discipline: a position is held while the scenario remains intact and is exited on invalidation or on accrued carry, not at a fixed price distance. The corollary is that the meaningful constraint belongs at the portfolio level, which is where the 30% limit sits.

Worked example: the yen expression of July 2026

The single macro view is expressed across several instruments simultaneously. In July 2026 the programme held staged short positions in GBPJPY, EURJPY and USDJPY at uniform size, entered over several weeks as the scenario developed and closed in a coordinated exit — five GBPJPY tranches were closed within the same minute on 30 July for a combined realised \$79,509. This is one position expressed in three instruments, not three independent trades, and the concentration figures that follow should be read in that light.

Attribution and Costs

Where the return came from, and what it cost to carry



Precious metals (XAU, XAG) **\$824,339 · 38% · 22 trades**

JPY crosses **\$725,329 · 34% · 46 trades**

Other FX crosses **\$559,123 · 26% · 127 trades**

Crude and index **\$31,593 · 2% · 12 trades**

Top 5 trades / gross wins **25.4%**

Largest single trade **+\$214,310 · XAGUSD · 30 Jan 2026**

Seventy-two per cent of net profit came from two macro expressions — precious metals and the yen complex — each executed across several instruments. This is a concentrated discretionary macro programme by design, not a diversified systematic strategy, and should be evaluated as such. Overlapping positions sharing a theme group into 47 independent clusters.

Cost structure

Item	Amount	% of gross winning P&L
Financing / swap	-\$282,633	11.7%
of which paid on losing positions	-\$135,509	5.6%
Commissions	-\$41,553	1.7%
Total frictions	-\$324,186	13.4%

Percentages are expressed against gross winning P&L of \$2,406,932, the same denominator used in the profit factor. Largest drags: GBPJPY -\$103,488 · GBPCAD -\$66,742 · EURJPY -\$41,096 · EURAUD -\$29,549. Financing scales linearly with assets and is the principal cost item at a 10.9-day average hold; institutional prime-of-prime terms would be expected to improve it materially versus retail swap rates.

Important disclosures. Past performance is not indicative of future results and involves the risk of loss of capital. The record shown is a single discretionary managed account funded with \$1,000,000 in July 2025 and traded for one client under trading authority. It is not the performance of a pooled vehicle, it has not been audited, and it is published with the account owner's consent. Returns are gross of fees unless expressly stated as pro-forma net. The record covers 13 months, 207 closed trades of which 23 were losses, and 47 independent position clusters — a limited sample from which to infer the distribution of losses. Every month in the sample was positive, which means the strategy's behaviour in a losing period has not been observed. Stops are maintained by a proprietary monitoring layer external to the broker rather than as resting orders, which introduces an operational dependency on that layer. Leverage is employed at a median of 9.2 times equity, with a recorded peak of 51.1 times, and materially amplifies both gains and losses. The exposure profile in the record was achieved at a venue not subject to retail leverage caps and is not reproducible under standard retail classification in the European Union. Results were achieved in market conditions prevailing between July 2025 and July 2026 and may not be repeatable. The bootstrap analysis is an illustrative reconstruction from the closed-trade record and is not a statement of what would have occurred. Published for information only. Not an offer to sell or a solicitation to buy any security or interest in any fund, and not investment, legal or tax advice. Any offering would be made solely pursuant to definitive documentation.