

Hedge-only drawdown budgets and return trade-offs

A specialised mandate with 10%, 15% and 20% design alternatives

Executive assessment. The proposed controlled-drawdown solution will concentrate on paired long/short hedge bots, supported by a cash reserve. Standalone directional grids and discretionary swing positions sit outside this mandate. This narrower scope removes a material source of legacy directional inventory and makes the risk and return of the selected systems easier to measure. It does not, by itself, establish a drawdown ceiling.

The earlier assessment considered containing a broader portfolio of hedge, grid and swing strategies. The revised proposal separates that broader business from a dedicated allocator mandate: returns would be sought through paired rebalancing and a deliberately bounded directional bias, with account-level controls governing gross exposure, net exposure, costs and margin.

Design budget	Hedge-only portfolio implication	Expected trade-off
10%	Smallest stress-sized gross and residual net exposure; greatest reserve within a common strategy.	Lower participation and less tolerance for adverse trends or execution delays.
15%	Intermediate sizing; proposed midpoint for a controlled comparison.	More participation than the 10% design, subject to a larger loss allowance.
20%	Highest stress-sized exposure of these designs, still using paired systems.	Greater return potential and greater exposure to losses; no grid or swing sleeve is added.

The investment case and its present limit

We favour specialisation because the earlier account reviews separated relatively balanced hedge positions from much larger unpaired directional exposures. Removing those exposures changes the return engine as well as the risk. The mixed-book return history therefore cannot establish what a hedge-only portfolio would earn at any of the three budgets.

The research sequence is to establish the selected hedge systems' net economics, test a controlled ratio/slope adjustment at unchanged sizing, and consider additional gross exposure only after the evidence supports it. A 15% design is a research midpoint, not an empirically optimal allocation. No annual return or guaranteed drawdown cap is established by this paper.

This revision changes the proposed mandate, not any live account. Evidence is dated through 1 October 2026; no fresh performance collection is implied. Drawdown means a decline from a cash-flow-adjusted portfolio high-water mark, including open PnL. Client examples are anonymised.

The options considered and why the scope changed

The original question was how to contain drawdown across a mixed strategy. Three approaches were considered; each remains relevant to portfolio research, although the revised mandate applies its controls only to the selected paired hedge systems.

Earlier option	Return and risk mechanism	Reason for the narrower proposal
Scale the whole mixed book	Reduce hedge, grid and swing exposure proportionally; retain their existing return drivers.	Lower sizing still retains unpaired grid and swing risk. It also reduces potentially efficient hedge activity.
Keep the mixed book and de-risk after losses	Intervene at portfolio thresholds, then re-enter under predefined rules.	Exit timing and re-entry can dominate outcomes; inventory may accumulate before the trigger. Execution can overshoot it.
Selectively redesign the mixed inventory	Remove inefficient or recovery-dependent positions while retaining selected hedge, grid and swing strategies.	Potentially better capital use, but disparate return sources make a single controlled-drawdown proposition harder to attribute and govern.

Evidence behind the specialisation

In the 19 September detailed review, 22 hedge pairs held approximately \$5,228 net long, while 11 long grids held \$32,607. Grids supplied approximately **86.2% of net direction**. Their combined open PnL was about -\$6,753, compared with +\$90 for the hedge pairs. Those are marked position observations, not lifetime profitability comparisons. [S5]

Across 70 connected accounts on 26 September, grids contributed approximately \$0.566 million of \$0.929 million net long exposure, or 60.9%. This heterogeneous aggregation is neither a performance composite nor a common collateral pool; its hedge-labelled category also includes intentional swing systems. [S8]

A static illustration of the scope change

At the detailed account's recorded \$33,921.90 equity, a synchronous 20% fall with unchanged quantities implied a 22.31% equity loss for its full net position. Retaining only the recorded hedge pairs at the same equity, with no redeployment of released capital, gives approximately **3.08%** under that single shock. The paired positions still carried roughly \$208,720 gross notional. [S5]

The 3.08% result is a counterfactual exposure calculation, **not a hedge-only backtest, maximum drawdown estimate or suggested leverage limit**. It excludes rebalancing, divergent prices, costs and liquidation mechanics. Removing the grids sacrifices their potential future gains as well as their risk; redeploying the reserve into larger hedges would change the calculation.

The specialised mandate excludes standalone directional grids and intentional one-sided swing systems even where the platform labels them as hedge bots. Internal price grids used to rebalance a paired hedge bot remain part of the strategy.

Ratio, slope and leverage are different controls

Paired positions can limit immediate net direction while rebalancing changes their exposure as prices move. Economic return depends on the path, the inventory retained, execution costs and funding. Completed profitable cycles can coexist with open losses; the mandate must assess both legs together rather than count rebalances as income.

Slant: the long/short balance at a particular price

For combined position-margin budget M , equal actual leverage L on linear perpetual legs and long notional share w : **gross notional** $G = M \times L$; **net long** $N = G \times (2w - 1)$. The examples below hold M at \$1,000. They are static sizing identities before fees and rebalancing. [S11]

Current ratio	Actual sized leverage	Gross notional	Net long	PnL at a fixed -10% move
55/45	5x	\$5,000	\$500	-\$50
60/40	5x	\$5,000	\$1,000	-\$100
65/35	5x	\$5,000	\$1,500	-\$150
55/45	10x	\$10,000	\$1,000	-\$100

Moving from 55/45 to 60/40 doubles the net bias at unchanged gross exposure; moving to 65/35 triples it. The additional upside participation has a corresponding downside exposure. Increasing sized leverage from 5x to 10x doubles gross as well as net exposure at the same margin budget. Changing a leverage setting while holding quantities fixed does not increase dollar PnL.

Slope: how the target changes through the range

Boundary ratios determine the permitted exposure schedule. The price-range width determines how rapidly that schedule changes with price; internal grid spacing, cadence, tolerance and order minimums determine whether and when trades occur. A steeper schedule can create larger inventory adjustments, but does not establish profitable turnover. With a long-heavy lower boundary, further declines can increase the target long bias.

Under Cornix guidance supplied to management, Range mode derives its starting target from the grid nearest the market and the boundary ratios; it has no independent current-ratio override that preserves both boundaries. Fixed mode uses a selected ratio without that range-driven schedule. Editing active bots can queue orders. These are product mechanics, not a new execution instruction.

Research preference

Test a modest ratio/slope change first while holding actual leverage, allocation, symbol, range, spacing and cadence constant, with an unchanged comparison cohort. Consider larger gross sizing only after repeatable after-cost economics and stress margin are demonstrated. This sequence follows the earlier exposure study; it is not proof that a larger slant produces a better risk-adjusted return. [S4, S11]

Eligibility depends on liquidity, attributable net results and stress behaviour. Major crypto and equity perpetuals may be compared with longer-tail assets; neither a familiar ticker nor a hedge label is evidence of low risk. Intentional one-sided swing configurations remain excluded.

Sizing the specialised mandate

A 10%, 15% or 20% budget is a proposed account-level risk objective. Its return consequence must be measured on a common, cash-flow-adjusted hedge-only record. Legacy mixed-account results cannot supply that calibration, and the strategy selection may change return per unit of risk in ways that simple proportional scaling cannot capture.

Budget	Conditional exposure / return contribution*	Gain required after a loss of this size
10%	50% of an assumed 20%-budget reference	11.11%
15%	75% of that same reference	17.65%
20%	100% reference	25.00%

*First-order sizing illustration only: $k = \text{budget} / \text{assumed reference drawdown}$. This assumes the same positive-edge hedge strategy, unchanged cost efficiency and approximately proportional exposure, with zero reserve yield and no fixed or investor fees. Compounding, lot sizes and different loss episodes break exact scaling. It is not a forecast; no numerical annual return is assigned.

Size against the exposure that can develop

A near-50/50 position today may move towards its long-heavy lower boundary during a decline. Stress tests must include that schedule, outstanding orders and potential rebalances, with account-specific leverage and maintenance-margin tiers. Equal percentage bot allocations do not guarantee equal notional or equal risk.

Same gross / equity: 2.0x	Net long / equity	Loss under one -20% price shock
Current or assumed ratio 55/45	0.20x	4% of equity
Current or assumed ratio 60/40	0.40x	8% of equity
Current or assumed ratio 65/35	0.60x	12% of equity

These static examples hold quantities fixed and assume identical percentage price changes. They exclude dynamic trading, costs and liquidation. They demonstrate why a larger ratio bias consumes the drawdown allowance even without more gross exposure; 2.0x is an illustration, not a recommended portfolio limit.

Execution, intervention and re-entry

Use separate limits for gross exposure, residual net exposure, symbol concentration and venue margin. Intervention should occur before the outer drawdown budget is exhausted, with enough allowance for executable depth, gaps and a failed or delayed leg. The earlier 8% / 12% / 16% trigger examples remain unvalidated and are not proposed operating thresholds.

Reducing exposure must also prevent automatic rebuilding by the bots and verify residual holdings. Closing a position does not reset the high-water mark: at 10% current drawdown, a 15% outer budget leaves only **5.56% of current equity** before costs. Available margin is a different measure. A validated portfolio-wide intervention mechanism is not established by this review.

Re-entry criteria and staged sizing must be specified before testing. Stops can crystallise temporary losses and miss a rebound; they can also prevent a deeper decline. A trigger does not guarantee an execution price. [P2]

What the existing record does and does not establish

Evidence	Recorded result	Interpretation
Ordinary matched hedge cohort [S4]	1,163 bots over 54.29 hours, 26-28 September: realised PnL increased ~\$3,631; combined reported PnL changed by ~\$3,763.	Profitable closed cycles can coexist with marked losses. This dated interval is not a current profitability verdict.
Systems becoming inactive [S4]	134 bots separately retained, with a reported PnL change of ~\$1,519.	Closed and discontinued systems must remain in performance measurement.
Reference hedge accounts [S9]	30 September: ordinary Binance PnL ~\$237; OKX ~+\$231. A separate swing position was excluded.	Symbol, leverage and allocation differences prevent a causal ratio or leverage comparison.
Longer-tail example [S10]	1000RATS: ~55 days and 428 rebalances; combined reported PnL ~\$50.54 on 30 September.	Activity and elapsed time do not alone establish a positive economic edge.

Combined reported PnL already includes open and closed trade changes; the realised increase must not be added again. Commissions and funding have not been fully reconciled to venue ledgers. Recent management reports of positive hedge-bot income are encouraging, but do not replace a continuous marked, after-cost record for the selected mandate. [P3]

Legacy equity histories remain background evidence

The earlier report displayed five account examples with seven monthly marks. Their largest sampled equity falls ranged from 29.59% to 45.34%, and August-October changes ranged from +40.89% to +75.80%. These were unadjusted account-equity observations, with unresolved transfers and mixed strategies. They are **not the drawdown or return record of the hedge-only proposal**. [S1-S3]

Of 2,161 requested quarterly/monthly observations, 1,445 had exact or nearest evidence, 332 preceded connection and 384 were unavailable. The five examples' marks ranged from five hours before to 31 hours after nominal boundaries; monthly sampling also misses intramonth extremes. A separate reconstruction identified a manager-confirmed \$7,604.60 withdrawal within an apparent fall, without independent ledger confirmation. [S1-S3, S7]

Accounting boundary

Use a dedicated account or independently attributable sleeve. Reconcile both legs, shared-symbol positions and closed bots to exchange fills, fees and funding. Cash movements are not profit; open PnL is already part of equity. Shared collateral with excluded grids or swings would still expose the mandate to their losses, so operational separation must match the stated investment scope.

A promising live configuration and a demonstrated controlled-drawdown product are different stages of evidence. The specialisation is the portfolio-design decision; continuous performance measurement and executable risk controls are the work required to substantiate it.

Validation, reporting and progression

Proposed next stage. Establish a separately attributable hedge-only reference portfolio and compare the 10%, 15% and 20% designs on the same accounting and execution basis. Treat 15% as the initial research midpoint. Progression depends on the measured results, not on filling an allocation target or raising leverage to meet a promised return.

Three stages of validation

- 1. Establish the reference.** Identify exact ordinary paired-bot cohorts, capital history and configuration dates, retaining discontinued systems. Reconstruct cash-flow-adjusted NAV including open PnL, fees and funding. Daily returns support performance measurement; intraday fills, prices and margin observations are needed for stress and execution testing. [P1, P3]
- 2. Isolate the return controls.** Compare unchanged controls with a modest slope/ratio variant at common actual leverage and allocation. Keep symbol, range, spacing and cadence fixed. Evaluate completed cycles, marked net PnL, turnover costs, adverse trends and minimum order constraints before considering a separate leverage experiment.
- 3. Test each risk budget.** Replay permanent exposure scaling and predefined de-risking/re-entry policies on common periods, including gaps, depth, partial fills, funding and maintenance margin. Evaluate outside the calibration sample and retain correlated market episodes. Confirm bot-level execution and prevention of unwanted re-entry before any controlled implementation.

Allocator reporting

Disclose trading return before investor fees and net-to-investor return under the agreed fee schedule, alongside maximum drawdown, time underwater, recovery time, gross and net exposure, concentration, fees/funding, stress margin and any breaches. Any annualised measure must identify the underlying period. Historical mixed-book results will remain separate from the specialised mandate record.

Source notes

S1-S3. Internal account-equity reconciliation, April-October 2026, compiled 1 October; transfer ledger incomplete.

S4, S11. Matched hedge-system study and slant/leverage analysis, 26-28 September 2026, covering 54.29 hours.

S5. Detailed portfolio review, 19 September 2026, 08:41 BST; sequential position snapshots.

S7. Anonymised cash-flow reconstruction, 1 October; specified withdrawal manager-confirmed only.

S8. Cross-account exposure review, 26 September 2026.

S9-S10. Reference-account and 1000RATS reviews, 30 September. All are dated evidence, not refreshed holdings.

Product guidance. Cornix support response supplied by management on Fixed/Range ratios and active-bot edits; deployment availability not rechecked in this revision.

P1. CFA Institute, [GIPS Standards Handbook for Firms, valuation and time-weighted returns](#). Accessed 3 October 2026.

P2. SEC Investor.gov, [Stop, Stop-Limit, and Trailing Stop Orders](#). Accessed 3 October 2026.

P3. Cornix, [Trade PnL on the Exchange vs. Trade PnL on Cornix](#). Accessed 3 October 2026.

Underlying account mappings, source observations and calculations are retained in the internal research record. Client identities are omitted from this discussion paper. The analysis has not been independently verified; no GIPS compliance claim is made.