humarb Research 08 · Methodological and Descriptive Empirical Preprint

Instrument Equivalence Across Perpetual-Futures Venues: Evidence Levels, Unit Normalization, and Reference Prices

Hang Liu · humarb Research, London, United Kingdom · https://humarb.ai · ORCID 0009-0008-1059-3926

Prepublication author manuscript prepared 4 November 2026 · Not peer reviewed
DOI pending Zenodo publication; after publication, cite the Zenodo concept DOI

Abstract

Perpetual futures with similar names need not represent the same economic claim. We propose an equivalence framework that separates seven evidence levels, normalizes contract units, and classifies pairs as exactly convertible, reference-basis pairs, or hold-outs. Using official public snapshots, we enumerate 523 Hyperliquid metadata entries across the default and ten named perp DEXs on 20 September 2026, of which 325 were not delisted. For five pre-IPO companies on 14 October 2026, canonical implied equity values show that OKX's native-price ratios of 0.1011 and 0.1031 become residuals of +1.15% and +3.08% after its documented tenfold adjustment. Across 26 eligible exact venue observations, the largest remaining absolute residual is 9.53%, indicating reference-price differences rather than a common unit error. Gate KIMI remains unclassified because its stated and example-implied share counts conflict. A BingX documentation example separately shows a 1,000-fold difference between ordinary and Coin quantity fields. The snapshots do not establish account eligibility, executable size, convergence, or net returns.

Keywords: perpetual futures; instrument equivalence; contract units; reference prices; pre-IPO markets; market microstructure

JEL classifications: G12; G13; G14; C58 (American Economic Association, 2026)

Corresponding author: Hang Liu, humarb Research, London, United Kingdom; humarb.ai; admin@humarb.ai.

Contents

1 Introduction **4**
1.1 Research question: when is a familiar ticker actually the same claim? 4
1.2 Evidence and principal findings 4
1.3 Contributions and relationship to humarb Research 01 and 02 5

2 Literature **6**
2.1 Perpetual-futures design presupposes a defined claim 6
2.2 Cross-market prices, basis, and limits to arbitrage 6
2.3 Instrument design supplies negative controls 7
2.4 Execution costs require quantity semantics 7
2.5 Evaluation methods: applicable standards and deliberate non-application 8

3 Institutional Background **8**
3.1 Builder-deployed markets make identity venue- and version-specific 8
3.2 Oracle transitions and closed external markets prevent trivial price comparisons . 9
3.3 Funding formulas have windows, multipliers, clamps, and dead zones 9
3.4 Reference-contract rolls and synthetic indices can diverge by construction 10
3.5 Pre-IPO rebasing changes units without necessarily changing valuation 10

4 Data **11**
4.1 Sources and observation windows 11
4.2 Inclusion, exclusion, and preservation rules 11
4.3 Missing evidence and inaccessible interfaces 13
4.4 Descriptive coverage and temporal recheck 13

5 Methods **14**
5.1 Evidence levels are not interchangeable 14
5.2 Canonical value and quote-currency conversion 14
5.3 Residual definitions and anchors 15
5.4 Classification rules 15
5.5 Quantity-aware depth walking 16
5.6 Identification assumptions and counterfactual interpretation 16
5.7 Symbol table 17

6 Results **17**
6.1 Public coverage stops well short of executable equivalence 17
6.2 The documented OKX factor explains most of the raw gap 18
6.3 Remaining residuals vary with the reference construction 19
6.4 Conflicts and negative controls remain hold-outs 20
6.5 Quantity fields can change depth by orders of magnitude 21

7 Robustness **22**

7.1 Early and late windows support scale interpretation, not persistence 22
7.2 Venue groups and company hold-outs 22
7.3 Share-count sensitivity is decisive for Gate KIMI 22
7.4 Alternative anchors and residual forms preserve the classification 23

8 Limitations **23**
8.1 The evidence is cross-sectional and incomplete 24
8.2 Raw-response authentication is not available for every historical observation . . . 24
8.3 Anchors, marks, and quote currencies affect measurement 24
8.4 Equivalence does not imply permission, liquidity, or common settlement 25
8.5 Documentation changes and conflicts are intrinsic to the research object 25
8.6 Required next steps 25

9 Conclusion **25**

Declarations **26**

A Derivation of Rebase Invariance and Reference-Share Sizing **30**

B Hyperliquid Coverage and Order Books **31**

C Additional Robustness Tables **31**

D Reproduction Guide and Data Schema **31**

E Document Conflict and Hold-Out Register **33**

1 Introduction

1.1 Research question: when is a familiar ticker actually the same claim?

Cross-venue comparisons of perpetual futures often begin with a string match. Two contracts called OPENAI, NVDA, GOLD, or SP500 are aligned, their prices are subtracted, and the difference is treated as a basis. This sequence compresses several distinct questions into one unsupported assumption. Does the symbol exist in a document or in current metadata? Is the market active? Is a two-sided book observable? Is the account permitted to trade it? Do the contracts reference the same economic claim, in the same unit, with compatible settlement and corporate-action rules? Finally, can the intended size actually be executed? A failure at any one of these stages changes the meaning of the calculated spread.

This paper asks a deliberately narrow research question: **what evidence is required before two perpetual futures listed on different venues can be treated as equivalent, and how should their prices and quantities be normalized before any residual is interpreted?** The question is upstream of price discovery, funding comparisons, and execution. It is not answered by no-arbitrage pricing alone. No-arbitrage theory generally starts from a defined underlying claim and a trading technology; our problem is to establish whether the two instruments supply that claim in the first place. Nor is it answered by observing a tight price ratio. A common price can conceal an undocumented unit, while a large price difference can be a harmless consequence of a documented rescaling.

The distinction has become more important as perpetual markets have moved beyond a single crypto-asset universe. Hyperliquid's builder-deployed perpetual markets allow deployers to define the oracle and contract specifications and to operate independent order books and margin arrangements (Hyperliquid, 2026b). Public catalogues now include builder-specific equities, indices, commodities, foreign exchange, and pre-IPO markets. These products trade around the clock even when their external reference markets are closed. In parallel, pre-IPO perpetuals refer to private-company valuations using assumed share counts that can be rebased. The same company name may therefore appear with prices that differ by a factor of ten for a documented reason, or by several percent because two venues construct different reference prices. The numerical similarity or dissimilarity is not self-interpreting.

We develop an equivalence protocol with three layers. The first is a seven-level evidence ladder: listed in documents, present in metadata, not delisted or active, two-sided book observed, account allowed, economic equivalence reviewed, and size executable. The levels are cumulative as evidence requirements, not interchangeable labels. In particular, a public order book does not demonstrate account permission, and account permission does not demonstrate that two instruments are economically equivalent. The second layer is unit normalization. For a pre-IPO contract whose native price is based on an assumed share count N_v, we define a canonical implied equity value in USD billions. The third layer is classification. A pair is exactly convertible only when a common reference and unit conversion are supported; it is a reference-basis pair when the company or asset identity is shared but the reference construction or settlement differs; and it is held out when identity, status, or units conflict or remain unavailable.

1.2 Evidence and principal findings

The empirical analysis uses three public-evidence modules. First, a Hyperliquid snapshot dated 20 September 2026 enumerates the default perp DEX and ten named DEXs. It contains 523 metadata entries, of which 325 are not marked delisted. These are entries, not 325 unique assets

or executable pairs. The snapshot also records ten representative two-sided books. Second, a pre-IPO snapshot dated 14 October 2026 covers five companies and reports official public marks across a broad venue catalogue. The numerical analysis uses 26 eligible exact venue observations plus one pseudo-anchor and one conflicting Gate KIMI line; grouped ranges are retained separately and are never disaggregated into invented observations. Third, BingX's official API documentation supplies a SHIB-USDT order-book example in which ordinary quantities and Coin quantities differ by exactly 1,000 at all ten displayed levels. This example isolates quantity-field semantics rather than price units.

The unit results are sharp but deliberately bounded. OKX's native OPENAI and ANTHROPIC marks are 0.10114781 and 0.10308438 times the corresponding Binance marks in the later pre-IPO window. Multiplying OKX by its documented factor of ten produces implied equity values of USD 1,656.7 billion and USD 2,062.1 billion. Relative to the Binance anchors, the residuals are only +1.147811% and +3.084383%. Thus, most of the apparent roughly 90% raw gap is a scale mismatch, not evidence of an economically mispriced company. A one-to-one position-size assumption would hedge only one-tenth of the intended exposure, leaving 90% unhedged before considering any residual reference risk.

The remaining residuals are not uniform. The largest absolute value among the eligible exact observations is −9.530398% for Aster POLYMARKET against the supplied cross-venue median. Lighter's OPENAI and ANTHROPIC residuals are +3.208377% and +4.089182%; Gate's OURA residual is +5.050092%; Bitget's OPENAI residual is −3.544172%. These magnitudes are consistent with different reference constructions and do not collapse into a single unit explanation. The two pre-IPO windows on the same day repeat the direction and broad scale of several gaps, but two windows cannot establish long-run persistence, mean reversion, or convergence.

Gate KIMI illustrates why a numerical fit must not override documentary evidence. The announcement states a ten-billion-share assumption, yet its three worked examples imply 8.59598854, 8.58369099, and 8.59106529 billion shares. Using ten billion produces an implied value of USD 77.86 billion and a +15.966637% residual against Binance MOONSHOT. Using the mean example-implied count produces USD 66.883673 billion and a −0.381780% residual. The latter looks close, but choosing it merely because it fits would be circular. We therefore keep the line on hold.

The robustness analysis performs four checks appropriate to descriptive evidence: an early-versus-late window comparison, venue and unit-evidence grouping, share-count sensitivity, and alternative median anchors and log residuals. It does not report predictive accuracy, statistical significance, annualized income, or trading profit. Those estimands require time-series and execution evidence that the snapshots do not contain.

1.3 Contributions and relationship to humarb Research 01 and 02

The paper makes four contributions. First, it converts a collection of informal checks into an explicit evidence ladder, preventing a lower-level observation such as "listed" or "has quotes" from being silently promoted to "economically equivalent" or "executable." Second, it provides a reproducible normalization for pre-IPO perpetuals and shows how documented share-count rebasing changes the interpretation of cross-venue prices. Third, it separates exact conversion, reference-basis pairing, and hold-out decisions, including negative controls such as an index versus an exchange-traded fund (ETF), WTI versus Brent, and reciprocal FX quotes. Fourth, it supplies data, a data dictionary, and code that recompute every reported point estimate and figure while preserving conflicting and unavailable evidence.

The paper is an upstream companion to two humarb Research papers. Liu (2026a) measures executable net carry using contract-identity controls, settled funding history, depth-walked costs, and data-quality checks. Liu (2026b) formalizes a broader net-carry framework with asynchronous settlement and liquidation-aware accounting. Neither paper's funding or cost calculations can repair an incorrect instrument map. Our contribution is to define and test the identity and unit layer before carry is computed. We do not reproduce their return or cost estimates and do not claim that an equivalent pair is a viable trading route.

The rest of the paper follows the required sequence. Section 2 positions the study in the perpetual-futures, cross-market, arbitrage-limit, instrument-design, and evaluation literatures. Section 3 explains the institutional mechanisms. Section 4 documents samples and provenance. Section 5 defines the estimands and classification rules. Section 6 reports results, Section 7 reports robustness, and Section 8 sets out limitations before the conclusion.

2 Literature

2.1 Perpetual-futures design presupposes a defined claim

Perpetual contracts were proposed as a way to provide cash-settled exposure without a sequence of dated maturities (Shiller, 1993). Modern pricing work derives no-arbitrage values and bounds under specified funding mechanisms and market frictions (Ackerer et al., 2026). Contract design matters for arbitrage behavior and market structure in Bitcoin futures (De Blasis and Webb, 2022). This literature supplies the pricing logic that operates after the underlying claim, funding formula, and trading constraints have been defined. Our analysis addresses an antecedent measurement problem: whether two venue-specific instruments define the same claim and express it in compatible units.

That distinction is particularly relevant for pre-IPO markets. A conventional perpetual model can take the reference price as exogenous. A pre-IPO venue instead creates a synthetic valuation signal from an assumed share count and a trade-driven or internal reference. The reference is not ownership in the private company and does not provide voting or shareholder rights (Aster, 2026). If the assumed count changes, the native price can be rebased even when the implied equity value is intended to remain economically continuous. A pricing model that uses the native price without the share-count state therefore risks treating a mechanical redefinition as a market event.

The paper also differs from funding-rate studies. Funding remains relevant after equivalence has been established, but it cannot identify equivalence. Hyperliquid's native crypto contracts can be USDT-denominated while margined and settled in USDC without a USDT/USDC conversion, which is a quanto feature (Hyperliquid, 2026a). Lighter calculates funding hourly with market-specific multipliers and clamps (Lighter, 2026a). These are cash-flow specifications layered on top of instrument identity. A shared ticker does not imply a shared collateral, funding, or settlement convention.

2.2 Cross-market prices, basis, and limits to arbitrage

Price-discovery research compares multiple trading locations for a common security. Hasbrouck (1995) decomposes contributions to price discovery, while Brogaard et al. (2019) shows that limit orders can contribute even without trades. In cryptocurrency markets, Makarov and Schoar (2020) documents cross-exchange price differences and the role of market segmentation. Deviations from text-book parity can also persist in large conventional markets because balance-sheet

and intermediary constraints matter (Du et al., 2018). These studies motivate careful treatment of a residual as something other than a free arbitrage profit.

The limits-to-arbitrage literature is equally important but answers a different question. Shleifer and Vishny (1997) explains why arbitrage capital can be withdrawn precisely when mispricing is large. Gromb and Vayanos (2002) models separately collateralized positions in segmented markets, and Gromb and Vayanos (2010) reviews the broader theory. We do not estimate arbitrage capital, a convergence speed, or an equilibrium price. Instead, we ask whether the object being priced is the same on both legs. If one leg is an index and the other an ETF share, or if one leg's unit is ten times the other's, a convergence model is misspecified before capital constraints are even introduced.

This ordering changes empirical practice. A researcher should first establish identity and units, then study price discovery or basis dynamics on a time series of validated pairs. Tight contemporaneous pricing is neither necessary nor sufficient for identity: it is not necessary when documented units differ, and it is not sufficient when two independently constructed references happen to be close. Our two same-day pre-IPO windows illustrate sign repetition, but we deliberately leave dynamic convergence to a later time-series study.

2.3 Instrument design supplies negative controls

The ETF literature provides a useful negative control. An ETF is an investment vehicle with a creation, redemption, and ownership structure that seeks to track an index (Lettau and Madhavan, 2018). Its market price can deviate from net asset value, especially when underlying assets are illiquid or prices are stale (Petajisto, 2017). Therefore, an S&P 500 index level and an SPY share price are related but not interchangeable units. A one-to-one mapping would confuse index points with a fund share and suppress fees, distributions, corporate actions, and the ETF arbitrage mechanism.

Commodity futures offer a second control. Futures returns depend on curves, inventories, and rolling between contracts (Gorton et al., 2013). WTI and Brent differ in grade and market structure even when both are quoted in dollars per barrel. Two perpetuals that reference oil must therefore be checked for the actual commodity, reference contract month, roll weights, and effective roll time. A same-name match is not enough, and a cross-grade match is not an identity map. These negative controls operationalize the principle that economic relatedness is weaker than instrument equivalence.

2.4 Execution costs require quantity semantics

The net-carry framework in Liu (2026a) and Liu (2026b) evaluates funding after fees, depth-dependent entry and exit costs, and margin risks. A depth walk requires base quantities at each price level. If an API returns both an ordinary quantity whose unit is undocumented and a Coin quantity explicitly denominated in the base asset, choosing the wrong array can distort the measured depth and volume-weighted average price. Contract size, quantity precision, minimum order size, and an asset-name prefix such as 1000SHIB are separate transformations and should not be collapsed into one multiplier.

Our BingX case is a documentation-level identification test, not an execution-cost estimate. It shows a 1,000-fold ratio in an official example and thereby rejects the universal assumption that ordinary and Coin arrays are interchangeable. It does not show that every BingX symbol has the same ratio, nor that the ordinary field always requires multiplication by contract size. The safe procedure is to prefer explicitly documented units and otherwise leave the transformation

unknown.

2.5 Evaluation methods: applicable standards and deliberate non-application

Forecast evaluation requires a forecast and a realized outcome. Point forecasts can be compared using MAE or RMSE and Diebold–Mariano tests (Diebold and Mariano, 1995). When many models or parameters are tried, White's Reality Check, Hansen's Superior Predictive Ability test, and cross-sectional multiple-testing adjustments address data snooping (White, 2000; Hansen, 2005; Harvey et al., 2016). Time-series inference would also require an explicit dependence treatment, such as a Newey–West covariance estimator with a stated lag (Newey and West, 1987).

This paper estimates none of those objects. It computes deterministic unit conversions, residuals, and classifications from snapshots. Consequently, there is no in-sample trading rule to take out of sample, no forecast on which to perform a Diebold–Mariano test, and no family of return-generating models requiring a Reality Check. Assigning p-values to two same-day windows would create a false sense of statistical evidence. The relevant validation is instead sensitivity to anchors, units, venues, and observation windows, together with full reporting of attempted classifications and hold-outs. If a future study turns validated pairs into a predictive or trading model, out-of-sample evaluation and multiple-testing control become mandatory rather than optional.

3 Institutional Background

3.1 Builder-deployed markets make identity venue- and version-specific

HIP-3 allows a deployer to define a market's oracle and contract specifications, operate it, set leverage limits, and settle or resume an asset through `haltTrading`. Each perp DEX has independent margining, order books, and deployer settings; collateral is not universally USDC by architectural assumption (Hyperliquid, 2026b). The default API universe is therefore incomplete coverage for a study of all HIP-3 markets, and deleting a prefix such as `xyz:` is not a valid identity rule. The prefix identifies the deployment namespace whose specifications and lifecycle must be retained.

Trade[XYZ] distinguishes HIP-3 perpetuals from HIP-4 event products (trade[XYZ], 2026a). A stock perpetual does not confer a share, a vote, a dividend claim, or a redemption right. Its economic meaning comes from the reference and settlement specification, not from a stock icon or symbol. Likewise, two issuers can use the same short name for different claims. A correct mapping key therefore includes venue, deployment, native identifier, underlying definition, quote currency, collateral, and a specification version or observation timestamp.

Market status is a separate dimension. Metadata can continue to contain an entry after it is marked delisted, and a document can continue to describe a market after its API status changes. A failed public request is a third state—evidence unavailable—rather than either active or delisted. These distinctions motivate the evidence ladder and prohibit silently dropping inactive or conflicting lines from the research record.

3.2 Oracle transitions and closed external markets prevent trivial price comparisons

Trade[XYZ] uses an externally derived fair price when external inputs are available. When they are unavailable, its oracle initializes from the last external price and advances using an internal mechanism based on order-book impact prices (trade[XYZ], 2026d). Let S_t denote the oracle, and define the impact-price difference as

$$IPD_t = \max(P_t^{impactBid} - S_{t^-}, 0) - \max(S_{t^-} - P_t^{impactAsk}, 0). \tag{1}$$

The documented continuous-time update can be written as

$$S_t = \beta_t S_{t^-} + (1 - \beta_t) x_t, \qquad \beta_t = \exp(-\Delta t^*/\tau), \qquad x_t = S_{t^-} + IPD_t, \tag{2}$$

with $\tau = 30$ minutes, $\Delta t^* = \min(\Delta t, c\tau)$, and $c = 0.1$ in the documentation accessed on 4 November 2026. When external data resume, the oracle returns to the externally derived price on the next tick. This mechanism means that a weekend book can be genuine and two-sided while its reference state differs fundamentally from a weekday externally anchored book.

Discovery bounds add another state variable. The instantaneous mark-price bound is $\pm(1/L)$ around a reference price, where L is maximum leverage. Re-anchoring permits a configured number of resets in each direction before a hard cap is reached; external pricing resets the reference and counters (trade[XYZ], 2026b). The bounds are not a universal static percentage and should not be read as a guarantee against liquidation or reopening gaps.

Lighter similarly transitions from external oracles to internal impact-price pricing when data become stale and returns to external prices when available. Its documentation distinguishes 30-minute index smoothing from two-minute mark smoothing. It also records a 10 July 2026 removal of last-oracle-relative caps for a specified list of markets, including XAU, XAG, WTI, BRENTOIL, several equities, and indices (Lighter, 2026b). Copying only the generic formula above that update would incorrectly apply an old constraint to those markets. The same page notes that internal prices continue to be validated against other venues. An equivalence review must therefore record both the mechanism and its effective date.

External trading sessions require equal care. Equity, futures, metals, and energy calendars differ, and holidays are venue- and asset-specific (trade[XYZ], 2026c). Time-zone conversion should use an IANA zone such as `America/New_York`, not a permanent UTC offset, because daylight-saving transitions change the offset. A chain that trades 24/7 can consequently show prices while the external reference is closed; that observation does not prove that the external price is updating or that two venues share the same internal-price state.

3.3 Funding formulas have windows, multipliers, clamps, and dead zones

Funding is a cash flow attached to a specified position and reference, not a generic annualized spread. Hyperliquid's native crypto contracts pay funding hourly, and some are USDT-denominated but USDC-margined without a currency conversion (Hyperliquid, 2026a). HIP-3 adds deployer configuration and possible fee shares. Lighter also pays funding hourly but computes a one-hour premium from minute samples and applies market-specific parameters (Lighter, 2026a).

For Lighter, let M be the funding premium multiplier, $\bar{p}$ the average minute premium multiplied by M, I the interest-rate component, C_s the small clamp before multiplication, and C_b

the big clamp. The documented calculation is

$$s = C_s M, \tag{3}$$

$$\widetilde{p} = \bar{p} + \mathrm{clamp}(I - \bar{p}, -s, s), \tag{4}$$

$$f_{1h} = \frac{\mathrm{clamp}(\widetilde{p}, -C_b, C_b)}{8}, \tag{5}$$

where $\mathrm{clamp}(x, a, b) = \max(a, \min(b, x))$. When the premium lies within the interest-centered small-clamp region, the funding rate defaults to $I/8$, producing a dead zone rather than a continuously varying quote. The documentation lists common parameters but directs researchers to `orderBookDetails` for current market values. Raw API multiplier fields must not be reinterpreted as final hourly rates without applying the documented formula and units.

A cross-venue funding comparison must align the accrual window, settlement timestamp, sign convention, position notional, quote and collateral currencies, and market-specific multiplier. None of these steps validates the underlying mapping. They come after equivalence, just as execution costs come after quantity units. This paper therefore reports funding mechanisms as institutional background but does not annualize a funding differential.

3.4 Reference-contract rolls and synthetic indices can diverge by construction

For futures-linked perpetuals, the reference is a time-varying basket rather than a permanent spot asset. Trade[XYZ]'s roll schedule uses five daily transitions during the fifth through ninth business days, with updated weights taking effect when the oracle returns to external pricing (trade[XYZ], 2026e). The historical research snapshot recorded text differences between XYZ and Lighter schedules even when both displayed five 20% steps. Those schedules could not be treated as identical merely because the month codes matched.

The documents rechecked on 4 November 2026 make version control even more important. Lighter states that WTI, NATGAS, and XCU have used RedStone prices since 22 October 2026, so its historical rolling table no longer applies to them. BRENTOIL and GASOIL retain a five-step schedule, with the current configuration specified at 23:59 Europe/London (Lighter, 2026c). A current source can therefore invalidate a prior roll comparison without changing the historical fact that the earlier documents differed. Both the old observation and the new specification belong in the audit trail.

Synthetic equity indices create a related problem. A venue can use cash indices during regular hours and futures-derived estimates during extended hours, while another uses a proprietary futures–spot premium model. "Synthetic spot" is a family of constructions, not a unique price. Before comparing levels, researchers need the contract, source, expiry, discount or premium parameter, session, and staleness rule. The full algorithm is unnecessary only if the pair is rejected at an earlier identity stage—for example, an index level versus an ETF share.

3.5 Pre-IPO rebasing changes units without necessarily changing valuation

Pre-IPO perpetuals use an estimated share count to turn a per-share-like price into an implied valuation. Binance's OpenAI announcement lists an estimated one billion shares (Binance, 2026). Gate explains that an USD 800 billion valuation corresponds to a price of 800, which implies a one-billion-share scale (Gate, 2026). Aster documents one billion reference fully diluted shares for OPENAI, ANTHROPIC, POLYMARKET, and MOONSHOT, and 320,945,459 for OURA (Aster, 2026). The latter count is also present in Oura's S-1/A as common stock outstanding

immediately after the offering (Oura Inc., 2026).

OKX executed a platform-initiated adjustment ratio of ten for OPENAIUSDT and ANTHROPICUSDT on 30 June 2026 (OKX, 2026b). Its product documentation describes rebasing from an estimated share count to an actual disclosed count with value neutrality as the design principle (OKX, 2026a). The economically continuous quantity is therefore implied valuation, not the unadjusted native price. Contract-size adjustments and later transitions to a standard listed-equity perpetual must be treated as specification versions.

There is no common pre-IPO settlement venue. Documents describe conversion to a listed-equity reference if an IPO occurs, but cancellation, delay, or termination can involve venue-specific discretionary treatment (OKX, 2026a; Aster, 2026). Two legs need not converge if no listing occurs, and conversion helps only if both legs ultimately reference the same listed security in compatible units. This is why our classification uses "reference-basis pair" rather than implying a convergence date or a common terminal payoff.

4 Data

4.1 Sources and observation windows

The study uses official public documents and read-only public interfaces. No account was connected, no order was placed, and no private or user-level data were accessed. Table 1 summarizes the three historical modules and the separate recheck. The Hyperliquid module is a single-day snapshot dated 20 September 2026. Metadata were obtained by sending `perpDexs` and then `meta` for each named DEX and the default universe to the official `info` endpoint; collateral token indices were resolved using `spotMeta` (Hyperliquid, 2026c). Representative books came from `l2Book` using complete identifiers such as `xyz:GOLD`.

The pre-IPO module reports public catalogues and tickers fetched on 14 October 2026. Its broad inventory describes 21 perpetual venues, but the numerical analysis in this paper is deliberately narrower: five companies, 26 eligible exact venue observations, one cross-venue-median pseudo-anchor, and one conflicting Gate KIMI observation. The exact marks were observed at 08:33–08:42 UTC, within the broader 08:30–08:43 collection window. Four range summaries preserve additional venues for which the source did not preserve exact rows in the task package. An earlier 07:46–08:05 UTC window supplies rounded ratios for temporal sign checks. No observations were interpolated between the windows.

The quantity-unit module uses the SHIB-USDT example embedded in BingX's official order-book documentation (BingX, 2026). It contains five bid and five ask levels, each with an ordinary array and a Coin array. The example's embedded timestamp is 1,702,719,083,983 milliseconds. It is documentation evidence, not a current quote and not a sample from which venue-wide field semantics can be estimated.

4.2 Inclusion, exclusion, and preservation rules

The coverage counts include delisted entries in total metadata and then report non-delisted counts separately. This prevents a survivor-only denominator. Delisted or inactive pre-IPO lines identified in the source—including specified Hyperliquid `vntl:` and `para:` lines, inactive Lighter SPACEX, and an Aster SPACEX line absent from the active exchange catalogue—are not used as current executable pairs. They remain relevant as counterexamples to ticker-based mapping. A delisted stale price is not treated as a basis observation.

Table 1. Data modules, coverage, and analytical use

Module	Observation date / window	Coverage preserved	Analytical treatment
Hyperliquid coverage	2026-09-20	Default plus ten named perp DEXs; 523 metadata entries and 325 non-delisted entries	Counts by DEX and collateral; entries are not unique underlyings or approved pairs
Hyperliquid books	2026-09-20 12:54:41–12:54:43 UTC	Ten full coin identifiers with best bid and ask	Evidence of a two-sided book only; Sunday reference-state qualification
Pre-IPO exact marks	2026-10-14 08:33–08:42 UTC	Five companies; 26 eligible exact venue observations, one pseudo-anchor, one conflict	Deterministic canonical conversion; conflict retained as hold-out
Pre-IPO ranges and early window	2026-10-14 07:46–08:05 and 08:33–08:42 UTC	Four grouped ranges and six rounded early ratios	Sign and range checks only; no disaggregation or full-precision claims
BingX official example	Embedded document example; accessed 2026-11-04	SHIB-USDT, five bids and five asks, ordinary and Coin arrays	Quantity-ratio identity check; no live-market inference
Public recheck	2026-11-03 19:47–19:53 UTC	Hyperliquid: 556 metadata entries and 347 non-delisted entries	Temporal stability check only; not spliced into the historical snapshot

Sources: official public APIs and documents cited in Sections 3 and 4. All counts are reproduced from the package CSV files. The recheck uses a different date and is not a time-series observation of a fixed contract panel.

For pre-IPO point calculations, an observation is eligible when the source preserves an exact native mark, a company identifier, a venue, a unit factor or stated consensus factor, and an anchor rule. The pseudo-anchor for POLYMARKET is the supplied median of MEXC and LBank, not an additional market observation. Gate KIMI fails the unit-evidence requirement because its stated count conflicts with its worked examples; it is retained in the sensitivity analysis but excluded from the eligible residual summary. Range-only rows are excluded from point estimates but retained for sign and magnitude checks.

No winsorization, outlier deletion, or variance-stabilizing transformation is applied. Large residuals are economically informative and remain visible. Price-consensus units are not relabelled as documented. For these rows, a factor of one is a working normalization supported by nearby prices, not proof that no alternative contract definition exists. The analysis records this lower evidence tier and avoids claiming that the unit is "unambiguous."

Negative mappings are preserved rather than treated as missing. SP500 index points are not SPY shares; WTI is not Brent; a gold token is not automatically a gold perpetual; and USD/JPY is the reciprocal, not the negative, of JPY/USD. Single-venue names such as Gate ANDURIL, KALSHI, and NEURALINK cannot form a cross-venue pair in the snapshot. They are hold-outs, not zero-residual observations.

4.3 Missing evidence and inaccessible interfaces

Missingness is coded by cause. Bybit's pre-IPO catalogue was unreachable from the original measurement host with HTTP 403, so its coverage is unknown rather than empty. Deribit's support page and index constituents were likewise unread in the source study. Several venues had no methodology page establishing a share count; their factor-one status remains price-consensus evidence. The Gate KIMI conflict is not treated as missing and filled by the best-fitting factor.

A recheck was attempted from the writing environment on 3–4 November 2026. Hyperliquid, Lighter, and Gate returned usable public responses. Binance, OKX, KuCoin, MEXC, Aster, Bybit, and BingX were unavailable from that host through connection failures or HTTP 403 at the time of the request. No challenge was bypassed. These failures do not reverse the historical observations, but they prevent a full current revalidation of every venue and are disclosed as a publication limitation.

The task package preserved numerical observations but not the original historical response bytes. The reproduction package therefore distinguishes *computational reproduction* from *raw-response authentication*. Every point estimate in the paper can be recomputed from the preserved CSV observations and code. The historical bytes themselves cannot be independently hashed from the package. For BingX in particular, previously summarized BTC and RSR depth-walk results are not used as reproduced findings because their raw responses were not supplied; the official SHIB example provides the primary unit case instead.

4.4 Descriptive coverage and temporal recheck

The default Hyperliquid universe accounts for 234 of the 523 historical metadata entries and 178 of the 325 non-delisted entries. The named DEXs collectively account for the remaining entries, with `xyz`, `para`, `mkts`, and `io` retaining non-delisted markets in the snapshot. Several historical DEXs have entries but zero non-delisted markets. Collateral-token resolution shows USDC in the surviving groups discussed here, while delisted groups also include USDH, USDE, and USDT0. The architecture therefore cannot be summarized as "all HIP-3 markets use USDC."

The fresh Hyperliquid recheck returned 556 metadata entries and 347 non-delisted entries,

compared with 523 and 325 in the historical snapshot. These counts demonstrate that catalogue composition changes over time; they do not provide a matched-entry survival rate because the recheck has not been converted into a versioned contract panel. The current count is used only to reinforce the need for observation dates and lifecycle versions. Current Lighter documentation separately records changes to reference sources and price-cap arrangements, as discussed in Section 3. Historical and current evidence are never blended into a single synthetic time series.

5 Methods

5.1 Evidence levels are not interchangeable

Figure 1 defines the evidence ladder used throughout the paper. "Listed in documents" records an official specification or announcement. "Present in metadata" requires a current public catalogue entry. "Not delisted/active" uses the venue's status field at the observation time. "Two-sided book observed" requires both bid and ask arrays in a preserved response. "Account allowed" is account- and jurisdiction-specific and is not tested here. "Economic equivalence reviewed" requires underlying identity, units, reference, settlement, and corporate-action checks. "Size executable" additionally requires a stated notional, synchronized books, complete costs, and permission to trade both legs.

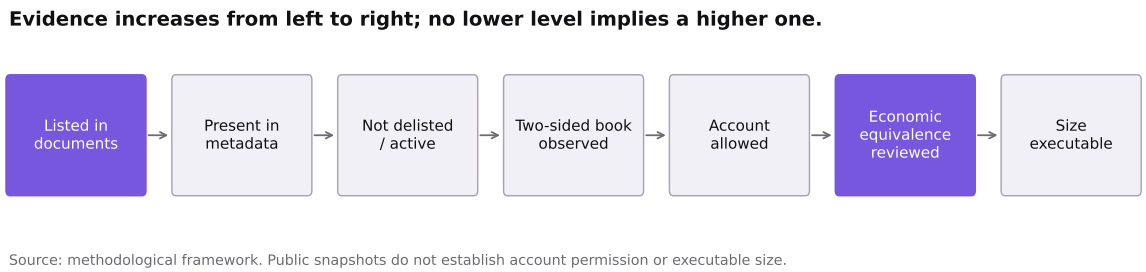


Figure 1. Seven evidence levels for cross-venue instrument comparison. A lower-level observation does not establish a higher-level claim. Source: methodological construction; public snapshots in this study reach neither account permission nor executable size.

The ordering is a reporting discipline rather than a claim that every research workflow must occur in exactly this sequence. Economic review can begin before an account check, but a positive account check cannot substitute for it. Evidence also has a timestamp and a specification version. A contract that passed identity review before a rebase, roll-source change, or settlement conversion must be reviewed again.

5.2 Canonical value and quote-currency conversion

Let i denote a pre-IPO perpetual on venue $v(i)$ at time t_i. Its native mark is P_i in quote currency Q_i, its documented assumed share count is N_i shares, and X_i is USD per unit of Q_i. The canonical implied equity value in USD billions is

$$V_i = P_i \frac{N_i}{10^9} X_i. \tag{6}$$

In the source study's descriptive convention, $X_i = 1$ for USD, USDT, USDC, and USD1 lines. This is a numéraire assumption, not a claim that stablecoins are risk-free or always trade at one

USD. Residuals can therefore include a stablecoin or quote-currency basis. HKD-denominated lines are not silently assigned $X_i = 1$; they require an explicit FX conversion and belong to a separate quote-currency class until then.

The transformation in Equation (6) is meaningful only if N_i is the share count embedded in the venue's price definition. A token index, API contract multiplier, minimum order size, and quantity precision do not substitute for N_i. If the count is undocumented, a factor inferred from nearby prices is labelled price-consensus evidence rather than documented evidence. If documents conflict, V_i is reported under alternative scenarios and the pair is held out.

For two native pricing units A and B that reference the same valuation and have no additional contract multiplier, the reference-share hedge ratio is

$$q_B = q_A \frac{N_B}{N_A}, \tag{7}$$

where q_i is the number of native pricing units. This follows by equating valuation sensitivities $q_A/N_A = q_B/N_B$. Thus, one Binance OPENAI unit corresponds to ten OKX OPENAI units after OKX's tenfold rebase, before allowing for residual basis. If native order quantities are contracts with multipliers κ_A and κ_B, the executable ratio becomes $q_B = q_A(\kappa_A/\kappa_B)(N_B/N_A)$. Minimum quantities, rounding, margin, and residual price differences must be handled separately; Equation (7) is not an order instruction.

5.3 Residual definitions and anchors

For company c, let $A(c)$ be an explicitly stated anchor observation. The arithmetic and log residuals for observation i are

$$r_i = 100\left(\frac{V_i}{V_{A(c)}} - 1\right), \qquad r_i^{\log} = 100\log\left(\frac{V_i}{V_{A(c)}}\right). \tag{8}$$

Binance is the source-study anchor where available. POLYMARKET uses the supplied median of MEXC and LBank, USD 23.85 billion, because Binance does not provide an anchor in the exact table. The median is a pseudo-observation and is not counted as a venue in sample totals.

Anchor choice is a measurement decision, not a claim that Binance is fair value. The robustness section recomputes residuals against the median of actual eligible venue values for each company and compares arithmetic and log forms. Log residuals are additive under inversion and more symmetric for multiplicative discrepancies, while arithmetic residuals are easier to interpret as percentage deviations from the anchor. Neither is a forecast error or a return.

The raw native-price ratio is retained when it reveals a unit problem. For OKX, we compute $\rho_i = P_{OKX,i}/P_{Binance,i}$ before conversion and compare it with r_i after conversion. A documented scale factor predicts $\rho_i \approx N_A/N_B$ up to reference basis. If that prediction fails materially, the unit explanation is not accepted merely because it is convenient.

5.4 Classification rules

The classification is deterministic at the conceptual level and does not embed a production threshold or an automated trading rule.

1. **Exact conversion.** The instruments reference the same economic claim and compatible settlement, the unit transformation is documented, and any residual reflects measurement or timing rather than an independent reference definition. A venue whose index is explicitly

constructed from another venue's converted price may satisfy the reference condition, but its own mark still needs checking.

2. **Reference-basis pair.** Company or asset identity is shared, but reference providers, internal pricing rules, collateral, quote currency, roll construction, or terminal settlement differ. The residual is a reference basis, with no convergence claim.

3. **Hold out.** Identity or status conflicts, the share count is internally inconsistent, only one venue exists, the interface is unavailable, or executable quantity semantics are unknown. A close price fit cannot override this decision.

This classification distinguishes documented units from price-consensus units and exact observations from range summaries. It also records rejected negative controls. There is no universal residual cutoff for "close enough," because such a cutoff would depend on the research purpose, timestamp synchronization, volatility, costs, and account permissions. Reporting a threshold without those inputs would create false precision.

5.5 Quantity-aware depth walking

For price levels $j = 1, \ldots, J$, let p_j be the ask price and a_j the explicitly documented base-coin quantity. For a quote budget B, the filled base quantity at level j is

$$x_j = \min\left(a_j, \frac{R_{j-1}}{p_j}\right), \qquad R_j = R_{j-1} - p_j x_j, \qquad R_0 = B. \tag{9}$$

If the budget is exhausted, the volume-weighted average price is $\text{VWAP} = B/\sum_j x_j$; if it is not exhausted, the result is unavailable at that book depth rather than extrapolated. The same logic applies to bids with the appropriate execution direction. A counterfactual transformation $a'_j = sa_j$ can be used to quantify a hypothesized unit error, but s must not be identified with contract size unless the documentation says so.

This paper applies the formula conceptually because the original BTC and RSR response bytes were not supplied. The official SHIB example instead provides a direct algebraic check $a_j/q_j = 1{,}000$ for every displayed level, where q_j is ordinary quantity and a_j is Coin quantity. That ratio is not generalized to other symbols or API versions.

5.6 Identification assumptions and counterfactual interpretation

Four assumptions determine what the residuals mean. First, the underlying and share-count definitions are stable within each observation window; an intraday rebase would require separate versions. Second, documented N_i is correct and consistently interpreted; failure produces a hold-out or scenario analysis. Third, timestamps are close enough that market movement is second-order for the descriptive comparison; asynchronous marks would add a time basis. Fourth, $X_i = 1$ is adequate for the source convention; failure adds quote-currency basis.

The primary counterfactual is not a trading return but a measurement error. If a researcher treats OKX's post-rebase price as factor one, the native-price ratio appears close to 0.10. If the researcher also uses one-to-one native units instead of Equation (7), the position covers one-tenth of the intended reference-share exposure. For USD 10,000 of intended exposure, that is USD 9,000 left unhedged before residual reference risk. This exposure calculation is deterministic and is not annualized.

No causal effect, forecast, variance, confidence interval, or probability of convergence is es-

timated. The available evidence identifies algebraic consequences of unit choices and describes residuals under stated anchors. Stronger claims require a versioned multi-day panel with synchronized quotes, raw-response hashes, complete contract specifications, and execution evidence.

5.7 Symbol table

Table 2. Symbols and definitions

Symbol	Definition
$i, v(i), t_i$	Instrument, its venue, and its observation timestamp
P_i	Native mark price of instrument i
Q_i	Native quote currency of instrument i
N_i	Assumed reference share count, in shares
X_i	USD per unit of Q_i; set to one under the source's USD/USDT/USDC/USD1 convention
V_i	Canonical implied equity value in USD billions
$A(c)$	Stated anchor for company c
$r_i, r_i^{\log}$	Arithmetic and log residuals in percent
q_i	Native pricing-unit quantity; not necessarily an exchange order-contract quantity
κ_i	Optional order-contract multiplier, required only for executable sizing
p_j, a_j, x_j	Order-book price, documented base quantity, and filled base quantity at level j
B, R_j	Initial quote budget and remaining quote budget after level j
M, I, C_s, C_b	Lighter funding multiplier, interest, small clamp, and big clamp

6 Results

6.1 Public coverage stops well short of executable equivalence

Table 3 answers what the Hyperliquid enumeration establishes. The default universe is only part of public coverage. The named `xyz`, `para`, `mkts`, and `io` groups contain 108, 27, four, and eight non-delisted entries, respectively. Six other named groups contain 97 metadata entries but no non-delisted entries. Across all groups, 325 of 523 entries are not delisted, or 62.141491%. This proportion is a metadata statistic, not a coverage probability for unique assets.

Table 3. What does the 20 September 2026 Hyperliquid enumeration establish?

DEX group	Metadata	Non-delisted	Collateral in surviving group	Evidence limit
Default	234	178	USDC	No unique-underlying count
`xyz`	123	108	USDC	Books observed for selected symbols
`para`	35	27	USDC	Identity review still required
`mkts`	24	4	USDC	Units not inferred from names
`io`	10	8	USDC	Independent deployer
Six fully delisted named groups	97	0	USDH, USDE, USDT0	Retained as lifecycle evidence
Total	**523**	**325**	Not universal	Not an executable-pair count

The six fully delisted groups are `flx`, `vntl`, `hyna`, `km`, `cash`, and `abcd`. Collateral symbols are resolved from public token metadata. Full rows appear in Appendix Table 7. Source: official `info` snapshot, 2026-09-20.

Ten representative instruments returned two-sided books at approximately 12:54:41–12:54:43 UTC. These include `xyz:NVDA`, `xyz:AAPL`, `xyz:GOLD`, `xyz:SILVER`, `xyz:CL`, `xyz:BRENTOIL`, `xyz:EUR`, `xyz:JPY`, `xyz:SP500`, and `mkts:USTECH`. The observation supports only the fourth evidence level in Figure 1. It does not show that an account is eligible, that displayed depth covers a

stated order, or that two books refer to the same economic claim.

The Sunday timing matters. External equity and commodity sessions can be closed while the chain continues to display quotes. `mkts:USTECH` illustrates the residual identity problem: its bid and ask were 717.77 and 717.89, but a familiar technology-index name does not establish its scaling or underlying definition. We therefore do not equate it with US100 or QQQ. Appendix Table 8 preserves the complete bid–ask sample rather than converting it into provisional pairs.

6.2 The documented OKX factor explains most of the raw gap

Table 4 isolates the unit counterfactual. OKX OPENAI is 165.67 in native units while Binance is 1,637.90; OKX ANTHROPIC is 206.21 while Binance is 2,000.40. Without the share-count state, these look like radically different valuations. After applying $N_{OKX} = 10$ billion and $N_{Binance} = 1$ billion, the canonical values are close but not identical.

Table 4. OKX unit conversion: raw ratio versus canonical residual

Company	OKX native	Binance native	Raw OKX/Binance	Canonical OKX (USD bn)	Residual (%)
OPENAI	165.67	1,637.90	0.10114781	1,656.70	+1.147811
ANTHROPIC	206.21	2,000.40	0.10308438	2,062.10	+3.084383

Canonical values use the documented factor of ten for OKX and one for Binance. The residual is $100(V_{OKX}/V_{Binance} - 1)$. Source: official public marks, 2026-10-14 08:33–08:42 UTC; OKX rebase documentation (OKX, 2026b).

Figure 2 makes the two interpretations visually distinct. The raw ratios are approximately ten percent, whereas normalized values are 101.15% and 103.08% of the Binance anchors. The result does not mean that OKX and Binance are equivalent in every respect. Their marks and settlement remain venue-specific, and the residuals are reference bases. It means that the factor-of-ten scale is the dominant explanation for the raw gap.

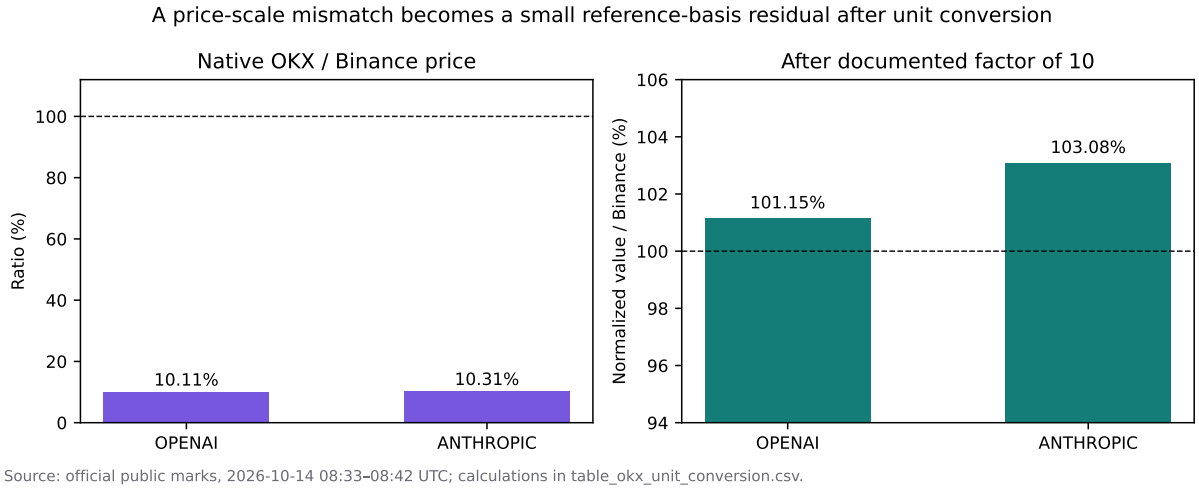


Figure 2. Native OKX/Binance price ratios before and after documented unit normalization. Source: official public marks on 2026-10-14; calculations in `table_okx_unit_conversion.csv`.

The position-size consequence is mechanical. Equation (7) requires ten OKX native units for one Binance native unit in the simplified reference-share setting. Using one of each would leave 90% of intended valuation exposure unhedged. At USD 10,000 of intended exposure, the unmatched component is USD 9,000 before any P&L from the small normalized residual. This

is a unit-error exposure, not an arbitrage loss estimate.

6.3 Remaining residuals vary with the reference construction

Figure 3 plots all non-anchor eligible exact residuals except the conflicting Gate KIMI line. The distribution is not centered on one alternative scale. It includes small positive and negative values as well as economically large deviations. The median absolute nonzero residual is 0.979323%, while the maximum is 9.530398%. These descriptive statistics summarize only the 26 exact eligible observations and should not be extrapolated to the broad 21-venue inventory.

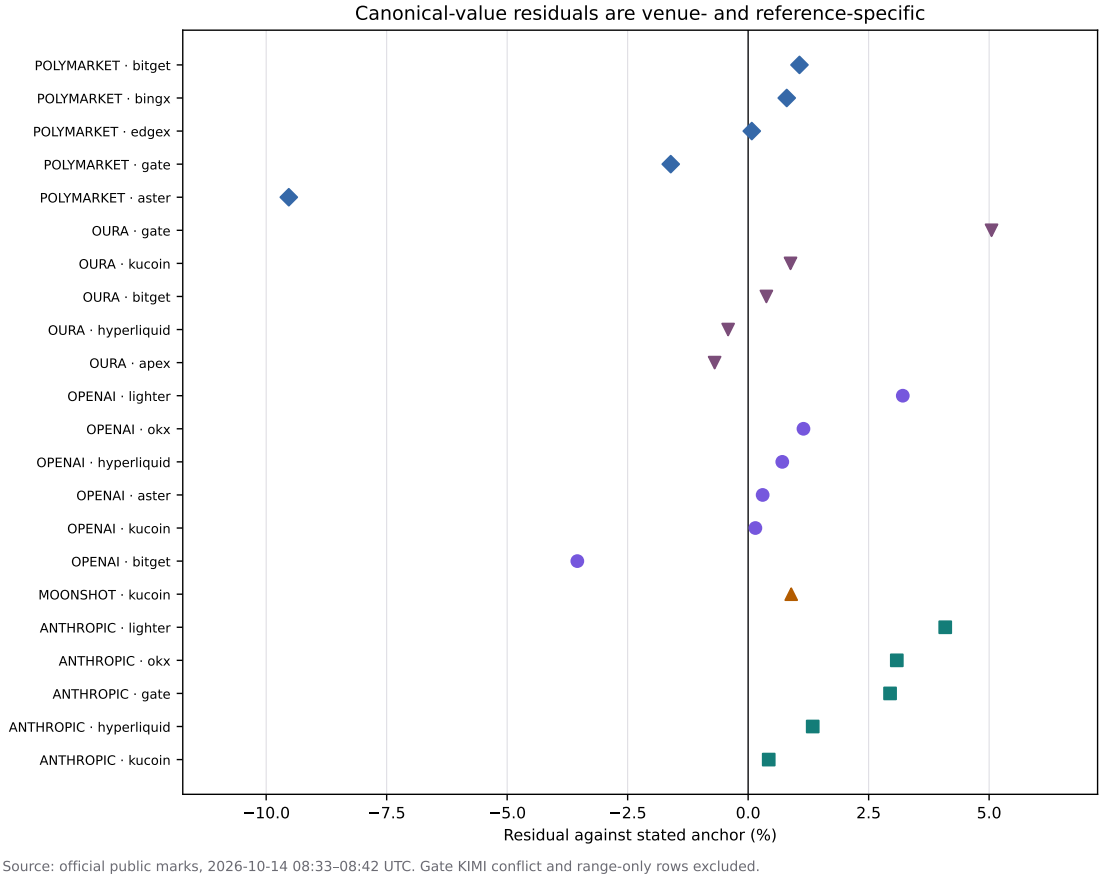


Figure 3. Canonical-value residuals against the stated anchors. Source: exact public marks on 2026-10-14 08:33–08:42 UTC. Gate KIMI, pseudo-anchors, delisted lines, and range-only summaries are excluded from this point plot.

Table 5 highlights the observations with distinct reference interpretations. Bitget OPENAI is −3.544172% and uses an own-book synthetic index in the source evidence. Lighter, which documents an internal oracle for these pre-IPO markets, is +3.208377% for OPENAI and +4.089182% for ANTHROPIC. Gate is +2.944411% for ANTHROPIC and +5.050092% for OURA. Aster POLYMARKET is −9.530398% and refers to tokenized pre-IPO exposure rather than an identical reference construction. These are reference-basis pairs, not alternative unit factors waiting to be fitted.

Small residuals can support an exact-conversion review when the reference is explicitly shared. KuCoin reports Binance-derived source fields, MEXC exposes index-origin fields, and Aster's OPENAI and ANTHROPIC index prices match Binance's marks to the cent in the source snapshot. However, an index link does not prove that the venue's own mark, settlement, or account eligibility is identical. OKX itself converts Binance and Gate inputs in its index while its mark

Table 5. Selected normalized residuals and their economic interpretation

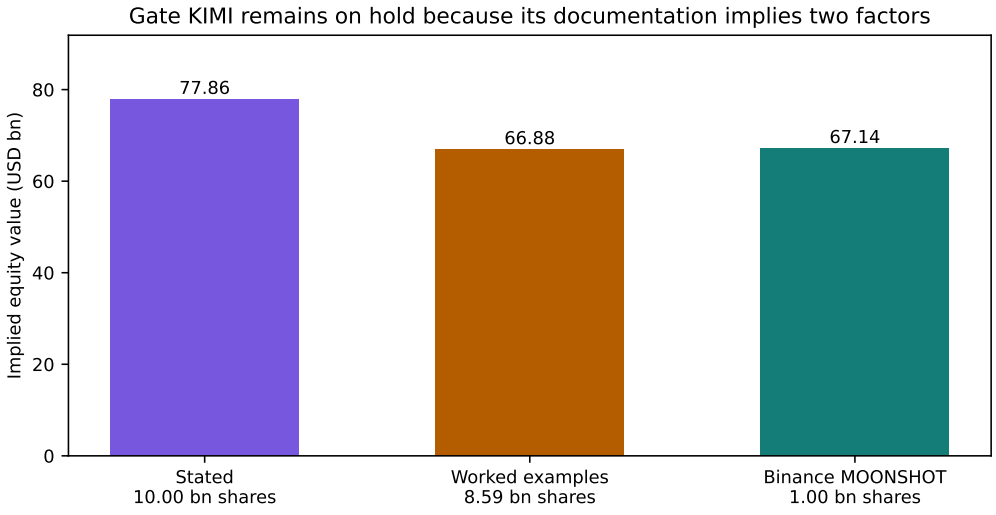
Company	Venue	Canonical value (USD bn)	Residual (%)	Reference interpretation
OPENAI	Bitget	1,579.850	−3.544172	Own-book synthetic index
OPENAI	Lighter	1,690.450	+3.208377	Internal oracle
OPENAI	Hyperliquid `io:OAI`	1,649.500	+0.708224	Deployer oracle; unit by price consensus
ANTHROPIC	OKX	2,062.100	+3.084383	Own trades after tenfold conversion
ANTHROPIC	Gate	2,059.300	+2.944411	Venue reference
ANTHROPIC	Lighter	2,082.200	+4.089182	Internal oracle
POLYMARKET	Aster	21.577	−9.530398	Tokenized pre-IPO exposure
OURA	Gate	16.490178	+5.050092	Venue reference
OURA	Hyperliquid `xyz:OURA`	15.632290	−0.415048	One-share specification

Anchors: Binance where listed; the supplied MEXC/LBank median for POLYMARKET. Values use the stated share-count factors and $X_i = 1$ under the source convention. Source: exact public marks on 2026-10-14.

is trade-driven; it therefore remains a reference-basis pair rather than an exact-conversion exemplar.

6.4 Conflicts and negative controls remain hold-outs

Gate KIMI is the clearest internal-specification conflict. Its stated factor produces USD 77.86 billion, 15.966637% above Binance MOONSHOT. The mean worked-example factor of 8.59024827 billion produces USD 66.883673 billion, only 0.381780% below it. Figure 4 shows both outcomes. The close second result does not justify selecting 8.59 as the true factor, because doing so would use the target price to choose the unit. The appropriate classification is hold-out pending clarification or a rebase notice.



Source: Gate announcement 100758 and public marks on 2026-10-14; sensitivity calculations in reproduction output.

Figure 4. Gate KIMI normalization under the stated and worked-example share counts. Source: Gate announcement 100758 and exact marks on 2026-10-14.

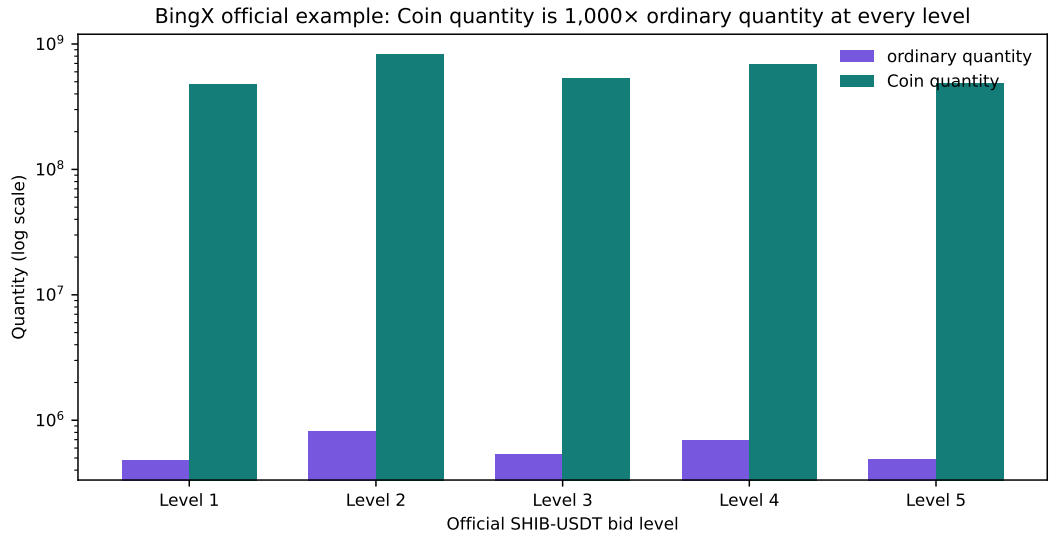
Other hold-outs have different causes. Delisted or inactive lines are lifecycle exclusions, even if a stale price remains visible. Single-venue ANDURIL, KALSHI, and NEURALINK have no cross-venue counterpart. Bybit coverage is unknown because the catalogue was inaccessible.

Methodology-undocumented venues remain at price-consensus unit evidence even when their values are close.

Negative controls prevent false positives. SP500 index points and SPY ETF shares are related but not one-to-one instruments. WTI and Brent can be studied as a relative-value pair but not as the same asset. A spot gold token or custody claim is not automatically a gold perpetual. USD/JPY must be reciprocated before comparison with JPY/USD; subtracting the native quotes would be meaningless. Finally, the source's SPCX and QNTX examples show that a venue's pre-market flag or legacy symbol can lag a conversion to listed equity, so flags and names cannot decide class membership alone.

6.5 Quantity fields can change depth by orders of magnitude

BingX documents ordinary `bids/asks` only as containing "quantity," while `bidsCoin/asksCoin` explicitly contain coin quantity. In the official SHIB-USDT example, the best bid price is 0.000009854, ordinary quantity is 483,909, and Coin quantity is 483,909,000. The ratio is exactly 1,000 at all five bid and all five ask levels. Figure 5 displays the bid-side values on a logarithmic scale.



Source: BingX official API documentation, embedded SHIB-USDT example (document timestamp 1702719083983 ms).

Figure 5. Ordinary and Coin quantities in BingX's official SHIB-USDT example. Source: BingX official API documentation; embedded example, not a live quote.

At that best bid, price times ordinary quantity is 4.768439 USDT if the ordinary value is counterfactually treated as base SHIB, whereas price times Coin quantity is 4,768.439286 USDT. The comparison illustrates the scale of possible depth mismeasurement; it does not establish the ordinary field's unit. The documentation does not instruct researchers to multiply ordinary quantity by contract size. Consequently, a depth walk should prefer the explicitly denominated Coin array and should not infer a universal conversion for other symbols or versions.

Contract size, quantity precision, minimum order quantity, and name prefixes such as 1000SHIB are separate concepts. A factor that corrects one need not correct the others. This result complements the pre-IPO analysis: price normalization and quantity normalization are distinct stages, and both must precede execution or cost comparisons.

7 Robustness

7.1 Early and late windows support scale interpretation, not persistence

The first check compares the rounded early-window ratios with exact later-window calculations. OKX's OPENAI raw ratio is 0.101 early and 0.10114781 later; ANTHROPIC is 0.103 and 0.10308438. Lighter's normalized OPENAI ratio is 1.033 early and 1.03208377 later, while ANTHROPIC is 1.042 and 1.04089182. Aster POLYMARKET is 0.900 and 0.90469602, and Gate OURA is 1.049 and 1.05050092. All six pairs lie on the same side of one in both windows.

These repetitions support the interpretation that the observed gaps are not a single-print artifact. They do not establish persistence beyond the same-hour windows, and they certainly do not identify mean reversion. The early values are rounded summaries, so we do not calculate confidence intervals or treat small early–late differences as precise changes. A multi-day panel remains necessary.

7.2 Venue groups and company hold-outs

The second check separates documented and price-consensus unit evidence. The documented group contains 18 eligible exact observations, with median absolute residual 0.886410% and maximum 9.530398%. The price-consensus group contains eight, with median 0.754531% and maximum 1.605870%. The smaller maximum in the latter group is not evidence that consensus units are more reliable: its composition differs, and consensus is precisely why these rows remain at a lower evidence tier. Group membership is reported rather than used to promote a factor to documentation.

Leaving out one company at a time shows how the largest residual depends on composition. Excluding OPENAI, ANTHROPIC, MOONSHOT, or OURA individually leaves Aster POLYMARKET's 9.530398% as the maximum. Excluding POLYMARKET reduces the maximum to Gate OURA's 5.050092%. Thus, the largest number is company-specific, but the conclusion that material reference bases remain after normalization does not rest on a single company. MOONSHOT has only two eligible exact observations after Gate KIMI is held out, so its alternative median is descriptive and receives little weight.

Reference-provider grouping gives a more substantive check than a unit-only split. Explicitly linked references, such as KuCoin's Binance-derived fields, belong to the exact-conversion review path. Independent internal or trade-driven references remain reference-basis pairs even when residuals are small. The classification therefore does not reverse when an anchor or a company subsample changes.

7.3 Share-count sensitivity is decisive for Gate KIMI

The third check varies the share count rather than hiding the conflict. Gate's three worked examples imply 8.59598854, 8.58369099, and 8.59106529 billion shares. Their mean is 8.59024827 billion. The stated ten-billion assumption gives a +15.966637% residual; the mean example assumption gives −0.381780%. The individual examples produce nearly the same implied count, so the conflict is not caused by selecting one unusually rounded example. It is a direct statement-versus-example inconsistency.

OKX provides the opposite kind of sensitivity. Factor one leaves an approximately 90% raw gap, whereas the documented factor ten reduces the residuals to +1.147811% and +3.084383%. The unit decision is therefore consequential in both cases, but only OKX has unambiguous

documentary support. These two examples show why parameter sensitivity must be combined with evidence tiers rather than used as a free calibration exercise.

7.4 Alternative anchors and residual forms preserve the classification

The fourth check replaces the source anchor with the median of actual eligible venue values for each company and reports both arithmetic and log residuals. The median anchors are USD 1,642.830 billion for OPENAI, 2,043.250 billion for ANTHROPIC, 23.868 billion for POLYMARKET, and 15.727130 billion for OURA. Under this alternative, OKX's OPENAI residual is +0.844275% and ANTHROPIC is +0.922550%; Lighter's are +2.898657% and +1.906277%. Aster POLYMARKET is −9.598626%, and Gate OURA is +4.851793%. Residual magnitudes change, but the unit-versus-reference interpretation and all hold-outs remain intact.

Log residuals tell the same algebraic story without privileging the anchor direction. For example, Aster POLYMARKET has arithmetic residual −9.530398% and log residual −10.015628% under the supplied anchor. OKX OPENAI has +1.147811% arithmetic and +1.141274% log. These differences are expected from the transformation, not model disagreement.

Table 6. Robustness checks and what they do—and do not—establish

Check	Result	Permitted interpretation
Early versus late window	Six rounded early ratios have the same sign relative to one as exact later ratios	Repetition within two same-day windows; no long-run persistence or convergence
Unit-evidence groups	Documented: 18 observations, median absolute residual 0.886410%; consensus: eight, 0.754531%	Different evidence tiers and composition; consensus is not documentation
Company hold-outs	Maximum remains 9.530398% unless POLYMARKET is removed; then 5.050092%	Material bases are not unique to one company, but maxima are composition-dependent
Share-count sensitivity	Gate KIMI: +15.966637% at ten billion versus −0.381780% at 8.59024827 billion	Conflict is decisive; best-fitting factor is not selected
Alternative anchors and log form	OKX median-anchor residuals: +0.844275% and +0.922550%; main signs and classifications retained	Anchor and functional-form sensitivity; no statistical equivalence test

All values are reproduced by `scripts/reproduce.py`. The robustness family is fully reported; no unsuccessful variant was hidden and no variant was selected to optimize a return.

Table 6 summarizes the checks. No statistical-significance claim accompanies them because the unit of analysis is a descriptive snapshot, not a time series of forecast errors. No transaction-cost-adjusted return is reported because neither execution nor a holding-period path is observed. The appropriate conclusion is narrower: the documented OKX conversion, the persistence of several same-hour reference bases, and the Gate KIMI hold-out are not artifacts of one anchor, one company, or one residual formula.

8 Limitations

8.1 The evidence is cross-sectional and incomplete

The most important limitation is temporal. Hyperliquid coverage is a single-day snapshot, and the pre-IPO numerical evidence consists of two same-day windows. There is no multi-day contract panel, no estimate of mapping-transition frequency, and no time-series model of basis dynamics. Two windows cannot establish stationarity, long-run persistence, convergence, half-life, or mean reversion. The paper's findings are algebraic and descriptive; they are not out-of-sample results.

The sample is also incomplete in coverage. The broad pre-IPO inventory includes 21 venues, but exact point calculations use only the preserved observations described in Table 1. Bybit coverage is unknown because the public interface returned HTTP 403 in the original study. Deribit constituents and several methodology pages were unavailable. The current writing-environment recheck reproduced only part of the public evidence, with connection failures or 403 responses for several venues. An inaccessible interface is not evidence of absence, and a current response is not retrospective proof of the historical state.

Before public posting, the release owner should repeat the complete official-document and public-interface review on the actual publication date. New listings, delistings, rebases, oracle-source changes, and settlement conversions should be versioned rather than silently overwriting the snapshots studied here. If a material term has changed, the historical result should be labeled with its original date and the current status should be reported separately.

8.2 Raw-response authentication is not available for every historical observation

The task package preserved de-identified numerical observations and document summaries but not all original response bytes. The reproduction package can recompute every reported point estimate, residual, ratio, and figure from those preserved values. It cannot independently prove that the historical bytes have not been altered or mis-transcribed. This is weaker than a full raw-data replication standard.

The missing raw responses matter most for the BingX BTC and RSR depth-walk summaries. We therefore exclude those earlier VWAP values from the paper's reproduced results. The official SHIB example establishes a documentation-level 1,000-fold field difference, but it does not replace a fresh multi-symbol depth study. A future package should preserve exact response bytes, UTC request times, schemas, and SHA-256 manifests at collection, together with fixed-decimal execution scripts.

8.3 Anchors, marks, and quote currencies affect measurement

Binance is an anchor, not a fair-value benchmark. The alternative-median analysis changes residual levels, and asynchronous marks can include ordinary price movement. Pre-IPO values are venue marks, not synchronized executable bid–ask pairs. No spread, depth, fees, market impact, transfer cost, borrow cost, or liquidation path is measured. Consequently, residuals cannot be interpreted as available arbitrage returns.

The canonical formula sets $X_i = 1$ for USD, USDT, USDC, and USD1 under the source convention. This simplifies comparison but can import stablecoin and quote-currency bases. Collateral and quanto treatment are not comprehensively observed across all venues. HKD lines require an explicit FX conversion, and USDC-versus-USDT collateral differences can matter for funding and margin even when native prices look aligned.

Price-consensus factors are another limitation. Nearby prices can make factor one plausible,

but they do not reveal a hidden contract definition or protect against a future rebase. Similarly, Gate KIMI's example-implied factor fits the MOONSHOT anchor closely, but choosing that factor would be endogenous. We retain uncertainty rather than reporting false unit precision.

8.4 Equivalence does not imply permission, liquidity, or common settlement

No account was connected and no jurisdictional eligibility test was conducted. Public API access can coexist with account, regional, or product restrictions. Lighter's terms, for example, reserve rights to limit or suspend API functionality and describe service restrictions (Lighter, 2026d). A two-sided book is not a promise of fills at the displayed price or size.

Economic equivalence itself is conditional. Equity perpetuals do not provide ownership rights, and corporate actions can be handled differently. Oil references can roll at different times or switch providers. Internal oracles can transition independently when external markets close. Pre-IPO contracts can be delayed, terminated, or converted on venue-specific terms. A shared company name therefore does not guarantee a common terminal payoff, particularly if an IPO is cancelled.

8.5 Documentation changes and conflicts are intrinsic to the research object

Official documents are authoritative evidence about stated rules, but they are versioned artifacts rather than immutable truths. Lighter's later documentation changes the treatment of some RWA price caps and reference sources. Trade[XYZ]'s pre-IPO specifications distinguish live and converted markets and warn that listing dates and public-filing figures can change. A static literature review cannot eliminate this version risk.

Conflicts can also occur within one document. Gate KIMI states ten billion shares while its worked examples imply approximately 8.59 billion. Resolving the contradiction requires a venue clarification or a later formal notice, not an econometric estimate. Likewise, a legacy symbol, a pre-market flag, or an old test-network page can conflict with current mainnet status. The appropriate response is a hold-out with preserved evidence, not an AI- or price-based adjudication.

8.6 Required next steps

A decisive extension should collect a multi-day, versioned panel for every candidate pair. It should retain delisted and unavailable states, timestamp native marks and books, hash raw responses, record quote and collateral currencies, and update reference-contract weights. Coverage and mapping-consistency statistics should be reported by asset class and venue, with all attempted mappings disclosed. Only after that panel exists should researchers estimate basis dynamics or evaluate predictive rules. At that stage, rolling out-of-sample tests, dependence-aware inference, multiple-testing control, and complete execution-cost accounting become necessary. Account-level permission and small-notional execution pilots, where lawful and authorized, are required for the final two evidence levels.

9 Conclusion

Similar tickers are not sufficient evidence of instrument equivalence. A defensible cross-venue comparison must establish what is listed, what is active, what is quoted, what the account may trade, what economic claim each contract represents, and whether the intended size is executable. These are different evidentiary states, and a public two-sided book establishes only one of them.

Units must be normalized before residuals are interpreted. The OKX examples show that a roughly tenfold native-price difference can be primarily a documented rescaling. The remaining residuals across pre-IPO venues show that normalized prices can still differ materially because reference prices are constructed independently. The appropriate object for those pairs is a reference basis, not a unit correction and not a presumed convergence trade.

A disciplined hold-out is as important as a successful mapping. Gate KIMI's conflicting share counts cannot be resolved by selecting the factor that best fits another venue. Index points are not ETF shares, WTI is not Brent, and reciprocal FX quotes require inversion. Quantity semantics deserve the same caution: BingX's official example demonstrates that two quantity arrays can differ by 1,000 without the ordinary field's unit being fully specified.

The resulting framework is an upstream complement to executable net-carry research. It supplies the identity and unit checks required before funding, execution costs, margin risks, or returns are calculated. It does not claim arbitrage feasibility, statistical significance, convergence, or investment performance. Its evidence is intentionally bounded by snapshot dates, interface accessibility, documentary consistency, and the preserved raw material.

Future work should apply the framework to a versioned multi-day panel with raw-response hashes and synchronized books. That extension would convert the present descriptive taxonomy into a basis for time-series inference and, eventually, lawfully authorized execution studies. Until then, the defensible conclusion is methodological: verify identity first, normalize units second, classify reference differences third, and leave unsupported equivalence unresolved.

Declarations

Conflict of interest

The author declares no competing financial or non-financial interests directly related to the empirical measurements reported in this paper. Venue names are used solely to identify public instruments and documents; no venue sponsored or reviewed the study before submission.

Funding

No external funding was reported for this work. The research was completed by humarb Research.

Data availability

The historical evidence consists of de-identified observations from official public read-only interfaces and official public documents, with snapshot dates and access limitations reported in Section 4. The accompanying reproduction package contains the preserved numerical observations, range summaries, data dictionary, and current Hyperliquid coverage check. Original historical response bytes were not included in the task package and this limitation is explicit. Public documents are cited with URLs and access dates. No private account, user-level, or undisclosed interface data are included.

Code availability

The reproduction package includes `scripts/reproduce.py`, which recomputes all point estimates, robustness tables, and five vector figures without network access, and `scripts/fetch_public.py`, which documents an optional refresh of selected official public endpoints. The README provides environment, runtime, lineage, output-map, and evidence-boundary infor-

mation. No production trading code, internal pipeline, private parameter, or account credential is included.

Generative AI use

Generative AI was used to assist literature discovery, language editing, code drafting, and internal consistency review. Official documents, DOI metadata, numerical calculations, and references were checked using the sources and reproduction scripts described in the paper. The author is responsible for the final text, analysis, and all remaining errors. AI assistance was not used to invent unavailable observations or to resolve documentary conflicts without evidence.

Preprint and peer-review status

This is a prepublication methodological and descriptive empirical manuscript. It has not undergone external peer review. The multi-role internal review described in the handoff materials is not external peer review. A Zenodo concept DOI and version DOI should be inserted only after they are assigned.

Non-investment-advice statement

This paper is for research and methodological information only. It does not constitute investment, legal, tax, or trading advice, a solicitation, or a recommendation to use any venue or instrument. Public quotes and marks are historical observations. Nothing in the paper guarantees execution, convergence, capital safety, or investment returns.

Acknowledgment and research-channel identification

This paper was completed by humarb Research, London, United Kingdom. humarb.ai is its research channel, and the research index is available at https://humarb.ai/research/. The author thanks reviewers of the numerical, methodological, documentary, and reproduction checks for their role-specific comments; responsibility for the final manuscript remains with the author.

References

Damien Ackerer, Julien Hugonnier, and Urban Jermann. Perpetual futures pricing. *Mathematical Finance*, 36(3):481–499, 2026. doi: 10.1111/mafi.70018. URL https://doi.org/10.1111/mafi.70018.

American Economic Association. JEL Classification System / EconLit Subject Descriptors, 2026. URL https://www.aeaweb.org/econlit/jelCodes.php?view=jel. Accessed 4 November 2026.

Aster. Pre-IPO Perpetuals, 2026. URL https://docs.asterdex.com/trading/perpetuals/pre-ipo-perpetuals. Official documentation; accessed 4 November 2026.

Binance. Margined Perpetual Contract Pre-IPO Trading (2026-05-26), 2026. URL https://www.binance.com/en/support/announcement/detail/94f88766c6c84d68a07f006bc076ba6a. Official announcement, 28 May 2026; accessed 4 November 2026.

BingX. USDT-M Perpetual Futures Market Data: Order Book, 2026. URL https://bingx-api.github.io/docs/#/en-us/swapV2/market-api.html. Official API documentation and embedded SHIB-USDT example; accessed 4 November 2026.

Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. Price discovery without trading: Evidence from limit orders. *The Journal of Finance*, 74(4):1621–1658, 2019. doi: 10.1111/jofi.12769. URL https://doi.org/10.1111/jofi.12769.

Riccardo De Blasis and Alexander Webb. Arbitrage, contract design, and market structure in bitcoin futures markets. *Journal of Futures Markets*, 42(3):492–524, 2022. doi: 10.1002/fut.22305. URL https://doi.org/10.1002/fut.22305.

Francis X. Diebold and Roberto S. Mariano. Comparing predictive accuracy. *Journal of Business & Economic Statistics*, 13(3):253–263, 1995. doi: 10.1080/07350015.1995.10524599. URL https://doi.org/10.1080/07350015.1995.10524599.

Wenxin Du, Alexander Tepper, and Adrien Verdelhan. Deviations from covered interest rate parity. *The Journal of Finance*, 73(3):915–957, 2018. doi: 10.1111/jofi.12620. URL https://doi.org/10.1111/jofi.12620.

Gate. Gate to Launch Pre-Market Perpetual Contracts for OPENAI, ANTHROPIC and Other Tokens, 2026. URL https://www.gate.com/announcements/article/50673. Official announcement 50673, 13 April 2026; accessed 4 November 2026.

Gary B. Gorton, Fumio Hayashi, and K. Geert Rouwenhorst. The fundamentals of commodity futures returns. *Review of Finance*, 17(1):35–105, 2013. doi: 10.1093/rof/rfs019. URL https://doi.org/10.1093/rof/rfs019.

Denis Gromb and Dimitri Vayanos. Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2–3):361–407, 2002. doi: 10.1016/S0304-405X(02)00228-3. URL https://doi.org/10.1016/S0304-405X(02)00228-3.

Denis Gromb and Dimitri Vayanos. Limits of arbitrage. *Annual Review of Financial Economics*, 2(1):251–275, 2010. doi: 10.1146/annurev-financial-073009-104107. URL https://doi.org/10.1146/annurev-financial-073009-104107.

Peter Reinhard Hansen. A test for superior predictive ability. *Journal of Business & Economic Statistics*, 23(4):365–380, 2005. doi: 10.1198/073500105000000063. URL https://doi.org/10.1198/073500105000000063.

Campbell R. Harvey, Yan Liu, and Heqing Zhu. … and the cross-section of expected returns. *Review of Financial Studies*, 29(1):5–68, 2016. doi: 10.1093/rfs/hhv059. URL https://doi.org/10.1093/rfs/hhv059.

Joel Hasbrouck. One security, many markets: Determining the contributions to price discovery. *The Journal of Finance*, 50(4):1175–1199, 1995. doi: 10.1111/j.1540-6261.1995.tb04054.x. URL https://doi.org/10.1111/j.1540-6261.1995.tb04054.x.

Hyperliquid. Contract Specifications, 2026a. URL https://hyperliquid.gitbook.io/hyperliquid-docs/trading/contract-specifications. Official documentation; accessed 4 November 2026.

Hyperliquid. HIP-3: Builder-Deployed Perpetuals, 2026b. URL https://hyperliquid.gitbook.io/hyperliquid-docs/hyperliquid-improvement-proposals-hips/hip-3-builder-deployed-perpetuals. Official documentation; accessed 4 November 2026.

Hyperliquid. Perpetuals: Info Endpoint, 2026c. URL https://hyperliquid.gitbook.io/hyperliquid-docs/for-developers/api/info-endpoint/perpetuals. Official documentation; accessed 4 November 2026.

Martin Lettau and Ananth Madhavan. Exchange-traded funds 101 for economists. *Journal of Economic Perspectives*, 32(1):135–154, 2018. doi: 10.1257/jep.32.1.135. URL https://doi.org/10.1257/jep.32.1.135.

Lighter. Funding, 2026a. URL https://docs.lighter.xyz/trading/funding. Official documentation; accessed 4 November 2026.

Lighter. RWA Pricing Mechanism, 2026b. URL https://docs.lighter.xyz/trading/real-world-assets-rwas/rwa-pricing-mechanism. Official documentation; accessed 4 November 2026.

Lighter. Futures Contract Price Rolling Mechanism, 2026c. URL https://docs.lighter.xyz/trading/real-world-assets-rwas/futures-contract-price-rolling-mechanism. Official documentation; accessed 4 November 2026.

Lighter. Terms of Service, 2026d. URL https://lighter.xyz/terms. Official terms; accessed 4 November 2026.

Hang Liu. Measuring executable net carry in cross-venue perpetual funding-rate arbitrage, 2026a. URL https://doi.org/10.5281/zenodo.23868361. humarb Research 01; published 2 November 2026.

Hang Liu. Framing net carry for cross-venue perpetual funding positions: A settled-history, depth-walked, and liquidation-aware method, 2026b. URL https://doi.org/10.5281/zenodo.23897013. humarb Research 02; published 3 November 2026.

Igor Makarov and Antoinette Schoar. Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319, 2020. doi: 10.1016/j.jfineco.2019.07.001. URL https://doi.org/10.1016/j.jfineco.2019.07.001.

Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610. URL https://doi.org/10.2307/1913610.

OKX. Pre-IPO Perpetual Futures, 2026a. URL https://www.okx.com/en-us/help/pre-ipo-perpetual-futures. Official product documentation; accessed 4 November 2026.

OKX. OKX to Execute Rebase on OPENAIUSDT and ANTHROPICUSDT Pre-Market Perpetual Futures, 2026b. URL https://www.okx.com/en-us/help/okx-to-execute-rebase-on-openaiusdt-and-anthropicusdt-pre-market-perpetual. Official announcement; adjustment ratio 10 on 30 June 2026; accessed 4 November 2026.

Oura Inc. Amendment No. 1 to Form S-1 Registration Statement, 2026. URL https://www.sec.gov/Archives/edgar/data/2133022/000119312526396051/d119865ds1a.htm. Filed 21 September 2026; accessed 4 November 2026.

Antti Petajisto. Inefficiencies in the pricing of exchange-traded funds. *Financial Analysts Journal*, 73(1):24–54, 2017. doi: 10.2469/faj.v73.n1.7. URL https://doi.org/10.2469/faj.v73.n1.7.

Robert J. Shiller. Measuring asset values for cash settlement in derivative markets: Hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931, 1993. doi: 10.1111/j.1540-6261.1993.tb04024.x. URL https://doi.org/10.1111/j.1540-6261.1993.tb04024.x.

Andrei Shleifer and Robert W. Vishny. The limits of arbitrage. *The Journal of Finance*, 52 (1):35–55, 1997. doi: 10.1111/j.1540-6261.1997.tb03807.x. URL https://doi.org/10.1111/j.1540-6261.1997.tb03807.x.

trade[XYZ]. About trade[XYZ], 2026a. URL https://docs.trade.xyz/overview/about-trade-xyz.md. Official documentation; accessed 4 November 2026.

trade[XYZ]. Discovery Bounds, 2026b. URL https://docs.trade.xyz/perpetuals/mechanics/discovery-bounds.md. Official documentation; accessed 4 November 2026.

trade[XYZ]. Holiday Closures, 2026c. URL https://docs.trade.xyz/perpetuals/specifications-and-schedules/holiday-closures.md. Official documentation; accessed 4 November 2026.

trade[XYZ]. Oracle Price, 2026d. URL https://docs.trade.xyz/perpetuals/mechanics/oracle-price.md. Official documentation; accessed 4 November 2026.

trade[XYZ]. Roll Schedules, 2026e. URL https://docs.trade.xyz/perpetuals/specifications-and-schedules/roll-schedules.md. Official documentation; accessed 4 November 2026.

Halbert White. A reality check for data snooping. *Econometrica*, 68(5):1097–1126, 2000. doi: 10.1111/1468-0262.00152. URL https://doi.org/10.1111/1468-0262.00152.

A Derivation of Rebase Invariance and Reference-Share Sizing

Suppose a venue rebases the assumed share count from N to $N' = kN$ while reducing the native price from P to $P' = P/k$, with unchanged quote-currency conversion X. Then

$$V' = P' \frac{N'}{10^9} X = \frac{P}{k} \frac{kN}{10^9} X = P \frac{N}{10^9} X = V. \tag{10}$$

This is the value-neutrality condition used in the paper. It explains why a native-price chart can show a mechanical discontinuity while implied equity value remains continuous. The condition is an accounting identity, not evidence that a venue actually executed a rebase correctly or that market marks remained continuous around the event.

For two instruments with native quantities q_A and q_B, the valuation sensitivities are q_A/N_A and q_B/N_B. Equating them gives $q_B = q_A(N_B/N_A)$. If order contracts have multipliers κ_A and κ_B, then order quantities must additionally satisfy $q_B = q_A(\kappa_A/\kappa_B)(N_B/N_A)$. The derivation ignores basis, fees, rounding, margin, and timing by design; those are separate execution constraints.

For an anchor value V_A, arithmetic and log residuals satisfy $r = 100(\exp(r^{\log}/100) - 1)$ and $r^{\log} = 100 \log(1 + r/100)$. The log form gives $r^{\log}_{B:A} = -r^{\log}_{A:B}$ when the comparison direction is reversed, whereas the arithmetic form is anchor-dependent. This property motivates the alternative-residual robustness check.

B Hyperliquid Coverage and Order Books

Table 7. Complete Hyperliquid DEX coverage snapshot, 20 September 2026

DEX	Metadata entries	Non-delisted	Collateral index	Collateral token
Default	234	178	0	USDC
`xyz`	123	108	0	USDC
`para`	35	27	0	USDC
`mkts`	24	4	0	USDC
`io`	10	8	0	USDC
`flx`	16	0	360	USDH
`vntl`	15	0	360	USDH
`hyna`	25	0	235	USDE
`km`	23	0	360	USDH
`cash`	17	0	268	USDT0
`abcd`	1	0	0	USDC
Total	523	325	—	—

Source: official `perpDexs`, `meta`, and `spotMeta` responses as preserved in the de-identified source study. Non-delisted entries are not unique underlyings, account-eligible markets, or executable pairs.

Table 8. Representative two-sided Hyperliquid books, 20 September 2026

Full coin	Bid	Ask	Evidence qualification
`xyz:NVDA`	220.12	220.15	Non-delisted; two-sided book
`xyz:AAPL`	333.57	333.59	Non-delisted; two-sided book
`xyz:GOLD`	4,370.5	4,370.6	Non-delisted; two-sided book
`xyz:SILVER`	66.215	66.216	Non-delisted; two-sided book
`xyz:CL`	97.53	97.535	API CL; display name WTIOIL
`xyz:BRENTOIL`	100.97	100.99	Non-delisted; two-sided book
`xyz:EUR`	1.1519	1.1520	Non-delisted; two-sided book
`xyz:JPY`	156.09	156.13	USD/JPY direction from specification
`xyz:SP500`	7,637.2	7,637.3	Index construction requires review
`mkts:USTECH`	717.77	717.89	Unit and scaling unreviewed

Observation window: approximately 12:54:41–12:54:43 UTC on a Sunday. Native prices are historical and must not be reused as current quotes. A book does not establish account permission, external-reference availability, economic equivalence, or executable size.

C Additional Robustness Tables

D Reproduction Guide and Data Schema

The reproduction package separates source observations in `data/` from generated outputs in `output/`. Running `python3scripts/reproduce.py` performs the following steps without network access:

1. Read and sum Hyperliquid coverage counts, asserting the 523-entry and 325-non-delisted totals.
2. Compute canonical pre-IPO values using Decimal arithmetic and apply the stated Binance or cross-venue-median anchors.

Table 9. Early-window ratios and exact later-window counterparts

Company	Venue	Metric	Early, 07:46–08:05 UTC	Later, 08:33–08:42 UTC
OPENAI	OKX	Raw ratio	0.101	0.10114781
ANTHROPIC	OKX	Raw ratio	0.103	0.10308438
OPENAI	Lighter	Normalized ratio	1.033	1.03208377
ANTHROPIC	Lighter	Normalized ratio	1.042	1.04089182
POLYMARKET	Aster	Normalized ratio	0.900	0.90469602
OURA	Gate	Normalized ratio	1.049	1.05050092

Early values are rounded summaries; later values are exact calculations from preserved marks. The table supports same-day sign checks only.

Table 10. Gate KIMI share counts implied by its worked examples

Example valuation (USD bn)	Example price (USD)	Implied shares (billions)
30	3.49	8.59598854
40	4.66	8.58369099
50	5.82	8.59106529
Mean	—	8.59024827
Separately stated count	—	10.00000000

Implied shares equal valuation divided by example price. The consistent examples conflict with the separately stated count, so no factor is selected by price fit. Source: Gate announcement 100758.

Table 11. Alternative median anchors and selected residuals

Company	Venue	Median anchor (USD bn)	Arithmetic residual (%)	Supplied-anchor residual (%)
OPENAI	OKX	1,642.830	+0.844275	+1.147811
OPENAI	Bitget	1,642.830	−3.833629	−3.544172
OPENAI	Lighter	1,642.830	+2.898657	+3.208377
ANTHROPIC	OKX	2,043.250	+0.922550	+3.084383
ANTHROPIC	Lighter	2,043.250	+1.906277	+4.089182
POLYMARKET	Aster	23.868	−9.598626	−9.530398
OURA	Gate	15.727130	+4.851793	+5.050092

Medians use actual eligible exact venue observations and exclude the POLYMARKET pseudo-anchor and Gate KIMI conflict. MOONSHOT has only two eligible exact observations and is not emphasized.

3. Generate arithmetic and log residuals, alternative median anchors, company hold-outs, and unit-evidence groups.
4. Recompute OKX raw ratios, Gate KIMI example-implied counts and scenarios, and all ten BingX quantity ratios.
5. Write the five SVG and PDF figures and `summary.json`.

The central source fields are `company`, `venue`, `native_symbol`, `native_mark`, `factor_n_over_1e9`, `unit_evidence`, `reference_type`, `anchor_rule`, and `observation_window_utc`. The canonical unit is USD billions. `unit_evidence=conflicted` prevents inclusion in the eligible residual summary. The `cross_venue_median` venue label denotes a pseudo-anchor, not a market. Range summaries use separate `native_mark_min`, `native_mark_max`, `residual_min_percent`, and `residual_max_percent` fields and are never expanded into exact observations.

The optional refresh script records URL, method, parameters, UTC request time, HTTP status, byte count, and SHA-256 hash. It does not bypass an access challenge. A refresh is a new evidence layer, not an overwrite of historical inputs. Full field definitions and lineage are in `data_dictionary.md` and `README.md`.

E Document Conflict and Hold-Out Register

Table 12. Representative conflicts and required classification

Object	Evidence issue	Required treatment
Gate KIMI	Ten-billion-share statement conflicts with 8.59-billion worked examples	Hold out; report both scenarios
Delisted or inactive lines	Metadata or stale marks can persist after trading status changes	Preserve lifecycle evidence; no executable-pair claim
Bybit and unread support pages	Public interface or page unavailable from measurement host	Coverage unknown; do not infer absence
Lighter historical versus current RWA rules	Later documents change cap treatment and reference sources	Version by effective and access dates
SPCX legacy symbols and QNTX flags	Names or pre-market flags can lag listed-equity conversion	Review class and underlying again
SP500 versus SPY	Index points differ from ETF shares	Reject one-to-one mapping
WTI versus Brent	Related energy curves but different commodities	Reject same-asset conversion
USD/JPY versus JPY/USD	Reciprocal quote directions	Invert before comparison; never subtract directly
BingX ordinary quantity	Unit not stated; Coin array explicitly base-denominated	Prefer documented Coin field; otherwise unit unknown

This register is methodological and does not contain production routing rules, display priorities, thresholds, or account-specific decisions.