

Managed-Futures Trend as an Unresolved Diversifier: Complete Trial Lineage and Forward Burden of Proof

Arhan Canli

Founder, System Architect, and Quantitative Researcher
Canli Capital / AlphaC Algorithms

VERSION	0.1.0
RELEASED	23 August 2026
REGISTRY KEY	alphatrend_managed_futures
STATUS	Working paper preprint / not peer reviewed

Research simulation and, where explicitly identified, Alpaca paper evidence. No funded performance, peer review, independent replication, external acceptance, or future return is claimed.

Short title: AlphaTrend: managed-futures trend lineage

Author: Arhan Canli, Founder, System Architect, and Quantitative Researcher, Canli Capital

Research system: ALPHAC / AlphaForge

Family key: managed_futures_trend

Status: public research record; not peer reviewed; not an investment solicitation

Evidence date: 2026-08-22

Abstract

AlphaTrend is ALPHAC's managed-futures trend program. It takes volatility-normalized directional positions from medium- and long-horizon price trends across liquid exchange-traded market proxies, then combines them under a covariance-aware risk budget. This paper reconciles all 21 charged hypothesis identities in the family rather than presenting the strongest backtest alone.

The result is promising as a diversifier but statistically unresolved. Nineteen identities have finite Sharpe estimates ranging from -0.238 to +0.345. Six identities retain complete walk-forward artifacts, seven retain source-bound summary reports, and eight retain immutable ledger summaries without full curves. Thirteen have persisted artifact-era DSR values; zero clear the 0.95 gate. The 17-market baseline had Sharpe 0.289 and 9.69% vol-matched maximum drawdown in a replica screen. Correlation whitening reduced Sharpe to 0.186, a single 110-session signal reached 0.271, and expanding to 33 ETFs reduced Sharpe to 0.128 while worsening vol-matched drawdown to 14.42%. AlphaTrend remains a paper evidence-collection sleeve, not an established forward result.

Claim boundary

This paper establishes the implemented hypothesis, complete family trial count, persisted results, evidence gaps, and historical decisions. It does not establish future return, an achieved forward Sharpe, expected maximum drawdown, capacity, live-money performance, or external attestation. Paper execution through Alpaca is clearly separated from historical simulation.

Economic hypothesis and literature


The falsifiable hypothesis is that the sign and magnitude of an asset's medium-horizon excess-price move contains information about its next return, and that combining trends across genuinely different markets improves risk-adjusted performance. Possible mechanisms include slow information diffusion, investor underreaction, positioning constraints, and compensation for bearing reversal risk. These mechanisms motivate a prior; they do not validate ALPHAC's implementation.

- Moskowitz, Ooi, and Pedersen document time-series momentum across futures and forward markets ([Journal of Financial Economics, 2012](#)).
- Hurst, Ooi, and Pedersen examine long historical evidence for trend following ([Journal of Portfolio Management, 2017](#)).
- Lempérière and coauthors study trend following over roughly two centuries and emphasize crisis behavior and risk management ([Journal of Investment Strategies, 2014](#)).

- Baltas and Kosowski analyze momentum and trend construction in futures markets ([SSRN working paper](#)).

ALPHAC implementation

The core signal blends three volatility-normalized price trends:



Weights are determined from information available through session t and become effective on the next session. The principal ETF implementation uses 17 proxies spanning equities, rates, commodities, and currencies, a 33-session volatility estimate, covariance-aware allocation, monthly rebalancing, gross and name caps, 6 basis points per side, and 50 basis points annualized short borrow. The full walk-forward path uses purged anchored-expanding windows. A separate futures-proxy experiment is retained as a distinct identity because its instruments and data path are different.

Trial lineage

Identity group	Charged identities	Evidence retained
Core signal/cadence/universe studies	14	Six full walk-forwards; eight immutable ledger summaries
Allocator and signal-horizon screen	3	Source-bound summary report
Fixed and selected breadth screen	4	Source-bound summary report
Total	21	Every identity remains in the family and union burden

The breadth count includes two post-hoc procedures: dropping the largest full-sample PnL contributor and greedily adding markets only when estimated effective breadth rises. These are not free diagnostics. They generated measured return configurations and are charged as selection identities. Rounded equality between the ARP baseline and breadth baseline is not used to collapse them because their exact return curves were not preserved together.

The machine-readable source of truth is [/glassbox/alphatrend family.json](#). It binds every hypothesis key to the immutable ledger, strongest available artifact, source hash, exact configuration, evidence grade, and missing measurements.

Complete walk-forward evidence

Configuration	Sharpe	Max DD	Turnover/yr	Artifact-era DSR	Decision
63/126/252 blend, 17 ETFs, 21-session rebalance	0.263	10.29%	6.44×	~0.000	Fail DSR
126-session signal	0.175	13.32%	5.47×	0.704	Fail DSR
252-session signal	0.182	11.52%	4.93×	0.700	Fail DSR
Blend, 5-session rebalance	0.278	10.91%	12.31×	0.798	Fail DSR
Blend, 10-session rebalance	0.311	10.19%	9.39×	0.816	Fail DSR
Real-futures proxy path	-0.238	16.03%	8.97×	~0.000	Fail DSR

The highest complete-artifact Sharpe is not adopted as proof. It was one of several cadence and signal choices, remained below the historical DSR gate, and increased turnover materially.

Allocator and signal-horizon screen

All three rows share 5,161 sessions and are historical summaries without daily curves or kurtosis. Maximum drawdown is shown after matching each stream to baseline realized volatility.

Variant	Sharpe	Realized vol	Vol-matched max DD	Turnover/yr	Artifact-era DSR
Baseline 63/126/252 blend	0.289	4.49%	9.69%	4.65×	0.000
Correlation-whitened allocator	0.186	2.75%	12.76%	7.79×	0.000
Single 110-session signal	0.271	4.56%	10.57%	5.31×	0.000

Correlation whitening did not improve the baseline. The 110-session signal was close but weaker. Neither is promoted.

Breadth experiment

Basket/procedure	Sharpe	N_eff	Avg. pair correlation	Vol-matched max DD	Decision
Baseline 17	0.289	13.52	0.127	9.69%	Control
Expanded 33	0.128	11.01	0.250	14.42%	Reject
Expanded minus largest contributor (SHY)	0.217	10.11	0.264	10.04%	Reject
Greedy N_eff-pruned 22	0.204	14.76	0.153	9.58%	Reject

More tickers did not mean more independent bets. The 33-market expansion raised average correlation, reduced effective breadth, reduced Sharpe, and worsened drawdown. The selected 22-market basket raised estimated N_eff relative to baseline but still reduced Sharpe. This is a useful null: breadth acquisition must be exposure-aware and independently validated, not counted by ticker quantity.

Uncertainty, missing evidence, and capacity

Eight identities have only immutable ledger summaries; two of those have non-finite Sharpe because their reduced instrument sets did not support a valid estimate. They remain charged and visible. Summary-only ARP and breadth reports preserve no exact daily curves or kurtosis, so current-union DSR cannot be honestly reconstructed for those rows. The family packet therefore reports original DSR only where it exists and never substitutes a rounded estimate.

No persisted AlphaTrend capacity sweep measures market impact, ETF creation/redemption effects, borrow availability, or stressed liquidity across AUM. Capacity is **unmeasured**. A gross cap and liquid ticker list are implementation controls, not capacity evidence.

Paper execution and forward burden

AlphaTrend is one of the broker-executed Alpaca paper sleeves. Its public paper record is useful only if it remains continuous, broker-reconciled, and configuration-controlled. The record is too short to establish a forward Sharpe or the program's expected maximum-drawdown objective. Historical simulations selected from the 228-identity union cannot be relabelled as forward evidence.

The family earns at most one sleeve slot. It remains in evidence collection because the mechanism is distinct and the paper program needs diversifiers, not because a backtest passed the governing standard. Any production change requires a versioned declaration; any future research variant must be registered before returns are inspected.

Reproduction and provenance

Run from the repository environment:



The audit reads only immutable experiment ledgers and persisted artifacts. It opens no holdout, records no trial, and fails if identity count, artifact Sharpe, source hash, or evidence mapping drifts. Authorship and research leadership are credited to Arhan Canli; that credit does not alter the statistical claim boundary.

References

1. Tobias J. Moskowitz, Yao Hua Ooi, Lasse Heje Pedersen (2012). *Time series momentum*. Journal of Financial Economics. <https://doi.org/10.1016/j.jfineco.2011.11.003>
2. Brian Hurst, Yao Hua Ooi, Lasse Heje Pedersen (2017). *A Century of Evidence on Trend-Following Investing*. The Journal of Portfolio Management. <https://doi.org/10.3905/jpm.2017.44.1.015>

3. Yves Lempérière, Cyril Deremble, Philip Seager, Marc Potters, Jean-Philippe Bouchaud (2014). *Two centuries of trend following*. The Journal of Investment Strategies.
<https://doi.org/10.21314/jois.2014.043>
4. Nick Baltas, Robert Kosowski (2011). *Momentum Strategies in Futures Markets and Trend-following Funds*. . <https://doi.org/10.2139/ssrn.1968996>