Crypto Perpetual Carry: Trial Lineage, Capacity Failure, and Live Evidence Boundary

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

VERSION	0.1.0
RELEASED	23 August 2026
REGISTRY KEY	alphaforge_crypto_carry
STATUS	Working paper preprint / not peer reviewed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: AlphaForge crypto carry: complete trial lineage

Author: Arhan Canli, Founder, System Architect, and Quantitative Researcher, Canli Capital
Research system: ALPHAC / AlphaForge
Family key: `crypto_carry`
Status: public research record; not peer reviewed; not an investment solicitation
Evidence date: 2026-08-24

Abstract

AlphaForge studies cross-sectional carry in USDT-margined perpetual futures, ranking contracts by the negative trailing mean of observed funding: expensive longs rank as shorts and negative-funding contracts as longs. This paper reconciles all 25 charged hypothesis identities.

The historical `carry_fund_21` walk-forward grew a simulated $100,000 to $138,236.27 from 2022-02-08 through 2026-06-01, with Sharpe 0.6766 and maximum drawdown 19.60%, but failed its DSR gate. A later eight-configuration matrix reported PBO 0.8818, no deflated winner, and a **NO-DEPLOY** verdict. Seven identities belong to a broader multi-factor engine, making their capacity sweep context rather than carry capacity; the sole carry-tilted arm had Sharpe 0.0618.

An exact-timestamp current-state replay on 2026-08-23 does not reproduce the selection: final equity is $103,335.16, Sharpe 0.1065, and maximum drawdown 20.98%. At the first decision, the historical 22 names included EOS while the current 21 did not; restoring EOS reproduces all ten historical order quantities. A zero-trial audit binds every surviving ledger: overlapping prices, position and funding marks, and funding rates agree, but holdings and state paths diverge. Because the historical run sealed neither its exact code nor derived-input snapshots, a unique additive attribution of the multi-year gap is structurally unidentifiable. Prospective private snapshots are now enforced but cannot repair that omission. The material correction therefore remains open and external submission is blocked.

The forward record is weaker: AlphaForge is a local simulated paper broker, not an Alpaca account or externally attested record. Its first configuration was signal-dead; funding cashflows were omitted until a forward-only repair; fewer than half of intended hourly cycles completed; and the stored universe remains frozen at 2026-06-01. Crypto carry is therefore an incumbent experiment with a plausible mechanism, severe selection and capacity evidence, known operational defects, and an unmet forward burden of proof, not a validated sleeve.

Claim boundary

This paper supports six claims only:

1. The funding-carry hypothesis and point-in-time implementation can be stated precisely.
2. Every charged family identity is enumerated in a public machine-readable manifest.
3. The selected positive artifact and later negative carry/blend evidence are published together.
4. A related multi-factor engine's capacity degradation is contextual evidence only and is not presented as carry capacity or investable capacity.

5. The paper record, defects, corrections, and unresolved production gaps are disclosed.
6. The selected artifact is preserved as a historical result, but its current exact reproducibility is disproved and the open correction is machine-verifiable.

It does **not** establish a forward Sharpe, expected maximum drawdown, future return, live-money performance, riskless arbitrage, or externally verified execution. It does not establish that the current live universe is the same strategy as the selected historical artifact.

Economic mechanism and falsifiable hypothesis

Perpetual futures have no expiry date, so exchanges use recurring funding transfers to pull the contract toward its reference index. When the funding rate is positive, long holders pay short holders; when it is negative, shorts pay longs. A cross-sectional carry strategy therefore takes the opposite side of expensive positioning: it shorts persistently positive-funding contracts and buys persistently negative-funding contracts.

The mechanism is compensation for constrained arbitrage, speculative demand, margin and liquidation risk, not a guaranteed convergence trade. It is falsified for this implementation if a point-in-time, net-of-fee, impact-aware walk-forward cannot remain positive after the complete search is charged, if the return disappears at useful capital, or if operational constraints make the intended decisions untradeable.

The literature supplies a prior, never proof of this implementation:

- He, Manela, Ross, and von Wachter show how funding anchors perpetuals to spot while margin, liquidation, and trading frictions prevent a riskless arbitrage interpretation (Fundamentals of Perpetual Futures).
- Ackerer, Hugonnier, and Jermann develop no-arbitrage pricing for perpetual futures and the funding mechanism (NBER working paper).
- Schmeling, Schrimpf, and Todorov document large, time-varying crypto carry and connect it to trend chasing and scarce arbitrage capital (Crypto Carry).
- Gornall, Rinaldi, and Xiao emphasize basis risk, constrained arbitrage capital, and speculative demand in perpetual markets (Funding Payments Crisis-Proofed Bitcoin's Perpetual Futures).
- Chi and coauthors find that basis is the strongest of the crypto-futures cross-sectional signals they study, while signal frequency matters (Journal of Futures Markets, 2023).

These sources do not validate ALPHAC's market data, universe, costs, trial selection, or paper execution.

Exact ALPHAC implementation

For instrument `i`, decision time `t`, and the last `K` funding settlements known by `t`, the historical feature is:

The negative sign makes a high positive funding rate unattractive to a long position. The factor uses a backward point-in-time as-of join: only settlements whose stored publication timestamp is available by `t` may enter the feature. `carry_fund_21` uses 21 settlements, historically about seven days at three settlements per day; `carry_fund_90` uses 90 settlements, historically about 30 days.

The `3 * 365` annualization is an implementation assumption, not a universal exchange rule. Binance exposes `fundingIntervalHours` because intervals can be adjusted. Its official market-data API separately exposes funding history and current interval information (Binance USD-M Futures market data). The ranking sign and settlement count remain meaningful when intervals change, but the annualized scale is then approximate unless the contemporaneous interval is applied. This limitation is disclosed rather than silently treating every period as eight hours.

The selected artifact uses rank allocation, 25 legs, weekly rebalancing (168 hourly bars), a 10-basis-point no-trade band, 6,048 training bars, 1,512 test bars, 72 purge bars, and 168 embargo bars across 58 Binance perpetual instruments. Other charged identities change the funding horizon, blend, cadence, allocator, no-trade band, or regime treatment. Each change remains a distinct hypothesis identity even though all belong to one economic family.

Trial lineage

The public manifest binds 25 distinct return identities to `crypto_carry`. The labels below sum exactly to the family total; they are an accounting view and do not collapse the trial denominator.

Machine-label group	Charged identities	Research question
`carry_fund_21`	10	Primary seven-day funding rank under cadence, allocator, band, and regime changes
`carry_fund_21 + carry_fund_90 + carry_z_252`	3	Whether three carry horizons improve robustness
`carry_fund_21 + carry_fund_90 + carry_z_252 + carry_mom_21_63`	6	Whether carry dynamics add information across the original and recovered legacy ledgers
`carry_fund_90`	2	Whether a slower funding window is more persistent
`carry_z_252`	2	Whether long-window normalization improves ranking
`carry_fund_21 + carry_fund_90 + carry_mom_21_63`	1	Whether a reduced dynamic blend improves carry
`carry_fund_21 + carry_fund_90 + mr_res_72`	1	Whether a carry-plus-residual-reversal tilt improves the multi-factor engine
Total	**25**	Every identity remains charged to the union search burden

Twenty-one identities use the original crypto-perpetual research profiles. Four recovered legacy identities were filed with their durable result artifacts: one original carry experiment, two shorter-window carry experiments, and the grand matrix's carry-plus-reversal tilt. Their recorded annualized

Sharpe observations range from -1.3945 to +0.6766. That range is not a confidence interval; it is direct evidence of specification sensitivity.

The machine-readable source of truth is `/glassbox/trial_packet_manifest.json`. Select identities where `research_family_key` equals `crypto_carry`; each row includes its exact hypothesis key, immutable first measurement, source ledger, paper binding, and missing packet sections.

Results and decisions

Selected historical survivor

> **Open material correction (2026-08-23).** The table below reports the immutable selected historical artifact, not current-code performance. A current-state replay reports Sharpe 0.1065 and 20.98% maximum drawdown, and the original run did not bind every derived input. The result is therefore labelled **historical artifact; not currently exactly reproducible; open correction**. It must not be described as validated or submission-ready.

The persisted `crypto_carry_wk` artifact covers 1,575 daily observations from 2022-02-08 through 2026-06-01. It reports:

Measure	Persisted value
Initial / final simulated equity	$100,000 / $138,236.27
Total return / CAGR	38.24% / 7.80%
Annualized Sharpe / volatility	0.6766 / 11.83%
Maximum drawdown	19.60%
Annual turnover	33.22x
Fees paid	$7,372.72
Net funding cashflow	$19,500.02
DSR / DSR gate	0.0386 / failed

Funding contributed roughly half of the artifact's dollar gain. The result is therefore mechanically dependent on accurate funding settlement and cannot be validated by a live ledger that records price and fees but omits funding. The full-history Sharpe also hides a severe 2022 tail: the recorded 2022 segment was approximately -1.63 Sharpe and the full artifact's maximum drawdown was 19.60% around the LUNA/FTX period.

Related system-level robustness study

The later grand backtest used a 2021-01-01 through 2026-06-01 window and eight distinct configurations. It must not be described as a pure carry replication. Seven configurations use the engine's default, broad crypto multi-factor signal set (`alpha_names: null`) while varying ML, regime, cadence, band, or allocator. Only the eighth explicitly tilts to `carry_fund_21 + carry_fund_90 +`

`mr_res_72`, and that arm reported Sharpe 0.0618. The seven multi-factor configurations belong to the separate `crypto_multifactor_engine` research family; the carry-tilted arm remains charged to `crypto_carry` because carry is its primary named input.

All four Block A multi-factor variants produced the same Sharpe 0.0424 and DSR 0.2112 at $1 million. Other multi-factor variants reported annualized Sharpes of -0.7901 for daily rebalancing, -0.1994 for the tighter 10-basis-point band arm, and -0.0946 for MVO. None cleared DSR 0.95.

Cross-sectional-combinatorial validation reported PBO 0.8818 against a gate below 0.20. There was no deflated winner and the persisted verdict is **NO-DEPLOY (honest null)**. Regime and ML labels did not rescue the result; the regime path was substantially inactive during cold-start periods, so equality of outcomes is not evidence that regimes never matter.

The selected carry artifact and broad multi-factor matrix are not interchangeable. The matrix is preserved because it contains one carry-tilted arm and valuable system-level stress evidence, but its baseline, PBO, winner decision, and capacity curve cannot be represented as a direct institutional re-grade of `carry_fund_21`. The carry-specific conclusion comes from the weak 0.0618 tilted arm, the recovered legacy carry experiments, the complete 25-identity lineage, and the failed selected-artifact DSR, not from relabelling a broader system as carry.

Capacity: related-system decay, not measured carry capacity

The grand matrix's capacity sweep varies capital for its default broad multi-factor configuration without adding a new return hypothesis:

Initial capital	Sharpe	DSR	Maximum drawdown	Final equity
$100,000	0.4009	0.4803	13.14%	$118,509.86
$1,000,000	0.0424	0.2112	13.59%	$993,219.39
$10,000,000	-0.3720	0.0460	22.69%	$8,043,578.96

This is evidence that the related engine is highly sensitive to market impact, not a measurement of pure carry capacity. The broad engine's edge is almost absent at $1 million and negative at $10 million under the committed model. Carry capacity is **unmeasured**: no persisted carry-specific capital sweep exists. It remains unestablished until a preregistered carry sweep and forward fills, order-book depth, venue fragmentation, and financing are observed at the intended scale.

Forward paper record and dated corrections

AlphaForge uses `ALPHAFORGE_PAPERBROKER`: locally simulated fills against live exchange order books with a local ledger. It is paper-only, not an Alpaca account, not live money, and not externally attested. ALPHAC's public composite is derived from its sleeves and is not itself a broker account.

The following defects materially limit interpretation:

- **Historical replay divergence, correction opened 2026-08-23.** A zero-new-trial replay on the current local state retained all 37,776 timestamps; the replay-minus-historical final-equity delta was -34,901.11 dollars. Sharpe fell from 0.6766 to 0.1065 and maximum drawdown rose from 19.60% to 20.98%. The old receipt wording incorrectly said the run isolated software drift and gave the non-exact outcome a `PASS` label. It has been re-sealed without rerunning as `REPLAY_EXECUTED_MATERIAL_CODE_OR_MUTABLE_INPUT_DRIFT_QUANTIFIED`. The first sizing difference is exactly attributable to mutable universe membership: EOS was present in the historical decision set and absent today. The corrected realized-vol overlay also changes the later path and bound 71 rebalances. The complete surviving path audit shows exact overlapping exogenous values but divergent quantities and risk states. It also proves why the omitted historical code, universe, metadata, and signal snapshots prevent a unique additive attribution of the terminal gap. Repeating the current-state replay cannot restore missing history. A prospective control now atomically freezes the exact signal frame, universe intervals, SCD2 metadata, execution partitions, resolved settings, dirty source tree, and lockfile before every persisted walk-forward. That control is deliberately not presented as retroactive evidence. The original artifact remains preserved; no historical number is silently overwritten.
- **Signal-dead start, corrected 2026-07-05.** From 2026-06-23 through 2026-07-05 the public explanation said carry was compressed. The actual cause was a wiring bug that blended equity-fundamental factors, undefined for crypto, into the signal and invalidated every cycle. The loop was then restricted to `carry_fund_21` with weekly rebalancing and signal-health logs.
- **Funding omitted, repaired forward-only.** The live paper broker initially moved cash only on fills, so the mechanism's funding cashflows never reached the account. A live settlement path was added with point-in-time publication filtering and tests. Historical marks were not restated, which preserves continuity but leaves the early record economically incomplete.
- **Structural venue availability.** As of the evidence date, the loop had completed fewer than half of its expected hourly cycles. The current machine-readable paper state publishes the exact rolling count. A strategy that is unmanaged for most intended cycles is not meaningfully continuous even when its observed marks are retained.
- **Commodity contamination remains unresolved.** The 2026-08-22 store still contains a 20-member universe effective from 2026-06-01. Four contracts, XAU, XAG, CL, and BZ, reference gold, silver, WTI, and Brent rather than cryptoassets. Removing them would change the strategy and therefore requires preregistration and remeasurement, not a silent filter.
- **Universe refresh remains unresolved.** The newest stored crypto membership decision is 2026-06-01; the July and August decisions are absent. Catch-up logic exists in `LiveLoop`, but the production constructor does not pass a universe refresher. The current live book therefore remains based on stale liquidity membership. No missed historical rebalance is reconstructed.

These corrections are evidence about process quality only if they remain visible and are pinned by tests. They are not evidence that the return process is profitable.

Selection, diversification, and governing targets

The family is one economic sleeve candidate, not 21 independent sources of alpha. Every return hypothesis remains charged to the union ledger. It also cannot receive diversification credit from old correlations measured before commodity-linked contracts overlapped AlphaTrend's metals and energy exposures. Correlation must be remeasured on the actual contemporaneous books and stressed in risk-off periods before any sleeve-count or diversification claim is made.

The governing objective is an honest forward Sharpe of 1.5, expected maximum drawdown near 11%, and a diversified book of up to 14 sleeves. Crypto carry has established none of those targets. Its selected historical maximum drawdown already exceeds the program objective, its institutional re-grade is a null, and its forward execution record is too short and discontinuous to estimate a reliable Sharpe.

Reproduction map

Primary implementation and evidence paths:

- `src/alphaforge/features/library/carry.py`: point-in-time funding features.
- `src/alphaforge/data/sources/ccxt_source.py`: funding interval metadata ingestion.
- `src/alphaforge/execution/paper.py` and `src/alphaforge/live/loop.py`: paper fills and live funding settlement.
- `docs/design/alphaDesign.md`: original signal and annualization contract.
- `var*/experiments.jsonl` and the four recovered `artifacts/**/experiments.jsonl` ledgers , immutable charged identities in the canonical durable union.
- `artifacts/walkforward/crypto_carry_wk/walkforward.json`: selected historical survivor.
- `artifacts/probe/crypto_carry_frozen_current_code_replay/replay_receipt.json`: truthful current-state non-reproduction receipt, metric deltas, and input-freeze audit.
- `artifacts/probe/crypto_carry_replay_drift/first_rebalance_attribution.json`: exact causal reconstruction of the EOS membership effect on every first-decision order quantity.
- `artifacts/probe/crypto_carry_replay_drift/full_path_attribution.json`: exhaustive surviving ledger comparison, code archaeology, and machine-readable non-identifiability proof.
- `artifacts/audit/walkforward_input_snapshot_protocol.json` and `src/alphaforge/validation/input_snapshot.py`: prospective pre-run snapshot enforcement, atomicity, private data-rights boundary, and tamper validation.
- `artifacts/publication/crypto_carry_replay_correction.json`: open material correction, publication decision, remediation, and zero-new-trial classification.
- `artifacts/grand_backtest/20260616T143620Z/matrix.json` and `verdict.md`: related multi-factor robustness, PBO, capacity, and no-deploy context; only `C_carry` is carry-tilted.
- `artifacts/research/trial_packet_manifest.json`: exact family/identity join and packet debt.
- `/glassbox/crypto_carry_selected_walkforward.json` : public selected-artifact payload.
- `/glassbox/crypto_carry_grand_matrix.json`: public grand-matrix payload.

- `/glassbox/crypto_carry_2022_tail.json`: reproducible UTC-daily 2022 tail derivation.

Reproduction requires the pinned project environment and the underlying market-data lake. Public artifacts permit claim auditing; they do not imply that every venue dataset can be redistributed. The full current-state equity Parquet remains local and is bound by its SHA-256 in the replay receipt; it is withheld from the public bundle under the conservative source-rights policy because it is a row-level derived series.

Packet completeness and legacy limitations

This family paper verifies shared identity, authorship, mechanism, literature, family accounting, and stable-publication sections for 25 identities. Its 2026-08-23 correction proves that the selected output is not currently exactly reproducible; this is a publication blocker, not a 26th return identity. It does not retroactively manufacture identity-level preregistrations, return series, exact environment snapshots, or complete rerunnable data bundles. The manifest must therefore continue to mark all 25 trial packets incomplete until every missing required section is proved. A family-paper binding is not a completed packet.

Decision

Research decision: retain AlphaForge only as an incumbent paper experiment under observation; do not describe the historical edge as validated; do not claim positive $10 million capacity; and do not admit this family under the current evidence.

Operational decision: fail closed on promotion while the live universe is stale, commodity classification is unresolved, and intended-cycle continuity remains structurally weak. Preserve all prior marks and corrections; repair forward without reconstructing decisions that never ran.

Publication decision: preserve the positive survivor, failed DSR, current-state replay, first-decision attribution, weak carry-tilted matrix arm, the broader engine's PBO 0.8818 and capacity decay, live execution provenance, and unresolved defects together under Arhan Canli's authorship. External submission is blocked while the material replay correction is open. A corrected release must present the historical and current-state results side by side, bind every derived input, and undergo independent scrutiny; no more favourable retrospective window may replace either result.

References

1. Damien Ackerer, Julien Hugonnier, Urban Jermann (2024). *Perpetual Futures Pricing*. . https://doi.org/10.3386/w32936
2. Maik Schmeling, Andreas Schrimpf, Karamfil Todorov (2022). *Crypto Carry*. . https://doi.org/10.2139/ssrn.4268371
3. Will Gornall, Juan Martin Rinaldi, Yizhou Xiao (2025). *Funding Payments Crisis-Proofed Bitcoin's Perpetual Futures*. . https://doi.org/10.2139/ssrn.5036933

4. Yeguang Chi, Wenyan Hao, Jiangdong Hu, Zhenkai Ran (2023). *An empirical investigation on risk factors in cryptocurrency futures*. Journal of Futures Markets. https://doi.org/10.1002/fut.22425
5. Songrun He, Asaf Manela, Omri Ross, Victor von Wachter (2022). *Fundamentals of Perpetual Futures*. arXiv. https://arxiv.org/abs/2212.06888
6. Binance (2026). *Get Funding Rate History*. Official documentation. https://developers.binance.com/docs/derivatives/usds-margined-futures/market-data/rest-api/Get-Funding-Rate-History