

Dual-Alpha 10/90 Governed Portfolio

Live-realistic, true next-open backtest on IBKR-sourced historical data.

Roberto Cocco & Amin Hydar Ali

STARTING CAPITAL

\$9,000

INITIAL ALLOCATION

10% Momentum / 90% CTR0166

HISTORICAL CAGR

14.30%

HISTORICAL MAX DRAWDOWN

-7.12%

Publication standard

A research brochure explaining the strategy architecture, portfolio governance, execution model, IBKR data lineage, reproduced backtest evidence, stress testing, Monte Carlo results, realistic expectations and limitations. It is not a guarantee of future performance and is not an audited live track record.

The governed portfolio improves drawdown control while preserving attractive return potential

The Dual-Alpha 10/90 Governed Portfolio combines two independent equity strategies – GAP_FILTERED Momentum E496 and Contrarian CTR0166 – inside one \$9,000 account. The portfolio begins with 10% allocated to Momentum and 90% to the lower-volatility Contrarian sleeve. Momentum is allowed to drift naturally, but when it exceeds a 20% portfolio ceiling at one of its native rebalance dates, it is trimmed back toward 10% at the next actual market open.

FINAL EQUITY

\$17,494.71

CAGR

14.30%

ANNUALIZED VOLATILITY

9.41%

SHARPE RATIO

1.468

SORTINO RATIO

1.850

MAXIMUM DRAWDOWN

-7.12%

SPY BETA

0.228

ANNUALIZED ALPHA VS SPY

10.72%

Interpretation

The governed portfolio sacrifices some headline return versus allowing Momentum to drift without limits, but it materially improves drawdown control and risk-adjusted performance. The result is more consistent with the intended core/satellite design.

Why the governance matters

- Without a Momentum cap, the 10/90 portfolio historically drifted to a maximum Momentum weight of 38.23% and experienced a -13.22% maximum drawdown.
- With the 20% ceiling, the maximum Momentum weight was 22.29%, average Momentum exposure was 11.59%, and maximum drawdown improved to -7.12%.
- Only one governance breach occurred in the primary run, so the overlay did not continuously rebalance the portfolio or dominate the strategy logic.

2. Data provenance: what is "real" in this backtest?

The price and volume inputs are real historical daily OHLCV records originally downloaded through the Interactive Brokers API and stored locally in parquet caches. During this governed backtest, the program deliberately makes no fresh API request: it reads the previously downloaded IBKR-sourced caches so the historical test window cannot silently change from one run to the next.

Real data, simulated trading

The market history is real IBKR-sourced data. The portfolio P&L is simulated. Historical bid/ask queues, exact opening-auction queue position, routing venue, and true historical market impact are not observable from daily OHLCV, so those execution frictions are modeled and then stress-tested.

Momentum data lineage

- Original universe: 500 stocks.
- Reconstructed frozen panel: 500 stocks x 1,506 dates, 31 August 2020 through 28 August 2026.
- The original Momentum source code was independently reproduced with zero internal implementation difference.
- The saved August 2026 result cannot be reproduced byte-for-byte because the local historical parquet files were subsequently refreshed; the current reconstructed Momentum result remains close to the original research result.

Contrarian data lineage

- Original stock-like universe: 719 securities after removing likely funds/ETPs.
- Frozen panel: 719 stocks x 2,514 dates, 6 September 2016 through 4 September 2026.
- Walk-forward OOS and diagnostic holdout metrics reproduce to floating-point precision.
- Cached daily gross-return and turnover streams match with a maximum absolute difference of 0.000e+00.

Important historical-universe limitation

The frozen universes were discovered from a current IBKR candidate universe rather than reconstructed from a fully point-in-time membership database. That means the research still carries survivorship/selection bias. This is the largest remaining research limitation and must be considered when interpreting expected future returns.

3. The two alpha engines

A. GAP_FILTERED Momentum E496 - satellite sleeve

Momentum E496 ranks stocks cross-sectionally using long-horizon price momentum combined with a penalty for unstable overnight gaps. The frozen winning specification uses 42 / 252 / 252 lookback inputs, a 5-day skip, a 90-day volatility lookback, Top 10 holdings, and a 42-trading-day rebalance cadence.

- Role in portfolio: return-seeking satellite.
- Historical characteristic: high return potential but materially higher volatility and drawdown than CTR0166.
- Portfolio control: initially 10% of account equity and governed by a 20% maximum Momentum weight.
- No leverage and no short selling in this combined account simulation.

B. Contrarian CTR0166 - core sleeve

CTR0166 is a stabilization/reversal strategy. It seeks longer-term losers that show evidence of recovery and stabilization, while applying liquidity, price, volatility, trend, history, and market-regime filters. It is the lower-volatility core of the portfolio.

- Original development walk-forward OOS: CAGR 14.33%, Sharpe 1.753, MaxDD -6.68%.
- Frozen diagnostic holdout: CAGR 18.56%, Sharpe 1.889, MaxDD -4.38%.
- Its original daily return and turnover streams reproduce exactly from the frozen cache.
- Rebalance cadence: 63 trading sessions with next-open execution.

Diversification logic

The design intentionally combines a higher-volatility trend/momentum source with a lower-volatility stabilization/reversal source. The objective is not to maximize either sleeve independently, but to create a portfolio whose return drivers are different enough to improve risk-adjusted performance.

4. How the code works - event-driven account simulation

The backtest is built as an account ledger rather than a weighted average of two return series. It maintains strategy-level cash, actual share quantities, account equity, order ledgers, commissions, modeled execution friction, and portfolio-governance transfers.

1. CLOSE T	2. SIZE	3. OPEN T+1	4. CLOSE T+1
Calculate frozen strategy signal using information available at the close.	Determine target shares from sleeve NAV and close-T information only.	Execute at the actual historical opening price plus modeled execution friction.	Mark cash and positions to the actual historical close; update account NAV.

Look-ahead control

The quantity decision is frozen using close-T information. The next day opening price is used only as the execution price. If an overnight gap makes the pre-planned buys unaffordable, the simulator scales buy orders down instead of borrowing cash or using the future opening price to recalculate the original target.

Portfolio accounting

- One starting account with \$9,000.
- Separate strategy sleeves for attribution, but one total portfolio equity value.
- No daily rebalancing between sleeves.
- Natural sleeve drift between native strategy rebalance dates.
- No leverage; gap-induced cash shortages reduce buys rather than create margin borrowing.
- Fractional-share primary simulation, plus an integer-share stress test.


Execution friction

The simulator charges commissions and models spread/slippage/market impact using order size, liquidity, and volatility information available at signal time. Because daily OHLCV cannot recreate the exact historical opening auction book, this is an approximation. Robustness is therefore judged not only at baseline assumptions but under 2x costs, 3x costs, an additional 5 bps per fill, and integer-share execution.


5. Portfolio governance: 10/90 with a Momentum risk ceiling

The portfolio starts at 10% Momentum and 90% CTR0166. The weights are then allowed to drift. Momentum is inspected only on its own native rebalance dates. If the Momentum sleeve exceeds 20% of total portfolio equity, it is reset toward 10% and the trim is executed at the next actual market open.

- The rule is risk governance, not a new alpha signal.
- CTR0166 is never forced to trade off-cycle just because Momentum is trimmed.
- Proceeds transferred from Momentum to CTR0166 remain as CTR cash until CTR0166 reaches its own normal rebalance.
- The rule is deliberately infrequent to preserve the original strategies and avoid artificial daily portfolio timing.



The 15%-25% ceiling region materially improves drawdown versus no cap. The 20% primary policy sits inside that stable region and triggered only one trim in the historical run.



Cap	CAGR	Sharpe	MaxDD	Max MOM	Trims
OFF	16.25%	1.293	-13.22%	38.23%	0
15%	13.80%	1.475	-6.80%	22.46%	2
20%	14.30%	1.468	-7.12%	22.29%	1
25%	14.97%	1.489	-8.71%	32.60%	1
30%	16.89%	1.464	-13.18%	35.25%	1

6. Primary live-realistic historical result

Test period	2021-09-01 to 2026-08-28
Starting capital	\$9,000.00
Final equity	\$17,494.71
CAGR	14.30%
Volatility	9.41%
Sharpe	1.468
Sortino	1.850
Maximum drawdown	-7.12%
Calmar	2.010
SPY correlation	0.417
SPY beta	0.228
Annualized alpha vs SPY	10.72%
Annual turnover	2.86x
Modeled friction	\$369.34
Orders	859

Trade-level and path behavior


- Closed position cycles: 367.
- Trade win rate: 56.1%.
- Average winning trade: +25.30%.
- Average losing trade: -10.03%.
- Longest underwater period: 570 trading sessions.
- Longest period of underperformance versus SPY: 860 trading sessions.
- Ending Momentum weight: 17.11%; average Momentum weight: 11.59%; maximum Momentum weight: 22.29%.
- The 20% governance rule was breached once; the historical transfer from Momentum to CTR0166 was \$1,255.26.

Why the long underwater periods matter

A strong CAGR and Sharpe do not mean smooth performance. The historical portfolio spent long stretches below a prior high and behind SPY. A live investor must be prepared for multi-year periods in which the strategy feels disappointing even if the long-run backtest remains intact.

7. Execution stress tests

A realistic strategy should remain economically viable when the assumed execution environment becomes worse. The governed portfolio remains profitable under all tested execution stresses.





Scenario	CAGR	Sharpe	MaxDD	Modeled friction
Baseline fractional	14.30%	1.468	-7.12%	\$369.34
2x costs	13.05%	1.366	-6.89%	\$733.40
3x costs	12.58%	1.319	-7.79%	\$1,090.20
+5 bps each fill	14.12%	1.451	-7.11%	\$454.67
Integer shares	10.09%	1.480	-6.67%	\$233.63

Interpretation

- Tripling modeled transaction costs still leaves a 12.58% CAGR and 1.319 Sharpe in the historical simulation.
- The integer-share stress falls to 10.09% CAGR, showing that small-account sizing matters, but the portfolio remains historically profitable and risk-adjusted performance remains strong.
- The results are not dependent on an unrealistically perfect zero-cost execution model.

8. Monte Carlo / block-bootstrap robustness

The program resamples 21-trading-day blocks from the historical daily return stream across 2,000 paths. Block resampling is used instead of independently shuffling individual days so that short-term volatility clustering and serial dependence are partially preserved.



Metric	p01	p05	p25	Median	p75	p95	p99
CAGR	3.85%	6.74%	11.08%	14.46%	18.31%	24.01%	28.96%
Sharpe	0.513	0.778	1.199	1.494	1.808	2.235	2.564
MaxDD	-14.71%	-12.32%	-9.45%	-7.98%	-6.91%	-5.85%	-4.74%
Final equity	\$10,859	\$12,448	\$15,175	\$17,613	\$20,766	\$26,242	\$31,874

Private quantitative research | Historical simulation - not audited live performance

Do not overinterpret the 100% positive-CAGR bootstrap rate

The bootstrap only rearranges blocks drawn from the same historical sample. It does not invent future crises, delistings, structural changes, strategy crowding, regulatory changes, data errors, or a permanent alpha decay. It is useful as a robustness test, not as a probability forecast for the real future.

9. Realistic expectations: what should an investor actually expect?

The historical 14.30% CAGR is the central backtest outcome, not a promise. A prudent deployment plan should assume that live returns may be materially lower because of survivorship/selection bias, market-regime change, execution differences, taxes, operational errors, and alpha decay.

Evidence-based planning ranges

Evidence	Historical result	How to use it
Baseline live-realistic backtest	14.30% CAGR	Central historical reference, not a forecast.
2x / 3x cost stress	13.05% / 12.58% CAGR	Shows resilience to materially worse execution costs.
Integer-share stress	10.09% CAGR	Useful lower-return operational reference for a small account.
Bootstrap p05	6.74% CAGR	A cautious within-sample downside planning reference.
Bootstrap p01	3.85% CAGR	A severe within-sample path reference; real life can still be worse.

Practical expectation

For planning purposes, it would be more prudent to budget around a materially lower return than the 14.30% backtest headline and to assume drawdowns can exceed the historical -7.12%. The observed stress tests and bootstrap suggest that a high-single-digit to low-double-digit outcome is a more conservative operational planning frame, but even that is not guaranteed.

What could make live performance worse?

- Survivorship/selection bias in the frozen historical universes.
- Future regimes unlike those represented in the historical sample.
- Opening-auction slippage materially larger than modeled.
- Fractional-share restrictions or routing differences at the broker.
- Alpha crowding, signal decay, or structural market changes.
- Missed orders, API outages, account constraints, corporate actions, or data corruption.
- Taxes and investor-specific cash flows, which are not part of the backtest.

Private quantitative research | Historical simulation - not audited live performance

10. Deployment protocol and conclusion

Recommended next stage: frozen paper trading

The research is now sufficiently mature to justify a frozen forward paper-trading stage, but not to claim live performance. The exact signal logic, 10/90 initial allocation, 20% Momentum ceiling, next-open order rules, sizing logic, and execution assumptions should be carried forward without parameter changes.

- Freeze all strategy parameters and portfolio-governance rules.
- Run in the IBKR paper account using the same order-generation engine intended for live deployment.
- Record every target, order, fill, rejection, partial fill, commission, and realized slippage.
- Compare paper fills with the backtest fill model and maintain a running slippage error distribution.
- Do not retune after a few disappointing trades; require several native rebalance cycles before judging forward behavior.
- Only consider real capital after the operational system reproduces the intended signals and orders reliably.

Overall assessment

The 10/90 governed portfolio is substantially more credible than the earlier synthetic weighted-return combination. Its strongest evidence is not any single performance statistic but the combination of: exact CTR0166 reproduction, exact Momentum implementation parity, true next-open account simulation, explicit cash/share accounting, natural sleeve drift, portfolio risk governance, execution-cost stress tests, integer-share stress, and block-bootstrap analysis.

Bottom line

The historical evidence supports continued development and paper deployment. It does not establish a guaranteed return. The strategy should be judged going forward by whether live/paper signals, orders, costs, and realized risk remain consistent with the frozen research design.

Research disclosure

This document summarizes a private quantitative research backtest. It is not investment advice, a solicitation, an audited performance report, or a representation that future returns will match historical simulations. Past simulated performance is hypothetical and subject to model risk. Historical universes retain survivorship/selection bias. Transaction costs and slippage are modeled. Real trading can lose money, including more than indicated by the historical maximum drawdown.

Source: DUAL-ALPHA 10/90 Governed live-realistic backtest output, run saved under C:
\\Quant_AI\\dual_alpha_live_realistic_9000_10_90_governed\\run_20260911_044058.