

Crypto Defensive Factors: Two Positive Summaries Without Admission Evidence

Arhan Canli

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

VERSION	0.1.0
RELEASED	23 August 2026
REGISTRY KEY	crypto_defensive
STATUS	Working paper preprint / not peer reviewed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: Crypto defensive factors: complete trial lineage
Author: Arhan Canli, Founder and Quantitative Researcher, Canli Capital
Family key: crypto_defensive · **System:** ALPHAC / AlphaForge
Status: public research record; not peer reviewed; not an investment solicitation
Evidence date: 2026-08-22

Abstract

This record reconciles two cross-sectional defensive-factor identities in crypto perpetual futures: 720-hour low beta and 720-hour low realized volatility. The immutable ledger summaries report annualized Sharpe ratios of 0.6295 and 0.6946 over 1,247 observations. Those are positive historical summaries, but they are not admission results. Neither identity retains a complete identity-matched curve, maximum drawdown, current-union deflated Sharpe ratio, capacity study, or broker-reconciled forward record. A later low-volatility reopening reproduced a lower Sharpe of 0.5229 and passed four diversification diagnostics, but its venue-reality gate remained outstanding and its DSR remained 0.04. The family therefore contributes zero sleeves.

Research question and mechanism

The test asks whether crypto assets with lower trailing market beta or lower realized volatility earn better subsequent risk-adjusted returns than their more aggressive peers. The economic prior is not that low risk mechanically causes high return. It is that leverage constraints, delegated benchmarks, and demand for lottery-like payoffs can raise the price of high-beta or high-volatility assets relative to defensive assets. Frazzini and Pedersen formalize the leverage-constraint channel in conventional markets ([Journal of Financial Economics](#)).

Transfer to crypto is uncertain. Perpetual funding, listing age, liquidation risk, exchange fragmentation, and unstable token liquidity can all make a low-volatility rank a disguised liquidity or maturity exposure. Low beta and low volatility are consequently treated as two implementations of one defensive hypothesis, not two independent sleeves.

Frozen implementations and trial accounting

Both identities use rank allocation, a 63-bar rebalance interval, a 10-basis-point no-trade band, an 8,760-bar training window, and a 2,184-bar test window. The recorded configuration universe contains 3,112 instrument identifiers. The signal names and immutable results are:

Identity	Signal	Observations	Sharpe	Skew	Kurtosis
56cada31ea70f66a	beta_lowbeta_720	1,247	0.6295	0.3642	9.9238
e204cb10f8ee937c	lowvol_720	1,247	0.6946	-0.1637	18.8212

The high kurtosis of both summaries matters. A Sharpe estimate can conceal concentrated gains or losses when returns are far from Gaussian. Without the exact identity-matched curves, the tails, drawdowns, event concentration, and dependence of the two estimates cannot be reconstructed honestly. The two identities remain charged to the 228-identity research union even though neither is admitted.

Reopening evidence and contradiction

A related low-volatility reopening is preserved at `artifacts/analysis/lowvol720_reopen/result.json`. It is useful because it tests portfolio overlap, but it is not substituted for the missing historical curve. The reopening reproduced candidate Sharpe at 0.5229 rather than the ledger summary of 0.6946. That discrepancy is disclosed rather than averaged away.

The reopening passed four preregistered diagnostics. Correlation with the existing book was 0.1243 at the selected lag, below the 0.15 kill threshold. Correlation with crypto carry was 0.0877, below 0.25. Regression on BTC and a dominance proxy produced an R-squared of 0.2039, below 0.50. Adding the candidate at the tested weight increased the common-window book Sharpe by 0.2687. These measurements support distinctness within that historical window. They do not establish return reliability, execution, or capacity.

The fifth gate, venue reality, was never completed. The reopening explicitly retained DSR 0.04. Passing correlation screens cannot repair a failed selection-adjusted probability or replace a forward execution record.

Evidence still required

Admission would require a newly frozen, point-in-time implementation with one declared return identity; exact daily returns; funding, fees, slippage, liquidation, and venue assumptions; a current-union DSR; drawdown and tail analysis; impact and capacity evidence; and a continuous broker or exchange-reconciled forward record. The low-beta and low-volatility variants would also need joint exposure attribution to show whether they contain more than one economic bet.

No such completion is inferred from the positive summaries. There is no Alpaca performance attributed to this family, and no live-money result is claimed.

Reproduction and decision

The machine-readable family record is `crypto_defensive_family.json`. It binds both hypothesis keys, configurations, first ledger summaries, source hashes, and the related reopening artifact. The global trial join is `trial_packet_manifest.json`.

Decision: NOT ESTABLISHED / zero sleeves. The historical signal is interesting enough to justify a properly frozen prospective test, but not strong enough to enter the book. Research, implementation, and the decision to preserve the contradiction were directed by **Arhan Canli**.

References

1. Andrea Frazzini, Lasse Heje Pedersen (2014). *Betting against beta*. Journal of Financial Economics. <https://doi.org/10.1016/j.jfineco.2013.10.005>.