Vasilis Kotopoulos

AI-native quant research · From strategy abundance to allocation conviction

NOVEMBER 2026

MANDATE ASSESSMENT · SAMPLE DELIVERABLE

A low-drawdown sleeve built for first-loss leverage

> First-loss capital allocator · Leveraged managed-account programme

VERDICT

Pass

Clears the drawdown screen and stays inside the hard stop at 2× leverage. Return per unit of risk above cash is below a 1.5 bar some allocators use.

PREPARED BY

Vasilis Kotopoulos

Case study against a publicly described mandate, not a past or current client. All results are simulated. Past and simulated performance is not indicative of future results. Not investment advice.

EXECUTIVE SUMMARY

Executive summary

VERDICT

Pass

Clears the drawdown screen and stays inside the hard stop at 2× leverage. Return per unit of risk above cash is below a 1.5 bar some allocators use.

THE MANDATE

A first-loss allocator matches a manager's own risk capital and levers the combined account, up to around 10×, with a hard stop at about −8.5% from starting NAV. It screens for drawdowns of 3–5% through the COVID crash, an attractive Sharpe ratio, money made every year, liquid strategies and a 2–3 year record.

−3.1%
worst fall, Mar 2007 – Sep 2025, incl. 2008 and 2020

1.56
Sharpe ratio (1.23 above cash)

19 of 19
calendar years ahead of cash

6.9%
a year unlevered, after costs

Key findings

1. **Inside the drawdown screen.** The worst fall over March 2007 – September 2025 was −3.1%, a period that includes the 2008 crisis and the COVID crash. Q1 2020 returned +1.2%.
2. **Positive against cash every year.** The sleeve beat cash in all 19 calendar years of the design period, with near-zero sensitivity to the stock market (beta −0.02).
3. **The hard stop holds at 2× leverage.** At 2× the worst fall from a peak is −6.0%, and the account never ended more than −4.0% below its value at the start of any year, so the −8.5% stop is never reached.
4. **Consistent in the held-back test year.** Sharpe 1.59 in October 2025 – October 2026, and 1.51 with double costs and late trades.
5. **Main caveats.** Measured above cash the Sharpe ratio is 1.23, below a 1.5 bar; a random-reshuffle stress points to a possible −6.1% fall unlevered; and the record is simulated, so it does not satisfy the 2–3 year live-record requirement.

RECOMMENDED NEXT STEP

A **fixed-fee, one-month paid trial**: we run the sleeve on paper every day against your exact risk rules (including the hard stop at your chosen leverage), report monthly in your format, and hand over full risk and attribution data. The trial also starts the live record your programme requires. No commitment beyond the trial.

SECTION 1

Mandate and constraints

The requirements below are taken from the allocator's published programme description. Where the programme gave a range or a qualitative aim, we fixed a specific number before testing.

Requirement	What the programme asks for	How we tested it
Structure	Manager capital is matched and levered, up to around 10×	Results shown unlevered and at 2×
Hard stop	About −8.5% from starting NAV	Worst fall at 2×, and worst level below each year's starting value
Drawdown screen	Worst fall of 3–5% through the COVID crash	Worst fall below 5% over 2007 – Sep 2025, including 2008 and 2020
Return quality	Attractive Sharpe ratio	Sharpe ratio, with a 1.5 reference bar; Sharpe above cash also reported
Consistency	Make money every year	Every calendar year ahead of cash
Liquidity	Liquid strategies	Liquid exchange-traded bond, gold, currency and commodity markets
Track record	2–3 years	Not met by a simulation; addressed by the proposed live trial
Return target	10–15%+ a year on the levered account	Unlevered return reported; leverage scales it

> **Pass rule, fixed in advance:** worst fall below 5% including 2020, a reported bad-case reshuffle drawdown, a strong Sharpe ratio, and the hard stop checked at 2× leverage.

SECTION 2

Approach

What we built

The same diversified risk premia sleeve we use for low-equity-risk mandates, run at a deliberately low level of risk:

- **Return sources:** a seasonal commodity strategy at the core, a defensive shelter switch (bonds and gold when stocks are weak, cash otherwise) and a small trend-following slice across alternative markets.
- **Volatility steering:** exposure is adjusted day to day so the sleeve runs at a very low, steady level of risk.
- **An automatic drawdown brake:** exposure is cut progressively as losses from the last high build up, and restored as the sleeve recovers. This is designed for a programme with a hard stop.

All parameters are fixed in advance. There is no machine learning and no discretionary input.

How we tested it

- **Rules locked in advance.** The candidate designs, the selection objective and the pass/fail rules were written down and time-stamped before any test was run.
- **Strict separation of design and test data.** All choices were made on March 2007 – September 2025 data, choosing among a small, pre-declared set of 19 versions using the mandate's own objective. October 2025 – October 2026 was held back and used once, only to check the final rules.
- **Realistic execution.** Signals are taken at the close and traded the next day; trading costs are charged on every position change; uninvested cash earns the T-bill rate.
- **Cost and timing stress.** Every result was re-run with double trading costs and with trades one day later, separately and together.
- **Luck adjustment.** Statistical confidence is adjusted for the roughly 2,000 ideas tested across our research to date, not just the versions tested for this mandate.
- **Statistical stress tests.** A block-bootstrap range for the Sharpe ratio, a random-reshuffle drawdown stress and a random sign-flip test on daily returns.
- **Benchmarks.** Results are compared with US stocks (S&P 500), cash, the mandate's own benchmark and a plain trend-following portfolio, to check the sleeve is not a repackaged version of something simpler.

SECTION 3

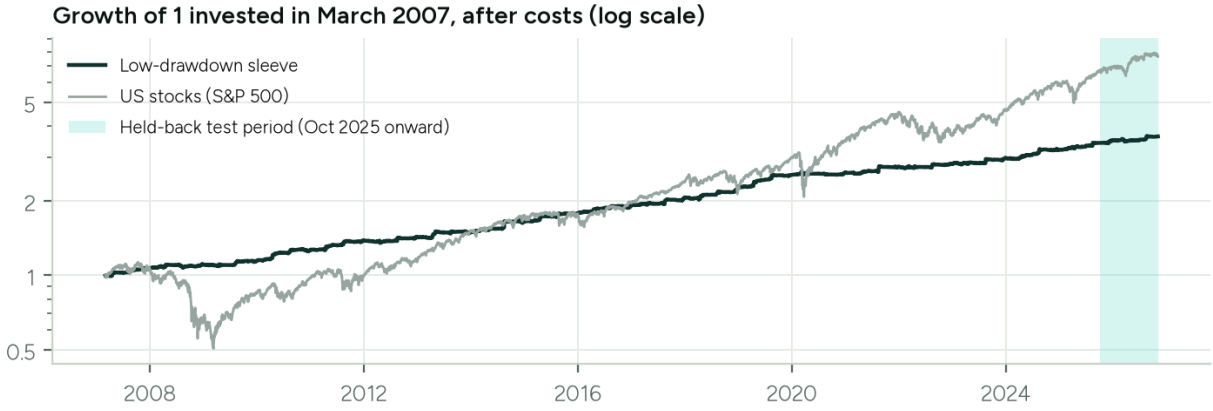
Results

Mandate scorecard

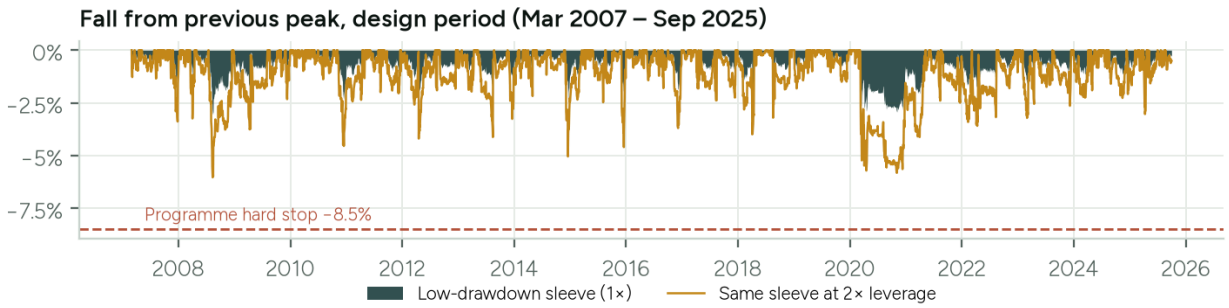
Requirement	Target	Result	Status
Worst fall incl. 2008 and 2020	3–5%	-3.1% (2007 – Sep 2025)	Met
Q1 2020	withstand	+1.2% (S&P 500 -19.4%)	Met
Calendar years ahead of cash	every year	19 of 19	Met
Sharpe ratio	attractive (1.5 reference)	1.56; above cash 1.23	Partly
Held-back test year	holds up	Sharpe 1.59; 1.51 with double costs and late trades	Met
Hard stop at 2× leverage	never hit	Worst fall -6.0%; never below -4.0% vs a year's start	Met
2–3 year live record	required	Simulation only	Not met

Design-period profile (March 2007 – September 2025)

Measure	Base case	Double costs + late trades
Annual return (unlevered)	6.9%	6.4%
Volatility	4.3%	4.2%
Sharpe ratio	1.56	1.47
Sharpe above cash	1.23	1.14
Worst fall	-3.1%	-3.1%
Calendar years ahead of cash	19 of 19	17 of 19
Return-to-worst-fall ratio	2.24	2.06
Months with a positive return	66%	—



Growth of 1 invested in March 2007, unlevered and after costs. The shaded area is the held-back test period.

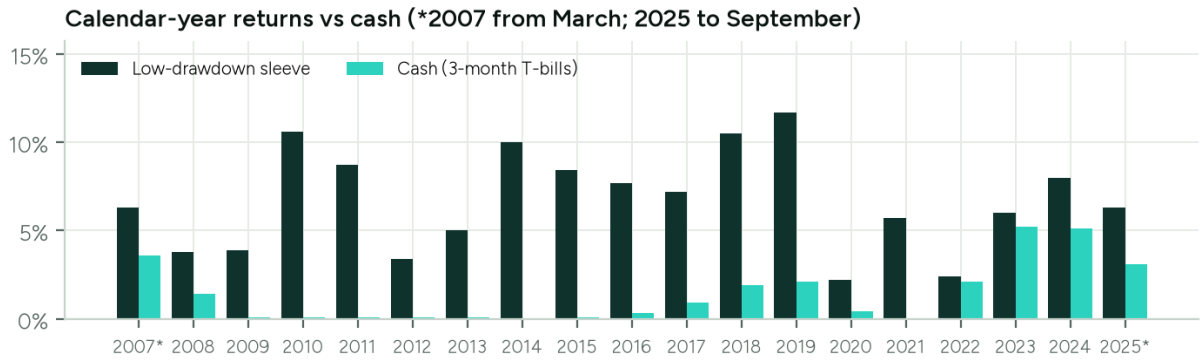


Fall from previous peak, unlevered (shaded) and at 2× leverage (line). Even at 2×, the deepest fall (−6.0%, August 2008) stayed well clear of the −8.5% stop.

Behaviour in stress and under leverage

Measure	Unlevered	At 2× leverage
Worst fall, 2007 – Sep 2025	−3.1%	−6.0%
Worst level below a year's starting value	—	−4.0%
2008 (calendar year)	+3.8%	—
Q1 2020	+1.2%	—
2022	+2.4%	—
Bad-case reshuffle fall (1-in-20)	−6.1%	—

2× figures scale daily returns by two and do not deduct any financing cost charged by the programme.



Calendar-year returns against cash, unlevered.

Held-back test year (October 2025 – October 2026)

Measure	Base case	Double costs + late trades
Sharpe ratio	1.59	1.51

Reported exactly as it came out; nothing was changed after seeing it. Only the figures above are reported for this period.

SECTION 4

Robustness and risk assessment

How much confidence the evidence supports

- **On its own terms the record is solid.** The block-bootstrap range for the design-period Sharpe ratio is 1.33 to 1.84, and a random sign-flip test was never matched in 1,000 trials.
- **After allowing for everything we have tested, it is not proven.** Adjusting for the roughly 2,000 ideas screened across our research, luck-adjusted confidence is about 17%, short of the 95% we would require. The pass rests on mandate fit and consistency, which is why we recommend a live paper trial before capital is committed.
- **Cost and timing.** Double costs and late trades cut the Sharpe ratio only modestly (1.56 to 1.47) and leave the worst fall unchanged, but two calendar years then fall slightly behind cash (17 of 19).

Key risks

- **Leverage and the hard stop.** The historical worst fall at 2× (−6.0%) leaves room under the −8.5% stop. But the random-reshuffle stress shows an unlucky sequence could produce about −6.1% unlevered, which at 2× would be roughly double and beyond the stop. Leverage should be set with that bad case in mind, not the historical one.
- **Modest unlevered return.** About 6.9% a year unlevered, before programme leverage. Returns above cash are thinner: the Sharpe ratio above cash is 1.23.
- **Concentration.** The seasonal commodity strategy is the largest single driver (monthly correlation 0.75), so diversification inside the sleeve is partial.
- **Track record.** The programme asks for 2–3 years of record; this is a simulation. A live paper record is needed.
- **Simulation limits.** Financing costs for leverage, market impact and operational frictions are modelled simply or not at all.

> **Bottom line.** The sleeve is built for exactly this structure: small, controlled drawdowns that leave room for leverage. The open questions are how much leverage the bad case allows and whether a live record confirms the simulation.

SECTION 5

Verdict and rationale

Pass Clears the drawdown screen and stays inside the hard stop at 2× leverage. Return per unit of risk above cash is below a 1.5 bar some allocators use.

Why it passes

- The worst fall of −3.1% over 2007 – Sep 2025, including 2008 and the COVID crash, sits inside the programme's 3–5% screen.
- It beat cash in every calendar year of the design period and was positive in Q1 2020.
- At 2× leverage it never comes near the −8.5% hard stop on the historical record.
- The held-back test year confirmed the risk-adjusted profile (Sharpe 1.59; 1.51 under stress).

What keeps it from being a stronger pass

- Sharpe above cash (1.23) is below a 1.5 bar some allocators apply.
- Luck-adjusted confidence is low (about 17%), and there is no live record yet.
- The bad-case reshuffle stress implies leverage should be set conservatively.

SECTION 6

Recommended next steps

1. **One-month paid trial (fixed fee).** We run the sleeve on paper every day against your exact policy rules, including the hard stop at your chosen leverage, send a monthly report in your format, and hand over full risk and attribution data. No commitment beyond the trial.

2. **Agree the leverage level.** Set leverage using the bad-case drawdown rather than the historical one, and confirm how the programme's financing costs apply.

3. **Build the record.** Continue the paper or small live record toward the programme's 2–3 year requirement, with monthly reporting against the same rules.

APPENDIX

Definitions in plain English

Term	What it means
Annual return	The average yearly growth rate over the period, compounded, after modelled trading costs.
Volatility	How much returns swing around, expressed per year. Higher means a bumpier ride.
Sharpe ratio (return per unit of risk)	Average return divided by volatility, per year. Above 1 is generally considered good for a single strategy. Unless stated otherwise, we show it on total return.
Sharpe above cash	The same ratio after first subtracting the return available on cash (3-month T-bills). It shows the reward for taking risk rather than holding cash, and is the stricter of the two.
Worst fall (maximum drawdown)	The largest drop from a previous high point to a subsequent low, before a new high is reached.
Leverage (2×)	Running the account at twice the sleeve's exposure. Gains and losses are roughly doubled.
Hard stop	The loss from starting account value at which the programme closes the account.
Stock-market sensitivity (beta)	How much the strategy tends to move when the stock market moves 1%. Zero means no systematic link; 1 means it moves with the market.
Design period	The historical window in which all choices were made. Nothing after 30 September 2025 was used to design or select anything.
Held-back test period	Data from 1 October 2025 onward, kept aside and used only once, to check the final, already-fixed rules. It is a simulation on unseen data, not live trading.
Double costs and late trades	A stress test that doubles every modelled trading cost and delays each trade by one extra day. A strategy that only works with perfect timing or very cheap trading usually fails this test.
Luck-adjusted confidence	Testing many ideas makes it likely that some look good by chance. This statistic (a "deflated Sharpe ratio") estimates the probability that the record reflects real skill, after allowing for the total number of ideas we have tested. We require 95% to call a result statistically proven.
Bootstrap range	We resample blocks of the historical returns 1,000 times and recompute the Sharpe ratio. The range shown covers the middle 90% of outcomes.
Random-reshuffle stress	We shuffle the order of historical daily returns 1,000 times and record the worst fall each time. The figure shown is a bad-case (1-in-20) outcome, a guide to how deep a fall could be in an unlucky sequence.
Sign-flip test	We randomly flip the sign of daily returns 1,000 times. If the real average return is rarely matched, the result is unlikely to be noise.

Important information

This document is a sample deliverable. It was prepared against a mandate described in public sources; the organisation that published that mandate is not a past or current client and did not commission, review or endorse this work.

All performance figures are simulated (backtested) and hypothetical. They were produced by applying fixed rules to historical market data, after modelled trading costs, and do not reflect actual trading or real money. Simulated results have inherent limitations: they are prepared with the benefit of hindsight, cannot fully account for liquidity, market impact, financing costs, operational constraints or the effect of real trading decisions, and may differ materially from results actually achieved. The held-back test period is also simulated, not live trading.

Past and simulated performance is not indicative of future results. Any investment can lose value, and leverage magnifies losses as well as gains.

This document is provided for information and discussion only. It is not investment advice, a recommendation, or an offer or solicitation to buy or sell any security, fund or strategy, and it does not take into account the objectives, financial situation or needs of any investor. Recipients should make their own independent assessment and seek professional advice before making any investment decision.