Freshness and Clocks in Multi-Venue Perpetual Market Data: Transport Capabilities, Staleness Practice and Timestamp Offsets

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

humarb Research 09 · Author manuscript v1
DOI: 10.5281/zenodo.23932021

Abstract

Cross-venue perpetual-market comparisons require data that are timely and timestamps whose meanings are understood. We audit public WebSocket capabilities at 22 venues and analyze 14,931 public funding-history records for BTC, ETH and SOL over 64 UTC dates in 2026. The new large-sample offset study aggregates to 4,977 venue-hour observations across Hyperliquid, Pacifica, Extended and edgeX. Median offsets after the nominal UTC hour are 37, 2,513, 849 and 0 ms; corresponding 95th percentiles are 88, 5,083, 1,593.15 and 0 ms. One Extended record at 1,096,849 ms is retained as an anomaly. A fixed calibration-period 99th-percentile-plus-1-ms rule attains later-period coverage of 99.35%, 100%, 96.61% and 100%, respectively, showing that historical tolerances need not transfer across venues or regimes. We also derive worst-case bounds for clock-difference exceedance shares when timestamps are missing. The results support venue-specific timestamp semantics, explicit anomaly reporting and partial identification rather than a universal staleness threshold.

Keywords: perpetual futures; market data; timestamp offsets; data staleness; WebSocket; partial identification

JEL classification: G12; G13; G14; C58

1 Introduction

A cross-venue market comparison is only as coherent as the clocks used to construct it. Two funding rates may appear side by side although one is a forecast, another is a settled rate, and a third is the timestamp of a database record rather than an economic event. A current order book may be easy to subscribe to but difficult to interpret when its update identifier does not connect to the preceding message. Conversely, an old exchange timestamp need not imply a broken feed: some channels repeat an unchanged state on a schedule, while others remain silent until the state changes. These distinctions are not implementation trivia. They determine whether an observation is eligible for a comparison, which nominal interval receives a funding cash flow, and what remains unknown

when a timestamp is absent.

This paper asks a focused measurement question: **what can public transport documentation and public funding histories establish about data freshness, timestamp semantics, and clock-offset uncertainty in multi-venue perpetual markets?** The question is not answered by a single number of seconds. Transport freshness concerns whether a stream exists, whether it carries the required fields, whether a local sequence remains continuous, and whether the connection is alive. Event alignment concerns the relation between a returned timestamp and the nominal settlement grid. Missing-clock inference concerns what can be said without assuming that missing observations resemble observed ones. Treating these as separate layers is the paper's organizing principle.

The empirical design has three components. First, a de-identified audit dated 2026-10-09 covers 22 perpetual-trading venues. It records public WebSocket availability for funding-related data, prices and order books, together with all-market subscription capabilities and selected single-host first-frame observations. Eighteen venues have a funding-related WebSocket route in that audit; nine have a full all-market funding channel and one has partial all-market coverage. Four have a material gap in the public WebSocket surface, although the nature of the gap differs: a missing funding channel is not the same as a request-for-quote system without a conventional order book. The matrix is used to describe capabilities, not to rank venues by latency. Its first-frame observations come from one host on one day and do not support a multi-location latency distribution.

Second, we collect a new public-history sample for BTC, ETH and SOL from 2026-09-01 through 2026-11-03 UTC. The frozen sample contains 14,931 records from Hyperliquid, Pacifica, Extended and edgeX. After grouping records by venue and nominal UTC hour and taking the median across assets, the primary dataset contains 4,977 venue-hour observations over 64 UTC dates. This aggregation matters: three assets often share exactly the same venue-generated timestamp, so counting them as independent clock draws would exaggerate the effective sample size. The unit, time field and record semantics are prespecified for each venue. In particular, edgeX contributes only records explicitly marked as settlements, and its observed four-hour grid is not forced into an hourly denominator.

Third, we study what happens when a two-legged comparison lacks one or both valid clocks. Let A be the number of pairs eligible for a clock test, n the number with two valid clocks, $m = A - n$ the number without them, and o the number of valid-clock pairs whose absolute clock difference strictly exceeds a research threshold τ. Without a missing-at-random assumption, the population exceedance share is partially identified by $[o/A, (o + m)/A]$, with no value assigned when $A = 0$. This is an identified set, not a confidence interval and not the complete-case proportion o/n. The distinction provides a disciplined way to report missingness rather than silently improving the denominator.

The new offset estimates show large and persistent venue differences. Median offsets after the nominal UTC hour are 37 ms for Hyperliquid, 2,513 ms for Pacifica, 849 ms for Extended and 0 ms for edgeX. The corresponding 95th percentiles are 88, 5,083, 1,593.15 and 0 ms. Hyperliquid's maximum observed offset is 139 ms, Pacifica's is 5,084 ms, and edgeX is exactly aligned throughout its 384 observed settlement events. Extended has a conventional upper tail below two seconds but also one extreme record at 2026-09-24 01:18:16.849 UTC, equivalent to an offset of 1,096,849 ms. That record is common to all three assets and remains in the data. Reporting only Extended's median would conceal it; reporting only its maximum would misrepresent the rest of the distribution.

A chronological diagnostic further shows why a tolerance estimated from the past is not a guarantee. For each venue we compute the empirical 99th percentile during the calibra-

tion period, 2026-09-01 through 2026-10-08, add 1 ms, and freeze the result before scoring observations from 2026-10-09 onward. Later-period coverage is 99.35% for Hyperliquid, 100% for Pacifica, 96.61% for Extended and 100% for edgeX. Extended therefore has 21 exceedances among 620 holdout venue-hours despite its ordinary full-sample tail appearing modest. The diagnostic involves one fixed rule per venue and no search over tolerances. It does not certify a production rule, forecast exchange behavior, or convert a resampling interval into a venue promise.

The contributions are fourfold. First, the paper provides a 22-venue transport-capability map that separates field availability, all-market reach and price semantics while explicitly bounding the latency evidence. Second, it supplies a response-level reproducible measurement of timestamp offsets using multiple assets and 64 UTC dates, with raw-response hashes and a venue-hour estimand that avoids pseudoreplication. Third, it evaluates a fixed event-assignment tolerance chronologically and retains an extreme observation that a conventional cleaning rule might erase. Fourth, it integrates missing-clock uncertainty into the same measurement framework through worst-case partial-identification bounds. The deliverable includes the public responses, normalized data, code, data dictionary, machine-readable tables and vector figures required to recompute every headline number.

Identification is descriptive and algebraic, not causal. There is no untreated venue, randomized clock, return forecast or counterfactual execution outcome. The conclusions rely on public-record integrity, correct semantic classification and transparent aggregation. If a venue changes the meaning of a field without changing its name, the recorded offsets can mix regimes; the paper addresses this risk through period and weekly checks but cannot eliminate it. If missing clocks are systematically associated with large clock differences, the complete-case share can be biased in either direction; the reported bounds deliberately remain agnostic. No claim is made about realized arbitrage, execution quality, trading profit, or an optimal production staleness threshold.

This study extends humarb Research (2026a) and Liu (2026a). Research 01 measures executable net carry after costs, and Research 02 develops a settled-history, asynchronous-settlement accounting frame. Both require inputs whose event times and availability are correctly represented. The present paper examines that upstream measurement layer: whether the necessary streams exist, what their timestamps mean, how observed offsets behave, and how missing clocks should affect reported coverage. It does not re-estimate net carry and does not validate any trading rule.

Section 2 positions the study in the perpetual-futures, low-latency, distributed-clock and partial-identification literatures. Section 3 distinguishes funding, price and transport mechanisms. Section 4 documents sources, coverage and collection failures. Section 5 defines the estimators and bounds. Section 6 reports transport and offset results. Section 7 presents chronological, weekly, asset, weighting and semantic checks. Sections 8 and 9 state the evidence boundary and conclusions.

2 Related literature

2.1 Perpetual-futures mechanisms and funding-time economics

Perpetual contracts have a long intellectual history in derivative design (Shiller, 1993), but cryptocurrency venues implement them through heterogeneous funding formulas, collateral conventions and settlement grids. De Blasis and Webb (2022) connect arbitrage,

contract design and Bitcoin futures market structure. Recent work studies the two-tiered organization of funding-rate markets (Zhivkov, 2026), temporal variation across centralized and decentralized perpetual venues (Zhivkov et al., 2026), and the mechanics and economic claims of decentralized perpetual futures (Shah and Pinsky, 2026). These contributions establish that contract design and funding dynamics are economically meaningful; they do not make a timestamp carried by a public response self-explanatory.

Funding timing has direct pricing content in some settings. Erez and Smirnov (2026) derive no-arbitrage bounds in decentralized perpetual markets, emphasizing the implications of funding design for pricing. Inan (2026) evaluates funding-rate predictability, while Werapun et al. (2026) and Pindza (2026) study funding-rate arbitrage and its risks across centralized and decentralized venues. A researcher addressing those questions needs to know whether an observation is a predicted rate, a rate that will apply at the next settlement, a rate already applied, or a snapshot written after the nominal event. Misclassification can create look-ahead information or shift cash flows between intervals. Our paper does not estimate a funding premium or an arbitrage return. It supplies the timestamp and transport audit that such studies require before economic interpretation.

This positioning also distinguishes the paper from conventional forecast evaluation. A forecast model maps information available at time t into a future funding outcome; a timestamp-offset measurement describes the relation between a historical record's time field and a nominal grid. The latter can affect the information set of the former, but it is not itself a forecast. Consequently, the chronological exercise below is a fixed-rule diagnostic, not a claim of predictive superiority. We do not compare forecast losses, estimate a Sharpe ratio, or report returns net of costs.

2.2 Fragmented markets, speed and the meaning of freshness

Cross-venue cryptocurrency markets exhibit segmentation and arbitrage frictions (Makarov and Schoar, 2020). Market-microstructure research also shows why speed can alter strategic behavior and market quality. Hasbrouck and Saar (2013) study low-latency trading, Budish et al. (2015) analyze the high-frequency trading arms race and frequent batch auctions, and Foucault et al. (2016) examine news trading and speed. In these settings, latency generally concerns the time required for information or orders to move through a trading system. That concept is related to, but not identical with, an offset stored inside an exchange's historical record.

Three clocks should therefore be distinguished. The source-event clock describes when an exchange says an economic event occurred or a rate was applied. The source-record clock describes when a record or snapshot was created. The client-receipt clock describes when a particular consumer received a response. A first-frame duration is dominated by connection and transport conditions; a historical `created_at` field is dominated by record-generation semantics; a settlement timestamp is intended to identify an event grid. A single measurement cannot substitute for all three. This separation prevents a low first-frame duration from being misread as a small settlement offset, and vice versa.

High-frequency econometrics offers a related warning about sampling and noise. Zhang et al. (2005) distinguish the timescales at which a latent process and microstructure noise operate. Our estimand is not integrated volatility, and funding-history timestamps are not transaction prices. The shared lesson is narrower: increasing the number of records does not automatically improve the relevant measurement if observation mechanisms differ or records are dependent. We address this by analyzing venue-hours rather than treating three assets with a common timestamp as three independent draws.

The public-practice review reinforces the absence of a universal freshness second count. Bybit documents depth-dependent push frequencies and a three-second repeated snapshot for unchanged level-one books (Bybit, nd). Binance's local-book instructions emphasize sequence continuity rather than a fixed age limit (Binance, nd). Pyth lets the caller specify the maximum acceptable age and exposes stale-price errors (Pyth Network, nd). Chainlink recommends controls for stale, delayed, missing or unexpected inputs and fallback or pause behavior (Chainlink, nd). These are different solutions to different failure modes. They support a taxonomy of sequence, heartbeat, age and status checks rather than a universal threshold imported from an oracle into an order book or funding stream.

2.3 Distributed clocks and externally visible timestamps

Distributed-systems research supplies precise reasons not to equate a timestamp with a complete ordering proof. Lamport (1978) formalize logical time and event ordering in a distributed system. Logical clocks can establish causal order without synchronizing every process to a common physical time. Market-data alignment requires a related but distinct operation: mapping externally generated records onto nominal UTC settlement intervals while acknowledging transport and record-creation delays.

Physical synchronization has its own error budget. Kopetz and Ochsenreiter (1987) analyze clock synchronization in distributed real-time systems, and Cristian (1989) develops probabilistic clock synchronization. NTP specifies a protocol and algorithms for network time (Mills et al., 2010); IEEE 1588 specifies precision clock synchronization for networked measurement and control (Institute of Electrical and Electronics Engineers, 2020); Yiğitler et al. (2020) review clock-discipline algorithms and protocols. This literature concerns estimating and disciplining clocks. Our public-history data do not provide the round-trip measurements or clock-discipline model needed to estimate a venue's clock error relative to a trusted time source. An offset after the nominal hour can include intentional scheduling, processing delay and record semantics; it is not identified as physical clock skew.

For that reason, the paper avoids the phrase "the exchange clock is fast by x ms." The defensible statement is "the returned timestamp lies x ms after the nominal UTC hour." A consumer should additionally retain its own receipt timestamp if it needs transport-age calculations. Even a well-disciplined client clock cannot separate source scheduling from source clock error using history alone.

2.4 Missing data and partial identification

Complete-case analysis is point-identified only under assumptions that relate missingness to the outcome. When clocks are missing because a venue does not expose a field, because a parser rejects an invalid value, or because a transport route differs by asset, missingness may be systematically related to the clock difference itself. The problem fits the partial-identification framework developed by Manski (2003, 2005) and reviewed by Tamer (2010). Rather than choose a convenient imputation, we derive the set of population shares compatible with the observed counts and all possible missing outcomes.

The bounds are intentionally assumption-light. They are exact accounting implications of $A = n + m$ and a binary exceedance indicator, not sampling confidence intervals. Their width is m/A, so missingness directly limits what can be claimed. Additional assumptions—such as missing at random, bounded venue-specific missing outcomes, or an instrumental restriction—could narrow the set, but they require separate justification

and are not imposed here.

Sampling uncertainty is also conceptually separate. Where we report resampling intervals for offset medians and upper quantiles, we use a seven-day moving-block bootstrap over UTC dates in the spirit of dependent-observation resampling (Künsch, 1989). Holdout coverage receives a Wilson binomial interval as a descriptive finite-sample summary (Wilson, 1927), while acknowledging that venue-hours within a date are dependent. Neither device turns a historical distribution into an exchange guarantee, and neither narrows the partial-identification bounds without an additional assumption.

2.5 Positioning

The resulting study sits between market microstructure, data engineering and measurement methodology. Its object is not the economic magnitude of speed, the no-arbitrage price of a perpetual, or a new clock-synchronization protocol. It is the public evidence needed to avoid false synchronization in empirical work. The transport matrix describes what can be observed, the offset study describes how returned event fields behave, and the bounds describe what remains indeterminate when clocks cannot be observed. This layered contribution is narrower than a trading study but necessary for making subsequent trading or pricing analyses auditable.

3 Institutional background

3.1 Funding mechanisms create several different time fields

A perpetual future has no expiry at which a conventional futures basis must converge. Instead, periodic funding transfers create an economic link between the contract and an underlying index or oracle. Although venue formulas differ, a stylized premium can be written as

$$P_t = \frac{\max(0, B_t - I_t) - \max(0, I_t - A_t)}{I_t}, \tag{1}$$

where I_t is an index or oracle price and B_t and A_t are impact bid and ask prices. A stylized funding rule is

$$F = \text{clamp}\big(\bar{P} + \text{clamp}(r - \bar{P}, -\delta, \delta), F_{\min}, F_{\max}\big)\,, \tag{2}$$

where $\bar{P}$ is an average premium, r an interest component, δ a damping bound, and $F_{\min}$ and $F_{\max}$ venue- or contract-specific caps. Equations (1)–(2) are institutional notation, not a claim that one formula describes every venue and not a disclosure of any production parameter.

Hyperliquid documents five-second premium sampling, an eight-hour-style rate paid hourly at one-eighth of the computed amount, and an hourly cap (Hyperliquid, nda). Its historical funding endpoint returns `time` in milliseconds and illustrates a timestamp close to, rather than exactly equal to, the nominal hour (Hyperliquid, ndb). Pacifica likewise documents five-second sampling and application of the average rate at the end of each one-hour interval; its public formula includes premium, interest and damping components (Pacifica, nda). Its history endpoint instead exposes `created_at`, described as a millisecond timestamp attached to the record (Pacifica, ndb). These definitions already warn against treating both fields as identical settlement instants.

Extended states that funding is calculated every minute but applied once per hour, and that history rows represent the one-hour rates applied to payments. The field `T` is

documented as the epoch-millisecond time when the rate was calculated and applied (Extended, nd). edgeX distinguishes `fundingTime`, associated with settlement, from `fundingTimestamp`, associated with calculation. It also marks settlement rows explicitly with `isSettlement` and exposes an interval field (edgeX, nd). The current contract metadata in our sample reports 240-minute intervals for BTC, ETH and SOL. Using every minute-level row, or assuming an hourly settlement grid, would therefore change the estimand rather than increase the sample.

The time fields used below consequently have different meanings. Hyperliquid's `time` and Extended's `T` are associated with applied funding history; Pacifica's `created_at` is a record timestamp whose equality with the settlement instant is not guaranteed; edgeX's `fundingTime` is analyzed only for explicit settlement rows. The offset is always measured relative to the nominal hour containing the returned timestamp. It is not a direct estimate of physical clock skew.

3.2 Transport capability is not a single binary variable

A venue can support WebSocket transport without offering every field required for a particular comparison. Funding-related data, a mark or oracle price, and order-book depth are separate capabilities. A venue can additionally require per-contract subscriptions, offer a partial all-market ticker without funding, or provide a full all-market channel. A large all-market frame can reduce subscription count but does not guarantee that every field updates on the same schedule.

Transport encoding also matters. A text JSON ticker and a binary Simple Binary Encoding feed can carry comparable economic content while requiring different parsers and schema-version controls. A request-for-quote surface is not a conventional central-limit order book: displayed quote tiers do not automatically supply the same depth-walk interpretation. Oracle-based trading systems similarly require a different execution-cost model. The transport matrix therefore records "unavailable for this purpose" rather than collapsing every nonconforming design into a single technical failure label.

The same caution applies to price names. A mark price, index price and oracle price may be related, but they can enter funding, margin and liquidation calculations differently. Hyperliquid, for example, documents distinct oracle and mark-price constructions and approximately three-second validator updates (Hyperliquid, ndc). The matrix marks venues whose available stream exposes an oracle or index price rather than a field explicitly named mark price. This notation preserves semantic uncertainty instead of silently translating every price into one common field.

3.3 Freshness mechanisms answer different failure questions

Public systems use at least six freshness mechanisms. First, *sequence continuity* asks whether the current update connects to the preceding local state. Binance's local-book rule uses the previous-final-update identifier for this purpose (Binance, nd). Second, *heartbeat and connection timeout* asks whether a transport remains alive; OKX documents connection-level timing rules in its API reference (OKX, nd). Third, *unchanged-state repetition* distinguishes an old timestamp caused by silence from a broken stream. Bybit's repeated level-one snapshot after three unchanged seconds is an example (Bybit, nd).

Fourth, *source-age limits* ask whether a value is too old for a particular consumer. Pyth's `getPriceNoOlderThan` makes the maximum age a caller argument and returns a

stale-price error when the applicable validity condition fails (Pyth Network, nd). Fifth, *status and quality flags* communicate whether a provider considers a value suspect even while a stream remains open; the reviewed LSEG community response illustrates this status-based approach (LSEG Developer Community, nd). Sixth, *semantic freshness* asks whether a field still represents the same economic object after a regime, schema or contract change. None of the first five mechanisms proves the sixth.

Open-source implementations mirror this diversity. The pinned Hummingbot source uses a 60-second user-stream inactivity comparison to select a five-second rather than 120-second polling interval, while its order-book tracker records the last diff timestamp (Hummingbot Foundation, 2026). Condor exposes a restart operation for a stale order book after a WebSocket disconnect (Hummingbot Foundation, nd). CCXT describes snapshot-and-delta maintenance, sequence checks, reconnections and keepalives as parts of a broader real-time data architecture (CCXT, 2026). These mechanisms should not be paraphrased as a universal "60-second order-book freshness rule": their scope and object differ.

A defensible monitoring design can therefore represent freshness as a vector rather than a single age. Let

$$\mathcal{F}_{v,k,t} = (C_{v,k,t}, H_{v,k,t}, S_{v,k,t}, A_{v,k,t}, Q_{v,k,t}, M_{v,k,t}), \tag{3}$$

where, for venue v and data kind k, C denotes sequence continuity, H heartbeat state, S unchanged-state behavior, A source and receipt ages, Q provider quality status, and M semantic validity. The vector is conceptual; this paper does not estimate a production score or prescribe product behavior. Its purpose is to make clear which failure a public mechanism can and cannot detect.

3.4 Settlement alignment is a separate problem

Even a perfectly fresh historical response can be difficult to assign to a settlement interval. Let $t_{v,a,j}$ be the returned timestamp for venue v, asset a and record j. The nominal UTC hour is $h_{v,a,j} = 3{,}600{,}000\lfloor t_{v,a,j}/3{,}600{,}000 \rfloor$, and the within-hour offset is $d_{v,a,j} = t_{v,a,j} - h_{v,a,j}$. A positive d does not imply a defective feed. It may reflect computation, database writing, intentional scheduling or a snapshot convention. Event assignment requires an empirical distribution and an explicit out-of-range action, not merely rounding every timestamp to the nearest hour.

Missing clocks add a final layer. If one leg lacks a valid timestamp, a pair cannot be assigned a measured clock difference. Treating that pair as having zero difference fabricates agreement; dropping it from both numerator and denominator can manufacture a favorable complete-case rate. The partial-identification bounds in Section 5 keep this uncertainty visible. Together, funding semantics, transport capability, freshness state and missing-clock inference define the institutional background for the measurements that follow.

4 Data

4.1 Transport-capability audit

The transport evidence is a de-identified audit of 22 perpetual-trading venues conducted on 2026-10-09 (humarb Research, 2026b). For each venue, the matrix records whether public WebSocket channels provide funding-related information, price information and

order-book data; whether an all-market funding subscription is available; selected first-frame durations; and a public documentation URL. Price semantics distinguish a venue-defined mark price from an oracle- or index-only stream. All-market status distinguishes full, partial and absent coverage.

The matrix is historical evidence about the audited access surface, not a permanent statement about venue support. Documentation, endpoints and schemas can change after the snapshot. Several entries also carry explicit access limitations. GRVT's documentation site returned an access-protection response in the original environment, so official software-development-kit definitions were used for field verification; edgeX documentation content was not fully retrievable on that date. The audit does not infer missing capabilities from inaccessible pages when an official alternative source is unavailable.

First-frame durations receive a narrower interpretation. They were observed from one measurement host, through a proxy, using Node 22's global WebSocket implementation. The durations combine DNS and connection setup, proxy routing, subscription scheduling and the source's decision about when to emit the first frame. They are not independent draws from multiple locations and cannot identify a venue's network latency. We report them only as supplementary single-day observations in the machine-readable matrix and do not use them in the paper's statistical comparisons.

4.2 Frozen public-history sample

The new timestamp study uses documented, unauthenticated public funding-history endpoints at Hyperliquid, Pacifica, Extended and edgeX. The target window begins at 2026-09-01 00:00:00 UTC and ends at 2026-11-03 20:00:00 UTC. The chronological split is 2026-10-09 00:00:00 UTC, chosen to align the calibration/holdout boundary with the transport-audit date rather than with an observed offset outcome. Assets are BTC, ETH and SOL, selected before analysis for common cross-venue availability. The design does not claim coverage of all contracts, all assets, or all trading venues.

Table 1 describes the frozen sample. Hyperliquid and Extended each contribute 4,596 records and 1,532 venue-hours. Pacifica contributes 4,587 records and 1,529 venue-hours; three nominal hours are absent between its first and last observations and are not imputed. edgeX contributes 1,152 settlement records and 384 four-hour events. All four series cover 64 UTC dates, but their last available timestamps at collection differ. That difference is reported as the response snapshot found at retrieval and is not interpreted as a publication-latency ranking.

Table 1: Frozen public-history sample and internal completeness. Records are asset-level rows; venue-hours are the primary analysis unit. Source: public APIs collected 2026-11-03 UTC.

Venue	Records	Venue-hours	Days	Gaps	Public time field and record rule
Hyperliquid	4,596	1,532	64	0	`time`; historical funding records
Pacifica	4,587	1,529	64	3	`created_at`; history-record timestamp
Extended	4,596	1,532	64	0	`T`; calculated and applied hourly rate
edgeX	1,152	384	64	0	`fundingTime`; `isSettlement=true`, four-hour grid
Total	14,931	4,977	64	3	BTC, ETH and SOL; no missing-value imputation

Hyperliquid's `fundingHistory` request is collected in seven-day windows using the docu-

mented inclusive millisecond boundaries (Hyperliquid, ndb). Pacifica's cursor-based endpoint is followed backwards until the sample start is crossed (Pacifica, ndb). Extended documents a 10,000-record maximum, but the observed responses contained at most 1,000 rows and no pagination object. The collector therefore pages backward using disjoint timestamp windows; this is a documented adaptation to the actual response, not an assumption that the first 1,000 rows are complete (Extended, nd). edgeX contract identifiers are read from current public metadata rather than copied from an old documentation example, and only rows explicitly marked as settlements enter the analysis (edgeX, nd).

The target window is a collection frame rather than a guarantee that every venue has already published its final target-hour record. Hyperliquid, Pacifica and Extended end at 19:00 UTC in the frozen data, while edgeX has a 20:00 UTC settlement. Internal completeness is consequently evaluated between each venue's first and last observed events, not by declaring the unequal last hours to be missing. This choice avoids turning retrieval timing into a false clock anomaly.

4.3 Pilot observations and failed refreshes

An earlier pilot dated 2026-10-26 contained 24 Hyperliquid observations and three observations each for Pacifica, Extended, Aster, ApeX and edgeX. It also contained three Paradex rate-snapshot timestamps with negative offsets. Those values motivated the larger design but are not pooled with the new sample. Paradex remains a semantic counterexample: a rate snapshot created before the hour must not be relabeled a settlement payment. The pilot values are included in a separate CSV with an explicit `usage` field so that no downstream script mistakes them for independent large-sample evidence.

The new collector attempted to refresh Aster and ApeX for the same three assets using their documented public history endpoints (Aster, nd; ApeX, nd). All six requests ended in connection-reset errors. These failures are stored in the request manifest and failure report. They establish only that the current measurement environment could not retrieve fresh responses from the documented public endpoints. They do not show that either venue lacks public history, that its timestamp is exactly aligned, or that its earlier three-row pilot is representative. No indirect aggregator, cached search excerpt or access-control workaround is substituted for the missing raw evidence. Accordingly, Aster and ApeX appear in the transport matrix and limitations but not in the new offset distribution or holdout tables.

4.4 Normalization and observation rules

Normalization preserves integer-millisecond arithmetic. For each returned timestamp t, the nominal hour is computed by integer floor division and the offset is $t - h$. Records are deduplicated by venue, asset, timestamp and public time field. No offset is rounded to seconds, no missing timestamp is replaced with the nominal hour, and no large offset is winsorized. Funding rates are retained for context but are not inputs to the offset estimators.

The primary dataset then groups asset rows by venue and nominal hour. Its offset is the median across available assets, and accompanying fields record the asset range and count. This choice preserves the complete asset-level file while recognizing that timestamps are often common across assets. Hyperliquid and Extended have identical offsets across all three assets at every observed event; edgeX settlement timestamps are likewise shared. Pacifica occasionally differs by 1 ms across assets. A record-weighted analysis would as-

sign three times as much weight to these shared mechanisms as a single independent observation deserves.

No generic "outlier removal" is applied. Extended's extreme 2026-09-24 event is retained in every full-sample statistic and marked separately in tables and figures. A conditional figure zooms into offsets of at most 6,000 ms for readability, but the caption states that this is a display restriction rather than sample deletion. The three internal Pacifica gaps remain missing; their presence is included in the sample-frame table rather than silently filled by adjacent offsets.

4.5 Provenance, documentation and reproducibility

The collection contains 58 request-manifest records, including the six explicit Aster and ApeX failures. Each entry stores the public URL, request body where applicable, UTC retrieval time, HTTP status when received, client-observed elapsed time, stored response path and SHA-256 hash. Fifty-eight stored response or error bodies pass the independent hash verification. Client elapsed time is provenance, not a multi-location latency measurement.

The reproduction package also contains normalized asset and venue-hour CSV files, official documentation snapshots with hashes, a data dictionary, machine-readable tables and the analysis script. Offline analysis regenerates every headline count, quantile, bootstrap interval, holdout result and vector figure without contacting an exchange. A new online collection would form a new snapshot and is not expected to reproduce the frozen data if a venue extends or revises its history.

Access was limited to public no-key interfaces. No user data, credentials, private endpoints or nonpublic infrastructure are included. The release does not provide instructions for evading a regional restriction or access-control response. These restrictions affect coverage and are therefore part of the data-generating process rather than a footnote external to it.

5 Measurement methods

5.1 Clock definitions and the primary estimand

Let t_{vaj} denote the integer-millisecond timestamp returned by venue v for asset a and event j. All timestamps are interpreted as Unix epoch milliseconds in UTC. The nominal hour containing the timestamp is

$$h_{vaj} = H \left\lfloor \frac{t_{vaj}}{H} \right\rfloor, \qquad H = 3{,}600{,}000, \tag{4}$$

and the within-hour offset is

$$d_{vaj} = t_{vaj} - h_{vaj}, \qquad 0 \le d_{vaj} < H. \tag{5}$$

This definition follows the supplied measurement material exactly: the reference is the beginning of the hour containing the record, not the closest hour. A record 18 minutes after 01:00 therefore has an offset of 1,080,000 ms, not a negative offset relative to 02:00. edgeX events retain their actual four-hour spacing; equation (4) merely provides a common millisecond reference and does not convert them into hourly observations.

For venue v and nominal event j, the primary offset is the median across available assets:

$$d_{vj} = \text{median}_{a \in \mathcal{A}_{vj}} \, d_{vaj}, \tag{6}$$

where $\mathcal{A}_{vj} \subseteq \{\text{BTC}, \text{ETH}, \text{SOL}\}$. The analysis set $\mathcal{J}_v$ contains one observation per venue-nominal-hour. I also retain $N_{vj} = |\mathcal{A}_{vj}|$, the minimum and maximum asset offsets, and the record count. This event-level estimand answers: "how far after the nominal hour is a typical venue event timestamp?" It does not answer: "how delayed is this data when it reaches a client?"

For each venue, let $Q_v(p)$ be the empirical quantile of $\{d_{vj} : j \in \mathcal{J}_v\}$ at probability p, computed with linear interpolation. I report $Q_v(0)$, selected lower quartiles, $Q_v(0.50)$, $Q_v(0.95)$, $Q_v(0.99)$, and $Q_v(1)$. Tail shares are defined as

$$S_v(x) = \frac{1}{|\mathcal{J}_v|} \sum_{j \in \mathcal{J}_v} \mathbf{1}(d_{vj} > x), \tag{7}$$

so an observation exactly equal to x is not counted as exceeding it. This strict inequality is used consistently in the missing-clock framework.

5.2 Resampling over dates rather than treating hours as independent

Venue-hours can be autocorrelated, and assets within an hour often share one timestamp. Conventional independent-and-identically-distributed standard errors would therefore misstate the sampling design. For descriptive uncertainty, I group all event observations by UTC date and apply a moving-block bootstrap with seven-day blocks and 2,000 replicates (Künsch, 1989). Dates are sampled circularly with replacement until the original number of dates is obtained; all venue-hours within each selected date are retained. The random seed is 20261103.

For each replicate b, I compute the median and 95th percentile of the resampled offsets. The reported 95% interval is

$$\left[\widehat{Q}^*_{0.025}(\theta_v), \widehat{Q}^*_{0.975}(\theta_v)\right], \tag{8}$$

where $\theta_v \in \{Q_v(0.50), Q_v(0.95)\}$ and $\widehat{Q}^*_p$ is the quantile across bootstrap replicates. These intervals describe sensitivity under the stated date-resampling scheme. They do not prove stationarity, eliminate structural breaks, or provide a venue guarantee. No Newey–West regression standard error is reported because there is no regression or mean-return estimand to which it would apply.

5.3 A frozen chronological tolerance diagnostic

To examine temporal transfer without optimizing on later observations, I split events at 2026-10-09 00:00 UTC. Calibration events $\mathcal{C}_v$ precede the split; holdout events $\mathcal{H}_v$ occur on or after it. For each venue independently, the research tolerance is

$$\tau_v = Q_{\mathcal{C}_v}(0.99) + 1 \text{ ms}. \tag{9}$$

The additional millisecond is a fixed numerical allowance, not a fitted production buffer. There is one rule per venue, no grid search, and no selection based on holdout coverage. The holdout coverage and exceedance count are

$$c_v = \frac{1}{|\mathcal{H}_v|} \sum_{j \in \mathcal{H}_v} \mathbf{1}(d_{vj} \le \tau_v), \qquad E_v = \sum_{j \in \mathcal{H}_v} \mathbf{1}(d_{vj} > \tau_v). \tag{10}$$

A 95% Wilson interval is reported around c_v as a finite-sample descriptive summary (Wilson, 1927). Because hourly observations within a date need not be independent, the Wilson interval is not presented as a fully dependent-data confidence statement; the date-block bootstrap and chronological exceedance counts remain the primary diagnostics.

Equation (9) is deliberately simple. It asks whether a transparent empirical rule survives a later period, not whether it is optimal for event assignment. A future implementation could use a more structured model with multiple candidate intervals, regime detection and an explicit unresolved state, but that would be a different estimator requiring its own out-of-sample evaluation. The present study does not promote τ_v as a recommended production threshold.

5.4 Event-assignment interpretation

A history timestamp can be assigned to the nominal hour in equation (4) only under a declared tolerance and mismatch policy. In this paper, an event with $d_{vj} \leq \tau_v$ is described as covered by the research tolerance. An event with $d_{vj} > \tau_v$ is an exceedance, not an observation to be silently forced into the hour. The extreme Extended record illustrates the distinction: relabeling 01:18:16.849 as exactly 01:00 would erase 18 minutes of observable discrepancy, whereas deleting the record would hide a failure of the simple rule.

This mismatch policy is conceptual rather than a production rule. It has three properties. First, the original timestamp is always retained. Second, an exceedance remains visible even when the bulk of the distribution is tight. Third, no alternative candidate hour is selected merely because it improves coverage. These properties prevent the evaluation from becoming a search for the most favorable alignment.

5.5 Partial identification when clocks are missing

Consider a population of A two-legged pairs that are eligible for a clock-difference test at a common reference time. For pair i, let t_{i1} and t_{i2} be the two source clocks and define the binary outcome

$$Y_i(\tau) = \mathbf{1}\left(|t_{i1} - t_{i2}| > \tau\right). \tag{11}$$

Let $R_i = 1$ when both clocks are present and valid, and $R_i = 0$ otherwise. The observed counts are

$$n = \sum_{i=1}^{A} R_i, \qquad m = \sum_{i=1}^{A} (1 - R_i), \qquad o = \sum_{i=1}^{A} R_i Y_i(\tau), \tag{12}$$

with $A = n + m$. The target population share is

$$\theta(\tau) = \frac{1}{A} \sum_{i=1}^{A} Y_i(\tau). \tag{13}$$

For $A > 0$, every missing pair has $Y_i \in \{0, 1\}$, so the sharp worst-case identified set is

$$\boxed{\theta(\tau) \in \mathcal{I}(\tau) = \left[\frac{o}{A}, \frac{o+m}{A}\right]}. \tag{14}$$

The lower bound assigns zero exceedances to every missing pair; the upper bound assigns one to every missing pair. Both extremes are logically possible without additional assumptions, so the bounds are sharp. Their width is

$$\text{width}\, \mathcal{I}(\tau) = \frac{m}{A}. \tag{15}$$

When $A = 0$, equations (13) and (14) are undefined. The correct output is "not applicable," never zero.

The complete-case proportion o/n answers a different question. It is defined only when $n > 0$ and describes valid-clock pairs rather than the full population. It can be rewritten as a weighted mixture of the complete-case rate and the unobserved missing-case rate. Assuming that the missing-case rate equals o/n yields the complete-case point estimate, but that is an additional assumption rather than a consequence of the data. The bounds in equation (14) deliberately avoid it.

For illustration, $A = 8, n = 5, m = 3, o = 2$ gives $\mathcal{I} = [0.25, 0.625]$, whereas the complete-case proportion is 0.40. A subpopulation with $A = 2, n = 1, m = 1, o = 1$ gives $[0.5, 1]$. If the same venue contributes disproportionately to m, the aggregate interval width is driven by that venue's missingness; a venue-by-missingness contingency table is therefore more informative than a single complete-case rate. The numerical figure in the paper uses synthetic counts only and contains no internal route population or production threshold.

5.6 Relation between identified sets and statistical intervals

Equation (14) is an accounting result for a finite observed population. It is not a 95% confidence interval, a Bayesian credible interval, or a tolerance interval for future pairs. Sampling uncertainty could be added if there were a defensible sampling design, but it would not remove the identification problem. Conversely, the bootstrap intervals for timestamp offsets concern observed public-history distributions and have no role in narrowing equation (14).

The distinction can be expressed by writing the population share as

$$\theta(\tau) = \frac{n}{A}\frac{o}{n} + \frac{m}{A}\theta_m(\tau), \qquad \theta_m(\tau) = \frac{1}{m}\sum_{i:R_i=0} Y_i(\tau), \tag{16}$$

when $n, m > 0$. The first term is observed, while $\theta_m(\tau) \in [0, 1]$ is unidentified. Missing-at-random would restrict the relation between observed and missing outcomes; the present framework instead reports what follows from $\theta_m = 0$ and $\theta_m = 1$.

5.7 Reproducible implementation

The collector stores raw bytes before normalization. The analysis script recomputes equations (4)–(15), emits all tables and figures, and verifies exact hashes and integer identities. The verifier additionally evaluates the partial-identification inequalities over 23,425 small-integer configurations, including both missing-outcome extremes. This is a finite exhaustive code check, not a substitute for the algebraic proof. Public collection failures, internal gaps and the large Extended anomaly remain visible in the outputs. All research tolerances are labeled as such and do not represent any live system setting.

6 Results

6.1 Public WebSocket capability is broad but heterogeneous

Figure 1 summarizes the 22-venue audit. Eighteen venues have a funding-related WebSocket route in the audited surface. Nine have a full all-market funding channel, while

MEXC has partial all-market coverage because its all-market ticker carries a fair-price field but not the funding rate. Four venues have a material gap: BingX lacks a funding WebSocket despite offering price and book channels; LBank has a connectable address but no documented usable public channel in the audit; GMX has neither a conventional WebSocket market-data route nor a central order book; Variational exposes REST and request-for-quote information rather than a conventional order-book stream.

These categories should not be read as a single quality ranking. BingX's missing funding channel can be addressed only through a different transport with its own polling and cache semantics. GMX and Variational require different market models rather than a WebSocket adapter. A REST fallback may be entirely appropriate for a slowly changing reference value but unsuitable for reconstructing a fast local order book. The matrix therefore answers whether a particular measurement architecture is feasible, not whether a venue is generally better or worse.

Figure 1. Public WebSocket capability differs by venue and field semantics

	Funding	Price	Order book	All-market funding
Binance	+	+	+	+
OKX	+	+	+	–
Bybit	+	+	+	–
Gate	+	+	+	+
Bitget	+	+	+	–
MEXC	+	+	+	P
KuCoin	+	+	+	+
BingX	–	+	+	–
LBank	–	–	–	–
Deribit	+	+	+	–
Hyperliquid	+	+	+	–
Lighter	+	+	+	+
dYdX	+	P	+	+
Paradex	+	+	+	–
Aster	+	+	+	+
Extended	+	+	+	+
ApeX	+	+	+	–
edgeX	+	P	+	+
GRVT	+	+	+	–
Pacifica	+	+	+	+
GMX	–	–	–	–
Variational	–	–	–	–

+ available P partial / different price semantics / partial all-market coverage – unavailable
Source: public documentation and single-host checks dated 2026-10-09. Capabilities are archival observations, not permanent guarantees.

Figure 1: Public WebSocket capabilities by venue. "P" denotes partial all-market coverage or oracle/index rather than mark-price semantics. Source: de-identified public documentation and single-host audit dated 2026-10-09. Capabilities are archival observations, not permanent guarantees or multi-location latency estimates.

All-market reach is similarly heterogeneous. Binance, Gate, KuCoin, Lighter, dYdX, Aster, Extended, edgeX and Pacifica have full all-market funding-related subscriptions in the

audit. Gate's `!all` behavior was observed but not documented on the reviewed page, so it is labeled as an observation rather than an official stability guarantee. KuCoin's all-market funding channel and Extended's funding stream have approximately minute-level source update behavior in the audit. A client's failure to receive a new value every few seconds on such a channel cannot be interpreted in the same way as silence on a millisecond order-book stream.

Price fields require separate treatment. dYdX and edgeX are marked partial in the price column because the reviewed channels expose oracle or index information rather than a field directly identified as mark price. This is not an assertion that the prices are unsuitable for every purpose. It is a warning that a common schema must preserve the original meaning before using the values in funding, margin or liquidation calculations. The replication table supplies venue-specific notes and public documentation URLs for these cases.

6.2 Timestamp offsets differ by orders of magnitude

Table 2 reports the primary venue-hour estimates. Hyperliquid has a median offset of 37 ms and a 95th percentile of 88 ms; all 1,532 observed venue-hours fall between 0 and 139 ms. Its seven-day block-bootstrap interval is 35–38 ms for the median and 83–94 ms for the 95th percentile. This is a narrow distribution around a small positive offset, not proof of zero clock error.

Pacifica has a median of 2,513 ms and a 95th percentile of 5,083 ms. Its offsets range from 397 to 5,084 ms, and 76.78% of venue-hours are more than two seconds after the nominal hour. The distribution contains visible level changes rather than one tightly concentrated point. The median's date-block interval, 2,114–3,524.01 ms, is correspondingly wide. The interval reflects resampled daily compositions and should not be read as a promise that future record-creation delays remain inside it.

Extended has a median of 849 ms and a 95th percentile of 1,593.15 ms. Its 99th percentile is 1,821.83 ms, so 99% of venue-hours remain below approximately 1.83 seconds. Nevertheless, the maximum is 1,096,849 ms. edgeX is exactly aligned for every one of its 384 settlement events: minimum, median, 95th percentile, 99th percentile and maximum are all zero. This is an exact description of the returned settlement field in the sample, not an estimate of network latency or a claim about non-settlement calculation timestamps.

Table 2: Offset distribution after the nominal UTC hour, in milliseconds. Unit: venue-hour median across BTC, ETH and SOL. Bootstrap intervals use 2,000 seven-day moving-block replicates over UTC dates. Source: frozen public histories.

Venue	Hours	Median	Median 95% interval	p95	p99	Maximum
Hyperliquid	1,532	37	[35, 38]	88	121	139
Pacifica	1,529	2,513	[2,114, 3,524.01]	5,083	5,083	5,084
Extended	1,532	849	[849, 937]	1,593.15	1,821.83	1,096,849
edgeX	384	0	[0, 0]	0	0	0

Figure 2 shows the full empirical cumulative distributions and an ordinary-range detail. The horizontal axis uses a symmetric-log scale in the full panel so that edgeX's zero, Hyperliquid's tens of milliseconds, Pacifica's seconds and Extended's extreme event can coexist. The second panel conditions only the display on offsets at or below 6,000 ms; it does not alter the estimates in Table 2. The four curves are inconsistent with a single cross-venue clock-offset distribution.

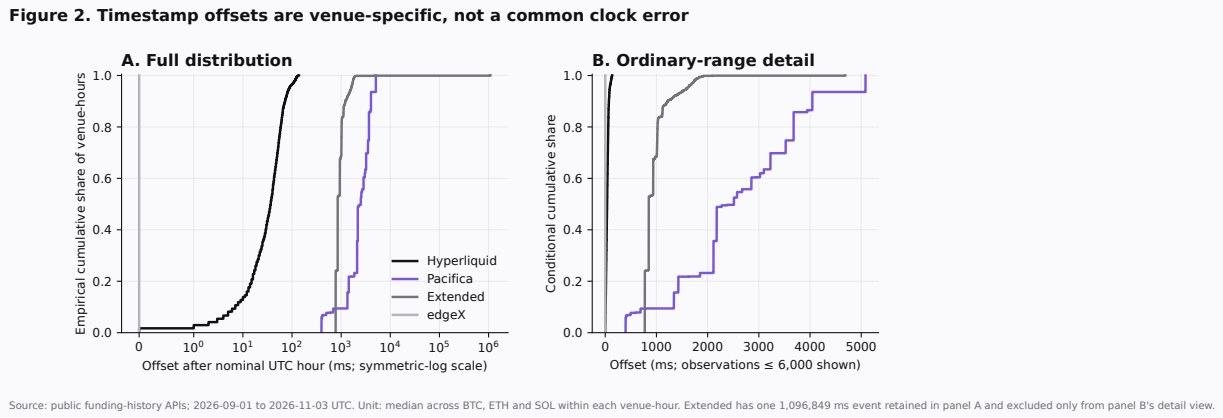


Figure 2: Empirical cumulative distributions of venue-hour offsets. The full panel retains Extended's extreme observation; the detail panel is only a visual restriction to offsets at most 6,000 ms. Source: public funding-history APIs, 2026-09-01 to 2026-11-03 UTC.

6.3 Daily paths reveal regime variation

Figure 3 plots daily medians on independent vertical scales. Hyperliquid fluctuates within tens of milliseconds and does not show a shift comparable to Pacifica's step changes. Pacifica moves among several approximately stable levels, including periods around 1.4, 2.1, 2.5, 3.7 and 5.1 seconds. These levels are observations of record behavior, not explanations of their cause. They could reflect scheduling or implementation changes, but public history alone cannot identify the responsible mechanism.

Extended spends much of the sample near 0.8–1.0 seconds, with intervals closer to 1.6 seconds and one large daily median around the extreme event. edgeX remains exactly zero on the settlement field. The independent vertical scales are necessary for readability but prohibit a visual comparison of slope magnitudes across panels; Table 2 supplies the common numerical scale.

The Extended anomaly occurs at 2026-09-24 01:18:16.849 UTC in all three assets. It is one of 1,532 venue-hours, or 0.0653% of the Extended event sample. Its presence has almost no effect on the median but dominates the maximum and any naive mean. We do not infer an outage, a clock-synchronization failure or a correction from one record. The appropriate conclusion is narrower: a rule calibrated to the ordinary two-second tail cannot represent every observed record, and an explicit unresolved or exceedance category is necessary.

6.4 Chronological holdout exposes tolerance failure

Table 3 and Figure 4 report the frozen chronological diagnostic. Hyperliquid's calibration 99th percentile is 121.89 ms, producing a tolerance of 122.89 ms. Four of 620 later venue-hours exceed it, for coverage of 99.35% and a Wilson descriptive interval of 98.35–99.75%. Pacifica's tolerance is 5,084 ms; all 620 later events are covered. edgeX's tolerance is 1 ms; all 156 later settlement events remain exactly aligned.

Extended provides the counterexample. Its calibration 99th percentile is 1,744.45 ms, so the frozen tolerance is 1,745.45 ms. Twenty-one of 620 holdout venue-hours exceed it, giving 96.61% coverage and a Wilson interval of 94.88–97.77%. The full-sample 99th percentile subsequently rises to 1,821.83 ms. A researcher who evaluated only on the calibration sample would therefore overstate the rule's later coverage. This is a temporal transport failure of the tolerance, not evidence of a trading loss or a venue malfunction.

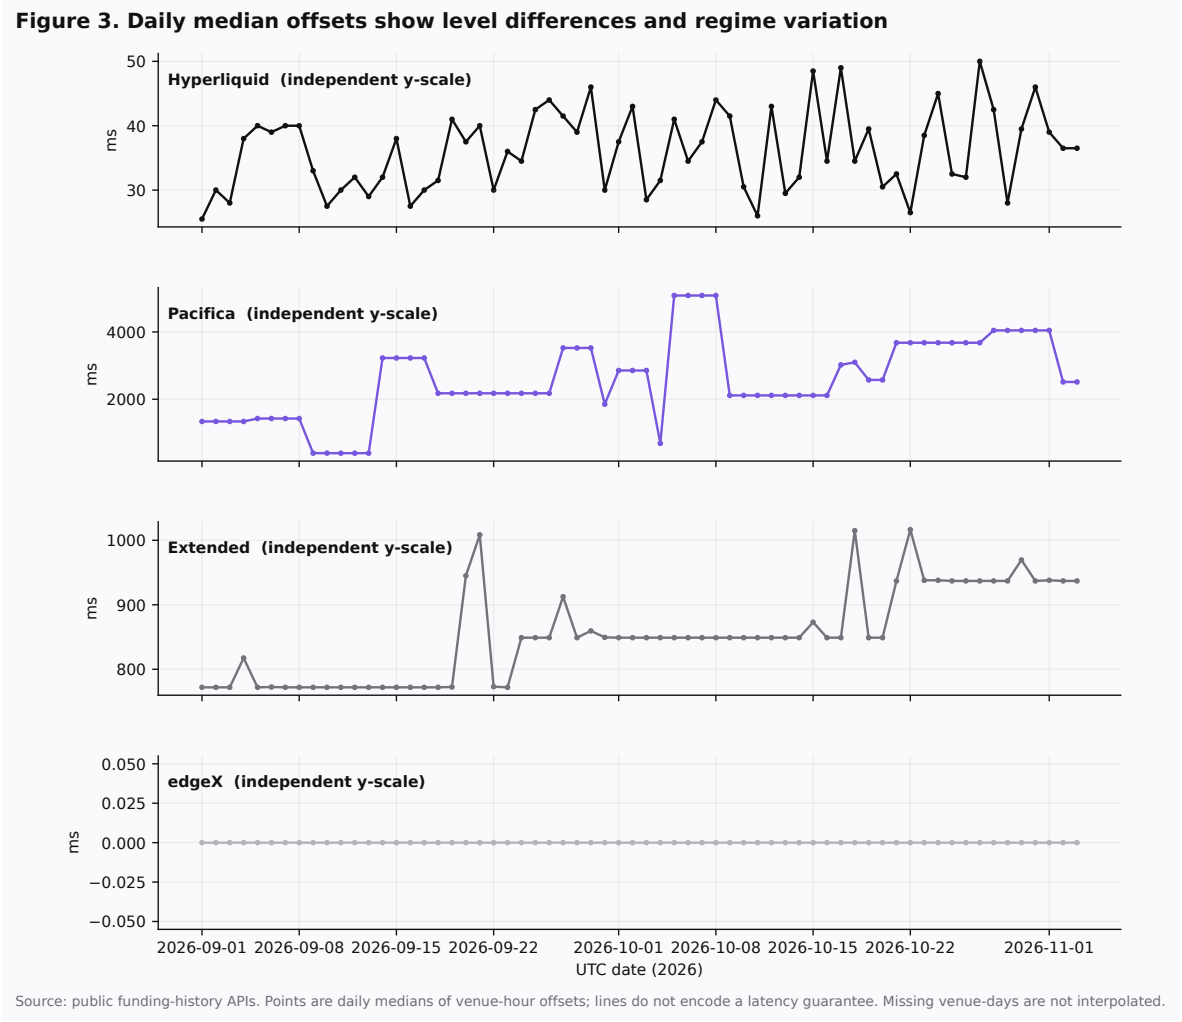


Figure 3: Daily median venue-hour offsets. Each panel has an independent vertical scale; missing days are not interpolated. Source: public funding-history APIs.

Table 3: Chronological holdout results. The tolerance is the calibration empirical 99th percentile plus 1 ms, fixed before later observations are scored. Coverage intervals are Wilson descriptive intervals, not future-performance guarantees.

Venue	Tolerance (ms)	Holdout hours	Exceedances	Coverage	Wilson 95% interval
Hyperliquid	122.89	620	4	99.35%	[98.35%, 99.75%]
Pacifica	5,084.00	620	0	100.00%	[99.38%, 100.00%]
Extended	1,745.45	620	21	96.61%	[94.88%, 97.77%]
edgeX	1.00	156	0	100.00%	[97.60%, 100.00%]

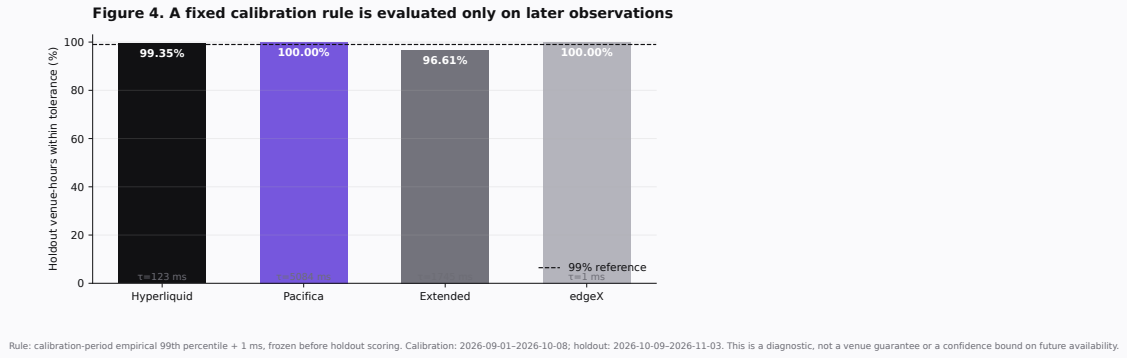


Figure 4: Coverage of a fixed calibration-period tolerance on later venue-hours. The dashed line is a 99% reference, not a target promised by any venue. Source: calibration observations through 2026-10-08 and holdout observations from 2026-10-09.

The four outcomes also demonstrate why coverage cannot be compared independently of scale. edgeX's 1 ms tolerance and Pacifica's 5,084 ms tolerance both attain 100% holdout coverage, but they imply very different event-assignment resolutions. A high coverage rate obtained with a wide tolerance is not equivalent to precise synchronization. The tolerance, exceedance count and underlying offset distribution must therefore be reported together.

6.5 Missing clocks produce visible identification width

Figure 5 illustrates equation (14) with synthetic counts. As the number of missing-clock pairs increases, the complete-case proportion remains a single line by construction, while the sharp population interval widens linearly. The upper bound rises because missing pairs may all exceed τ; the lower bound falls because they may all fall below it. The interval width is exactly m/A.

This result changes reporting practice even when no point estimate is available. If a report states only o/n, a reader cannot tell whether the full-population conclusion is narrow or nearly vacuous. Reporting A, n, m, o and the corresponding bounds preserves that information. If $A = 0$, the appropriate result is not a zero exceedance rate but an unavailable estimand. No empirical pair counts are used in this illustration.

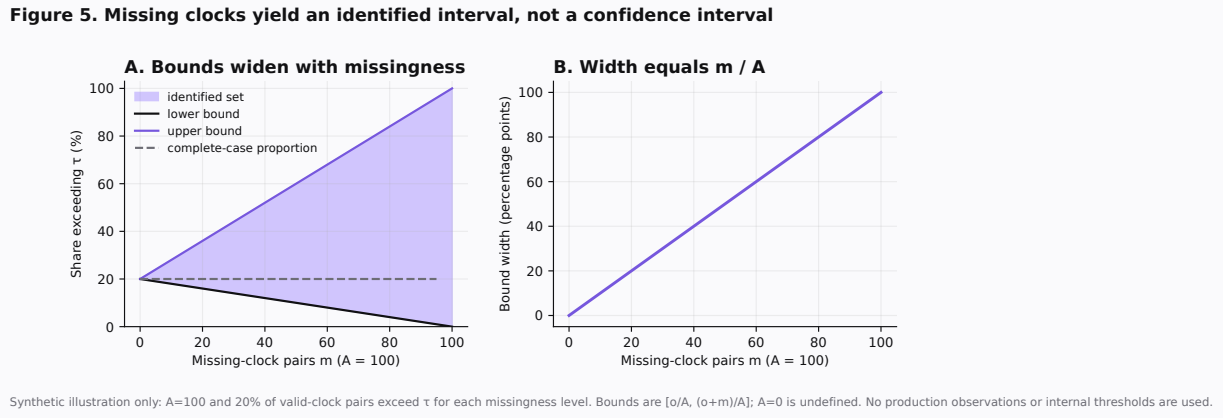


Figure 5: Synthetic partial-identification example with $A = 100$ and 20% of valid-clock pairs exceeding τ at each missingness level. Bounds are not confidence intervals and do not use production data.

6.6 Interpretation

The combined results support three practical measurement conclusions. First, transport availability must be assessed at the field and semantic level, not inferred from the existence of a WebSocket endpoint. Second, timestamp alignment requires venue-specific empirical distributions, an anomaly policy and a chronological check; one venue's milliseconds cannot be borrowed as another venue's tolerance. Third, missing clocks restrict identification in proportion to their population share. These conclusions concern data validity and interpretability. They do not establish that a trading strategy is profitable, that a venue is unreliable, or that a particular threshold should govern a live product.

7 Robustness and evidence boundaries

7.1 Weekly and calendar-subperiod checks

The first robustness check asks whether the headline levels are driven by one short interval. The sample is divided into ten ISO weeks and into the prespecified calibration and holdout periods. Weekly medians range from 34 to 42 ms for Hyperliquid, 400 to 5,083 ms for Pacifica, 772 to 937 ms for Extended, and remain zero for edgeX. Weekly 95th percentiles range from 63 to 105.95 ms, 1,427 to 5,083 ms, 1,257.45 to 1,701.20 ms, and zero, respectively. The full weekly table is released rather than selecting a favorable week for display.

The two calendar periods provide more detail. Hyperliquid's median changes from 36.5 ms in calibration to 37 ms in holdout, while its 95th percentile changes from 90 to 83.05 ms. Pacifica's median rises from 2,177 to 3,097 ms, but its 95th percentile falls from 5,083 to 4,046 ms because the later distribution has a different mixture of levels. Extended's median rises from 849 to 937 ms and its 95th percentile from 1,520.85 to 1,643.60 ms. These shifts reinforce the conclusion that a full-sample quantile is a summary of a time-varying mixture, not a stable physical constant.

Pacifica also illustrates why the median and upper tail need not move together. A period can contain more observations at intermediate levels while fewer observations occupy the highest level. Reporting only one quantile would conceal that compositional change. The daily paths, weekly estimates and chronological exceedance counts are therefore complementary rather than redundant.

7.2 Asset-specific estimates

The second check recomputes distributions separately for BTC, ETH and SOL. Hyperliquid and Extended have identical offsets across the three assets at every observed event, so their asset-specific quantiles are necessarily identical. edgeX settlement timestamps are likewise shared. Pacifica's median, 95th percentile and maximum are also identical across assets: 2,513, 5,083 and 5,084 ms. The only minimum-level distinction is 396 ms for ETH versus 397 ms for BTC and SOL.

These results do not justify treating the three assets as independent replications. They indicate the opposite: the timestamp mechanism is largely venue-level. The asset check is useful because a hidden asset-specific mapping error would be visible, but the near identity confirms the importance of the venue-hour estimand. All asset-level results remain available so that a future schema change affecting only one asset can be detected.

7.3 Record, event and day weighting

The third check varies the weight given to records and dates. Record-weighted medians and venue-hour medians are identical for all four venues: 37, 2,513, 849 and 0 ms. This equality follows from the nearly complete three-asset coverage and shared event timestamps. It should not be read as evidence that weighting is irrelevant in general; a venue with unequal asset coverage could produce a different record-weighted result.

Equal-day summaries change the interpretation more materially. The median of daily median offsets is 36.5 ms for Hyperliquid, 2,345.5 ms for Pacifica, 849 ms for Extended and zero for edgeX. The mean of daily 95th percentiles is 80.98, 2,696.82, 1,380.65 and 0 ms, respectively, compared with event-weighted 95th percentiles of 88, 5,083, 1,593.15 and 0 ms. Pacifica's difference is especially large because long runs at a high offset receive more weight in the event distribution than in an average of daily tail statistics.

Neither weighting is universally correct. Event weighting answers how often a timestamp drawn from the observed event population has a given property. Equal-day weighting answers how a typical UTC date behaves, regardless of whether a regime persisted for many hours. The paper's headline uses the former and discloses the latter. A consumer interested in daily monitoring outcomes may prefer the second; a consumer auditing every event should retain the first.

Table 4: Alternative weighting schemes, in milliseconds. Equal-day tail is the mean of daily 95th percentiles. Source: frozen public histories.

Venue	Record median	Venue-hour median	Median of daily medians	Equal-day tail
Hyperliquid	37	37	36.5	80.98
Pacifica	2,513	2,513	2,345.5	2,696.82
Extended	849	849	849	1,380.65
edgeX	0	0	0	0

7.4 Bootstrap block-length sensitivity

The fourth check changes the moving-block length from seven days to three and 14 days while retaining the same 2,000 replicates and seed. Hyperliquid's median interval remains [35, 38] ms to rounding for all three choices, and its 95th-percentile interval remains approximately [83, 94] ms. Extended's median interval is [849, 937] ms for three- and seven-day blocks and [773, 937] ms for 14-day blocks. Pacifica is more sensitive: the median interval is [2,114, 3,226] ms for three-day blocks, [2,114, 3,524.01] ms for seven-day blocks, and [1,428, 3,679] ms for 14-day blocks. edgeX remains degenerate at zero.

The Pacifica sensitivity is consistent with regime persistence. Longer blocks preserve longer runs of similar daily observations and therefore produce a wider range of resampled compositions. This is not a reason to select the block length that yields the most desirable interval. The paper retains seven days for the headline and reports the alternatives as evidence that serial dependence matters.

7.5 Semantic and record-selection checks

The fifth check concerns what is counted, not how a number is averaged. For edgeX, only `isSettlement=true` rows enter the main sample. Including minute-level calculation rows

would mix `fundingTime` and `fundingTimestamp` concepts and create a much larger but semantically different dataset. For Paradex, the three pilot rate snapshots remain outside the settlement-offset estimates because their negative offsets demonstrate that a snapshot can precede the nominal hour. Excluding these records is not outlier deletion; it is preservation of record-type boundaries.

The same principle applies to transport. A mark-price channel is not substituted with an oracle-price channel without a notation change, and a request-for-quote tier is not counted as an order-book level. A REST market-information response is not labeled a WebSocket observation. These classification decisions prevent apparently robust numerical results from being obtained by combining incompatible objects.

7.6 Missingness and threshold-boundary checks

The missing-clock bounds require no distributional assumption and are therefore robust by construction. For fixed A, n, m, o, increasing the strict threshold τ cannot create new exceedances: a pair with $|t_{i1} - t_{i2}| > \tau_2$ also exceeds every smaller $\tau_1 < \tau_2$. Both bounds are consequently nonincreasing in τ. A pair exactly equal to τ never counts as exceeding it. The verifier checks the lower-bound, upper-bound and width identities over 23,425 small-integer configurations and treats $A = 0$ separately as undefined.

These are semantic checks rather than evidence about any empirical pair population. No production threshold or real route count enters the synthetic figure. A future empirical application must report how A, n, m, o change with τ, including the complete-case count, rather than presenting one favorable threshold in isolation.

7.7 Why forecast, multiple-testing and trading-cost tests are not imported

The paper contains no predicted funding rate, fitted return model or model-selection exercise. The chronological rule is one fixed descriptive diagnostic per venue, and all robustness scenarios are reported rather than selected for statistical significance. Diebold–Mariano forecast-loss tests, White reality checks, Hansen superior-predictive-ability tests, Harvey–Liu–Zhu return-factor adjustments and backtest-overfitting probabilities therefore do not apply to the reported estimands. Applying their names without their required forecasts or return series would create a false impression of validation.

For the same reason, there is no per-$10,000 or annualized economic-significance calculation. Fees, bid–ask spreads, depth-dependent impact, transfer costs and time, borrowing costs, liquidation costs and actual fills would be indispensable before any return claim. The present study measures public data capabilities and timestamps. Its contribution is to make a later economic analysis less vulnerable to false synchronization, not to prejudge that analysis's profitability.

8 Limitations

8.1 Transport coverage and latency evidence

The transport matrix is a single-day, single-host audit. Proxy routing, regional access, connection setup and subscription scheduling affect first-frame observations. The study does not include repeated measurements from multiple geographic locations and therefore

cannot estimate a first-frame latency distribution, separate venue processing from network conditions, or rank venues by responsiveness. A future multi-location study should retain independent connection attempts, synchronized client clocks, failure counts and location labels. The present matrix is useful for capability mapping only.

The audit is also time-sensitive. Public channels, limits and schemas can change after 2026-10-09. An undocumented behavior observed on that date may later be documented or removed; an inaccessible documentation page does not prove that a capability is absent everywhere. I preserve these uncertainties in the evidence labels and do not convert a failed page retrieval into a universal statement about a venue.

8.2 Venue, asset and historical coverage

The large-sample offset analysis covers four venues and three large assets. Aster and ApeX could not be refreshed because of connection resets, and their three-row pilot observations are not generalized. The four measured venues are not a random sample of the 22-venue transport map. Consequently, the offset results cannot be extrapolated to the other 18 venues, to smaller assets, or to every contract on the measured venues.

Public history is also a revisionable source. An endpoint may later add, remove or correct records. The raw-response hashes prove what this study received at the stated collection times, not that a provider will always serve identical bytes. Internal completeness is measured between the first and last observed events; it cannot detect records missing before the first or after the last observation. Three internal Pacifica gaps remain unexplained, and unequal final hours are not treated as proof of publication delay.

8.3 Timestamp semantics and clock identification

An offset after the nominal hour is not physical clock skew. It can combine intentional scheduling, computation, record writing, caching and provider-specific timestamp semantics. Pacifica's `created_at` is explicitly a record timestamp, so its seconds-level offsets must not be relabeled settlement-engine delays. edgeX's exact zeros apply only to the analyzed settlement field; non-settlement calculation timestamps and client receipt latency can be nonzero. Hyperliquid's narrow historical offsets similarly say nothing about a consumer's WebSocket latency.

The study has no independent trusted timestamp inside the venue's infrastructure and no source-receipt round-trip series that would identify synchronization error. NTP or PTP on the client would improve the client's own clock discipline but would not separate these source-side components retrospectively. The paper therefore uses "observed timestamp offset" throughout rather than claiming to measure an exchange's clock error.

8.4 Regime dependence and anomaly interpretation

The date-block bootstrap depends on a block-length choice and a stationary-resampling intuition that regime changes can violate. The reported three-, seven- and 14-day alternatives show material sensitivity for Pacifica. The Wilson intervals around holdout coverage do not fully model within-day dependence and are included as descriptive finite-sample summaries only. One chronological split cannot establish future stability, and the fixed 99th-percentile-plus-1-ms rule is not externally preregistered as a production forecast.

The extreme Extended record is one observation. It demonstrates that an ordinary-tail tolerance can fail, but it cannot identify why the record occurred. Possible explanations include scheduling, record generation or an unusual operational event; the public response does not distinguish them. Removing the record would be unjustified, but using it alone to characterize the venue would be equally misleading.

8.5 Missing-clock bounds and empirical scope

The partial-identification bounds are worst-case accounting results for a defined finite population. They can be wide when missingness is substantial, and they do not include sampling uncertainty or a model of missingness. Additional restrictions could narrow them, but this study supplies no evidence for such restrictions. The numerical illustration is synthetic; it does not estimate the missing-clock rate in any real route population.

Finally, the research tolerances and classifications are not production settings. The paper does not evaluate execution, market impact, funding income, liquidation risk or profitability, and it does not recommend a trade or leverage. Public responses and documentation remain subject to their providers' terms. These limitations constrain the claims but also define the study's intended use: an auditable measurement foundation rather than a complete market-data production specification.

9 Conclusion

Multi-venue perpetual-market data cannot be synchronized by assumption. The public transport audit shows that funding, price, order-book and all-market capabilities vary independently, and that price and quote semantics can differ even when a stream exists. The large-sample histories show that timestamp offsets are likewise venue-specific: small positive offsets at Hyperliquid, seconds-level record offsets at Pacifica, a mostly sub-two-second distribution with an extreme anomaly at Extended, and exact settlement-grid alignment at edgeX. These are measurements of returned fields, not estimates of physical clock error or client latency.

The chronological diagnostic qualifies what those historical distributions can support. A tolerance frozen from calibration data remains close to its nominal coverage at Hyperliquid and attains full observed coverage at Pacifica and edgeX, but performs worse at Extended. High coverage obtained with a wide tolerance is not the same as precise synchronization, and a stable historical percentile is not a future guarantee. Extreme records and internal gaps belong in the evidence rather than in an invisible cleaning step.

When timestamps are missing, the appropriate response is also explicit. The bounds $[o/A, (o+m)/A]$ report every full-population exceedance share compatible with the observed counts, with width m/A and no value when $A = 0$. They prevent missing clocks from being converted into zero differences or a more favorable complete-case denominator. The identified set is distinct from a statistical confidence interval and should be presented with the underlying counts.

The resulting framework is deliberately modest. It does not prescribe a universal staleness second count, a production tolerance or a trading rule. It provides a reproducible way to document transport capabilities, measure timestamp behavior, evaluate temporal transfer and expose missing-clock uncertainty. That foundation can support more reliable subsequent research on perpetual funding, cross-venue pricing and execution without overstating what public market data establish.

Declarations

Author contribution and responsibility

Hang Liu designed the measurement framework, reviewed the public evidence, supervised the numerical analysis and approved the final manuscript. The author is responsible for numerical provenance, confidentiality boundaries and the final submission.

Corresponding author and research channel

Hang Liu, humarb Research, London, United Kingdom. Research website: `https://humarb.ai`; research channel: `https://humarb.ai/research/`. Contact: admin@humarb.ai. ORCID: 0009-0008-1059-3926.

Competing interests

The author reports no competing financial or non-financial interests directly related to this study. humarb Research conducts research on digital-asset market mechanisms. This affiliation does not alter the evidence boundaries or the negative empirical findings reported here.

Funding

This research received no external funding; it was conducted by humarb Research.

Data availability

The frozen public-history responses, normalized asset-level and venue-hour data, response hashes, public documentation snapshots and archival transport-capability matrix are included in the reproduction package (Liu, 2026b). Public endpoints and access dates are documented in the paper and package. Aster and ApeX refresh failures are preserved explicitly. Single-host raw WebSocket transport logs are not included; only the de-identified capability matrix and its documentary evidence are released.

Code availability

The reproduction package includes the public-data collector, offline analysis, figure generation, data dictionary and independent verification script. Running the offline workflow regenerates every reported numerical table and vector figure. No private source code, production parameters, internal status names, deployment details or user data are included.

Generative-AI use statement

Generative artificial-intelligence assistance was used for literature-discovery support, language drafting and editing, code assistance, and consistency review. Public data requests

and numerical analyses were executed and checked independently. References were verified against public metadata and official documentation. Hang Liu reviewed the outputs and assumes full responsibility for the manuscript, code and evidence boundaries.

Preprint status and license

This is a methodological and empirical preprint by humarb Research; it has not been peer reviewed. humarb.ai is the research channel of humarb Research. The paper and original figures are offered under CC BY 4.0. Third-party exchange responses and documentation remain subject to their respective providers' terms.

Non-investment-advice statement

This paper is for research and methodological information only. It does not provide investment, legal, accounting or tax advice; does not recommend any venue, trade or leverage; and does not promise data availability, execution, income or protection from loss.

References

References

ApeX (n.d.). Apex omni api documentation: Get funding rate history. https://api-docs.pro.apex.exchange/. Documentation checked 2026-11-03; public endpoint collection failed with connection resets.

Aster (n.d.). Aster api documentation: Get funding rate history. https://docs.asterdex.com/for-developers/aster-api/api-documentation. Documentation checked 2026-11-03; public endpoint collection failed with connection resets.

Binance (n.d.). How to manage a local order book correctly. https://developers.binance.com/docs/derivatives/usds-margined-futures/websocket-market-streams/How-to-manage-a-local-order-book-correctly. Accessed 2026-10-14; legacy documentation page rechecked 2026-11-03.

Budish, E., Cramton, P., and Shim, J. (2015). The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621. DOI: https://doi.org/10.1093/qje/qjv027.

Bybit (n.d.). Websocket public orderbook. https://bybit-exchange.github.io/docs/v5/websocket/public/orderbook. Accessed 2026-11-03.

CCXT (2026). Stop polling: Your trading bot is acting on stale prices. https://docs.ccxt.com/blog/realtime-market-data-with-websockets. Accessed 2026-11-03.

Chainlink (n.d.). Selecting quality data feeds. https://docs.chain.link/data-feeds/selecting-data-feeds. Accessed 2026-11-03.

Cristian, F. (1989). Probabilistic clock synchronization. *Distributed Computing*, 3(3):146–158. DOI: https://doi.org/10.1007/BF01784024.

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

edgeX (n.d.). Funding rate api. https://edgex-1.gitbook.io/edgeX-documentation/api-v2/public-api/funding-api. Accessed 2026-11-03.

Erez, E. and Smirnov, M. (2026). No-arbitrage bounds in decentralized perpetual markets. *Finance Research Letters*, 110:110598. DOI: https://doi.org/10.1016/j.frl.2026.110598.

Extended (n.d.). Extended api documentation: Get funding rates history. https://api.docs.extended.exchange/. Accessed 2026-11-03.

Foucault, T., Hombert, J., and Roşu, I. (2016). News trading and speed. *The Journal of Finance*, 71(1):335–382. DOI: https://doi.org/10.1111/jofi.12302.

Hasbrouck, J. and Saar, G. (2013). Low-latency trading. *Journal of Financial Markets*, 16(4):646–679. DOI: https://doi.org/10.1016/j.finmar.2013.05.003.

humarb Research (2026a). Measuring executable net carry in cross-venue perpetual funding-rate arbitrage. humarb Research 01. DOI: https://doi.org/10.5281/zenodo.23868361.

humarb Research (2026b). Public websocket capability audit: 22 perpetual venues. De-identified archival audit and machine-readable matrix supplied in the reproduction package. Single-host observations dated 2026-10-09; not a multi-location latency study.

Hummingbot Foundation (2026). Exchangepybase and orderbook-tracker source files. https://github.com/hummingbot/hummingbot/tree/9af100d6822da7d2d0291a906c730ef172284ee2. Pinned commit 9af100d6822da7d2d0291a906c730ef172284ee2; accessed 2026-11-03.

Hummingbot Foundation (n.d.). Restart order book tracker. https://condor.hummingbot.org/api-reference/market-data/restart-order-book-tracker. Accessed 2026-11-03.

Hyperliquid (n.d.a). Funding. https://hyperliquid.gitbook.io/hyperliquid-docs/trading/funding. Accessed 2026-11-03.

Hyperliquid (n.d.b). Info endpoint: Perpetuals—retrieve historical funding rates. https://hyperliquid.gitbook.io/hyperliquid-docs/for-developers/api/info-endpoint/perpetuals. Accessed 2026-11-03.

Hyperliquid (n.d.c). Robust price indices. https://hyperliquid.gitbook.io/hyperliquid-docs/trading/robust-price-indices. Accessed 2026-11-03.

Inan, E. (2026). Predictability of funding rates. *International Review of Financial Analysis*, 119:105464. DOI: https://doi.org/10.1016/j.irfa.2026.105464.

Institute of Electrical and Electronics Engineers (2020). Ieee standard for a precision clock synchronization protocol for networked measurement and control systems. DOI: https://doi.org/10.1109/IEEESTD.2020.9120376.

Kopetz, H. and Ochsenreiter, W. (1987). Clock synchronization in distributed real-time systems. *IEEE Transactions on Computers*, C-36(8):933–940. DOI: https://doi.org/10.1109/TC.1987.5009516.

Künsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *The Annals of Statistics*, 17(3):1217–1241. DOI: https://doi.org/10.1214/aos/1176347265.

Lamport, L. (1978). Time, clocks, and the ordering of events in a distributed system. *Communications of the ACM*, 21(7):558–565. DOI: https://doi.org/10.1145/359545.359563.

Liu, H. (2026a). Framing net carry for cross-venue perpetual funding positions: A settled-history, depth-walked, and liquidation-aware method. humarb Research 02. DOI: https://doi.org/10.5281/zenodo.23897013.

Liu, H. (2026b). Reproduction package for freshness and clocks in multi-venue perpetual market data. Public-response snapshots, normalized data, analysis scripts, data dictionary and verification report included with this preprint. Frozen public-history sample collected 2026-11-03 UTC.

LSEG Developer Community (n.d.). Ability for a consumer to identify if prices published by a provider are stale. https://community.developers.refinitiv.com/discussion/comment/7389. Public community response reviewed 2026-10-14; current page supplies limited text to the study environment.

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

Manski, C. F. (2003). *Partial Identification of Probability Distributions*. Springer. DOI: https://doi.org/10.1007/b97478.

Manski, C. F. (2005). Partial identification with missing data: Concepts and findings. *International Journal of Approximate Reasoning*, 39(2–3):151–165. DOI: https://doi.org/10.1016/j.ijar.2004.10.006.

Mills, D., Burbank, J., and Kasch, W. (2010). Network time protocol version 4: Protocol and algorithms specification. Technical Report RFC 5905, Internet Engineering Task Force. DOI: https://doi.org/10.17487/RFC5905.

OKX (n.d.). Okx api v5 documentation. https://www.okx.com/docs-v5/en/. Connection rules accessed 2026-10-14; direct current-page retrieval was unavailable in the study environment.

Pacifica (n.d.a). Funding rates. https://docs.pacifica.fi/trading-on-pacifica/funding-rates. Accessed 2026-11-03.

Pacifica (n.d.b). Get historical funding. https://docs.pacifica.fi/api-documentation/api/rest-api/markets/get-historical-funding. Accessed 2026-11-03.

Pindza, E. (2026). Centralized–decentralized exchange funding rate arbitrage as a basis trade: Risk decomposition, stress testing, and portfolio construction under venue uncertainty. *Digital Finance*, 8(3). DOI: https://doi.org/10.1007/s42521-026-00213-3.

Pyth Network (n.d.). Price feeds: getpricenoolderthan. https://api-reference.pyth.network/price-feeds/evm/getPriceNoOlderThan. Accessed 2026-11-03.

Shah, S. and Pinsky, E. (2026). Perpetual futures in decentralised finance: Mechanics, economic claims, and the drivers of trading volume. *International Journal of Financial Studies*, 14(7):178. DOI: https://doi.org/10.3390/ijfs14070178.

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

Tamer, E. (2010). Partial identification in econometrics. *Annual Review of Economics*, 2(1):167–195. DOI: https://doi.org/10.1146/annurev.economics.050708.143401.

Werapun, W., Karode, T., Suaboot, J., Arpornthip, T., and Sangiamkul, E. (2026). Exploring risk and return profiles of funding rate arbitrage on cex and dex. *Blockchain: Research and Applications*, 7(4):100354. DOI: https://doi.org/10.1016/j.bcra.2025.100354.

Wilson, E. B. (1927). Probable inference, the law of succession, and statistical inference. *Journal of the American Statistical Association*, 22(158):209–212. DOI: https://doi.org/10.1080/01621459.1927.10502953.

Yiğitler, H., Badihi, B., and Jäntti, R. (2020). Overview of time synchronization for iot deployments: Clock discipline algorithms and protocols. *Sensors*, 20(20):5928. DOI: https://doi.org/10.3390/s20205928.

Zhang, L., Mykland, P. A., and Aït-Sahalia, Y. (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association*, 100(472):1394–1411. DOI: https://doi.org/10.1198/016214505000000169.

Zhivkov, P. (2026). The two-tiered structure of cryptocurrency funding rate markets. *Mathematics*, 14(2):346. DOI: https://doi.org/10.3390/math14020346.

Zhivkov, P., Todorov, V., and Georgiev, S. (2026). Temporal dynamics of market microstructure in cryptocurrency perpetual futures: Econometric evidence from centralized and decentralized exchanges. *International Journal of Financial Studies*, 14(5):103. DOI: https://doi.org/10.3390/ijfs14050103.

A Transport-capability appendix

A.1 Full archival matrix

Table 5 reproduces the capability matrix used in the paper. “Partial” in the price column means that the reviewed stream exposes an oracle or index price rather than a field directly identified as mark price. In the all-market column it means that the all-market ticker does not include the funding field. First-frame duration is the selected single-host observation from 2026-10-09 and is not a latency estimate.

Table 5: Public WebSocket capability matrix, 2026-10-09. All URLs are supplied in the reproduction CSV.

Venue	Funding	Price	Book	All-market	First frame (ms)	Evidence note
Binance	Yes	Yes	Yes	Full	1,126	Separate market/public routes; sequence-aware book
OKX	Yes	Yes	Yes	No	1,460	Port 443 used after proxy failure on 8443
Bybit	Yes	Yes	Yes	No	1,690	Ticker snapshot with funding fields
Gate	Yes	Yes	Yes	Full	2,400	`!all` observed but not documented on reviewed page
Bitget	Yes	Yes	Yes	No	1,109	Per-contract subscription; multiple books channels
MEXC	Yes	Yes	Yes	Partial	1,100	All-market ticker has fair price but no funding
KuCoin	Yes	Yes	Yes	Full	2,070	Funding first frame 51,600 ms; minute-level source cadence
BingX	No	Yes	Yes	No	1,490	Funding requires REST; first frame shown for depth
LBank	No	No	No	No	—	Connection opens; no documented usable channel
Deribit	Yes	Yes	Yes	No	570	Multi-channel subscribe; no full-market channel
Hyperliquid	Yes	Yes	Yes	No	931	`activeAssetCtx` funding and price context
Lighter	Yes	Yes	Yes	Full	961	Read-only stream parameter required
dYdX	Yes	Partial	Yes	Full	2,070	`nextFundingRate` and oracle price
Paradex	Yes	Yes	Yes	No	—	SBE schema parameters required for public channels
Aster	Yes	Yes	Yes	Full	833	Binance-like all-market mark-price stream
Extended	Yes	Yes	Yes	Full	1,200	Funding first frame about 36,000 ms; minute-level cadence
ApeX	Yes	Yes	Yes	No	1,140	Instrument ticker with funding and mark price
edgeX	Yes	Partial	Yes	Full	1,900	Funding and index price in all-market ticker
GRVT	Yes	Yes	Yes	No	1,010	Official SDK used when documentation site blocked
Pacifica	Yes	Yes	Yes	Full	1,000	All-market prices channel
GMX	No	No	No	No	—	Oracle-based market; no conventional order book
Variational	No	No	No	No	—	REST/RFQ; quote tiers are not order-book depth

A.2 Capability distinctions that affect replication

Three distinctions matter when reading the matrix. First, an available funding stream can have a slow source cadence even if the transport itself remains connected. KuCoin and Extended therefore require source-cadence awareness before interpreting a lack of new funding values. Second, a full-market subscription can contain many instruments in one frame; this reduces connection count but does not ensure atomic updates across every field. Third, a venue with no conventional order book cannot be assigned a depth-walked execution cost by treating an oracle or RFQ quote as a book level.

The matrix does not disclose internal field mappings or production fallback rules. Public endpoint and channel names appear only where needed to make the capability evidence auditable. Detailed public URLs and archival notes are in `table-10-transport-capability.csv` in the reproduction package.

B Methods appendix

B.1 Symbol table

Table 6: Symbols and definitions.

Symbol	Definition
v	Venue index
a	Asset index: BTC, ETH or SOL
j	Nominal event index
H	One hour in milliseconds, $3{,}600{,}000$
t_{vaj}	Public timestamp in Unix epoch milliseconds
h_{vaj}	Beginning of the UTC hour containing t_{vaj}
d_{vaj}	Asset-level offset $t_{vaj} - h_{vaj}$, in milliseconds
d_{vj}	Primary venue-hour offset: median across available assets
$\mathcal{A}_{vj}$	Assets observed for venue v and event j
$\mathcal{J}_v$	Set of observed venue-hour events for venue v
$Q_v(p)$	Empirical quantile of venue-hour offsets at probability p
$S_v(x)$	Share of venue-hour offsets strictly greater than x ms
$\mathcal{C}_v$	Calibration events before 2026-10-09 UTC
$\mathcal{H}_v$	Holdout events on or after 2026-10-09 UTC
τ_v	Frozen research tolerance $Q_{\mathcal{C}_v}(0.99) + 1$ ms
c_v	Holdout coverage with offsets at most τ_v
E_v	Number of holdout offsets strictly above τ_v
A	Population of pairs eligible for a clock-difference test
R_i	Indicator that both clocks for pair i are present and valid
$Y_i(\tau)$	Indicator that the absolute two-clock difference strictly exceeds τ
n	Valid-clock pair count, $\sum_i R_i$
m	Missing- or invalid-clock pair count, $A - n$
o	Valid-clock pairs exceeding τ
$\theta(\tau)$	Full-population share exceeding τ
$\theta_m(\tau)$	Unobserved exceedance share among missing-clock pairs
$\mathcal{I}(\tau)$	Sharp identified interval $[o/A, (o+m)/A]$
$\mathcal{F}_{v,k,t}$	Conceptual freshness-state vector, not a production score

B.2 Proof of sharp missing-clock bounds

Assume $A > 0$. From equations (12)–(13),

$$A\theta(\tau) = o + \sum_{i:R_i=0} Y_i(\tau).$$

There are m terms in the sum and each lies in $\{0, 1\}$. Therefore

$$0 \le \sum_{i:R_i=0} Y_i(\tau) \le m.$$

Division by A gives

$$o/A \le \theta(\tau) \le (o+m)/A.$$

Assigning $Y_i = 0$ to every missing pair attains the lower bound; assigning $Y_i = 1$ to every missing pair attains the upper bound. Hence the interval is sharp, and subtracting its endpoints yields width m/A.

When $n, m > 0$, the mixture representation is

$$\theta(\tau) = \frac{n}{A}\frac{o}{n} + \frac{m}{A}\theta_m(\tau),$$

where only $\theta_m(\tau)$ is unidentified. Setting $\theta_m(\tau) = o/n$ produces the complete-case rate o/n, but this equality is an assumption. When $m = 0$, both bounds equal $o/A = o/n$. When $n = 0$ and $A = m > 0$, $o = 0$, the complete-case rate is undefined, and the sharp interval is $[0, 1]$. When $A = 0$, the population share itself is undefined.

For two thresholds $\tau_1 < \tau_2$, $Y_i(\tau_2) \le Y_i(\tau_1)$ for every pair because absolute differences above the larger threshold are also above the smaller one. Consequently, $o(\tau_2) \le o(\tau_1)$, and both bounds are nonincreasing in τ for fixed A, n, m. Equality with the threshold never counts as an exceedance.

B.3 Reproduction algorithms

The public-history analysis follows this deterministic sequence:

1. Store each raw public response and compute its SHA-256 hash before normalization.
2. Parse the documented time field, convert it to an integer without floating-point rounding, and filter to the target window and record type.
3. Compute h_{vaj} by integer floor division and $d_{vaj} = t_{vaj} - h_{vaj}$.
4. Deduplicate by venue, asset, timestamp and source time field; retain internal gaps and anomalies.
5. Group by venue and nominal hour, compute the median asset offset, and store asset count and range.
6. Compute full-sample quantiles, tail shares, calendar subsets and alternative weighting summaries.
7. Resample UTC dates in seven-day moving blocks for the headline intervals; repeat with three- and 14-day blocks for sensitivity.
8. Compute each calibration 99th percentile, add 1 ms, freeze the tolerance, and score the later period once.

9. Verify response hashes, primary keys, expected record counts, headline estimates and missing-clock algebraic identities.

The partial-identification illustration uses only constructed counts. Given A, n, m, o, it emits o/A, $(o+m)/A$, m/A, and o/n when $n > 0$. It emits "not applicable" when $A = 0$ and leaves the complete-case value undefined when $n = 0$. Missing differences are never filled with zero.

B.4 Interpretation of numerical precision

All source timestamps and offsets are integers in milliseconds. Quantiles can be fractional because of linear interpolation between ordered observations. Percentages are rounded only for display; coverage is computed from integer exceedance and denominator counts before rounding. The extreme Extended offset is not an artifact of floating-point conversion: the source integer corresponds exactly to 2026-09-24 01:18:16.849 UTC.

C Additional empirical tables

C.1 Calibration and holdout distributions

Table 7: Offset quantiles by calendar period, in milliseconds.

Venue	Period	Hours	Median	p95	p99	Maximum
Hyperliquid	Calibration	912	36.5	90.00	121.89	132
Hyperliquid	Holdout	620	37	83.05	118.81	139
Pacifica	Calibration	909	2,177	5,083	5,083	5,084
Pacifica	Holdout	620	3,097	4,046	4,046	5,083
Extended	Calibration	912	849	1,520.85	1,744.45	1,096,849
Extended	Holdout	620	937	1,643.60	1,852.81	2,094
edgeX	Calibration	228	0	0	0	0
edgeX	Holdout	156	0	0	0	0

C.2 Bootstrap block-length sensitivity

Table 8: Moving-block bootstrap intervals by block length, in milliseconds.

Venue	Block days	Median 95% interval	p95 95% interval
Hyperliquid	3	[35, 38.01]	[83, 94]
Hyperliquid	7	[35, 38]	[83, 94]
Hyperliquid	14	[35, 38]	[83, 93]
Pacifica	3	[2,114, 3,226]	[3,680, 5,083]
Pacifica	7	[2,114, 3,524.01]	[3,679, 5,083]
Pacifica	14	[1,428, 3,679]	[3,525, 5,083]
Extended	3	[849, 937]	[1,494.19, 1,643]
Extended	7	[849, 937]	[1,490.56, 1,651]
Extended	14	[773, 937]	[1,471.90, 1,655]
edgeX	3	[0, 0]	[0, 0]
edgeX	7	[0, 0]	[0, 0]
edgeX	14	[0, 0]	[0, 0]

C.3 Internal gaps and the retained large anomaly

The three internal Pacifica gaps occur at nominal hours 2026-09-05 00:00, 2026-09-09 00:00 and 2026-10-04 00:00 UTC. They are not replaced with neighboring observations. Extended's large anomaly is 2026-09-24 01:18:16.849 UTC for BTC, ETH and SOL. It produces one venue-hour offset of 1,096,849 ms and remains in calibration and full-sample results. These events answer different questions: missing rows reduce observable coverage, whereas the large timestamp remains observed but exceeds an ordinary tolerance.

C.4 Pilot-only observations excluded from new estimates

Table 9: Pilot observations from 2026-10-26. These rows are contextual and are not pooled with the frozen large sample.

Venue	Rows	Time field	Observed offsets (ms)	Current use
Hyperliquid	24	`time`	1 to 118	Superseded by 1,532 venue-hours
Pacifica	3	`created_at`	2,114 to 2,115	Superseded by 1,529 venue-hours
Extended	3	`T`	849 to 1,692	Superseded by 1,532 venue-hours
Aster	3	`fundingTime`	0 to 1	No refresh; no distribution claim
ApeX	3	`fundingTime`	0	No refresh; no distribution claim
edgeX	3	`fundingTime`	0	Superseded by 384 settlement events
Paradex	3	snapshot creation	$-14{,}657$ to $-4{,}657$	Semantic snapshot counterexample only

C.5 Public staleness practices

Table 10: Public freshness mechanisms reviewed in this study. These mechanisms are not interchangeable production thresholds.

Source	Public mechanism	Reported cadence or condition	Evidence boundary
Binance	Local-book sequence continuity	Previous final update must connect to the next event	No universal staleness second count supplied
Bybit	Snapshot/delta and repeated unchanged snapshot	L1 10 ms; L50 20 ms; L200 100 ms; L1000 200 ms; unchanged L1 snapshot after 3 s	Cadence is not an error threshold
OKX	Connection heartbeat and disconnect rule	Connection-level 30 s rule in archival review	Current page not directly retrieved in study environment
Hyperliquid	Validator price-update description	Oracle and mark prices approximately every 3 s	Does not define client stale-age policy
Hummingbot	User-stream inactivity and REST polling interval	60 s inactivity comparison selects 5 s rather than 120 s polling	Not an order-book expiry threshold
Pyth	Caller-specified maximum price age	`age` argument; stale-price error under applicable validity condition	Caller must choose asset-appropriate age
Chainlink	Feed-quality controls and circuit breakers	Detect stale, delayed, missing or unexpected inputs	No universal second count prescribed
LSEG community	Provider quality status	Suspect data can be marked while stream remains open	Archival public community evidence
CCXT / Condor	Sequence checks, reconnection and tracker restart	Snapshot/delta maintenance and explicit restart operation	Operational recovery, not a freshness guarantee

The public review supports a mechanism-based interpretation. A sequence gap can require snapshot recovery even when messages arrive frequently. A heartbeat can remain healthy while a source value is intentionally unchanged. A source-age argument appropriate for an oracle need not apply to a level-one order book. A provider quality flag can communicate suspect data without deleting it. These distinctions explain why the paper does not recommend a universal number of seconds.

D Data dictionary and reproduction guide

D.1 Normalized asset-level fields

Table 11: Fields in `funding_timestamp_offsets.csv`.

Field	Type, unit and interpretation
`venue`	Public venue name
`asset`	BTC, ETH or SOL
`timestamp_ms`	Integer Unix epoch milliseconds returned by the public source

Field	Type, unit and interpretation
timestamp_utc	ISO 8601 UTC rendering of the timestamp
nominal_hour_ms	Beginning of the containing UTC hour in epoch milliseconds
nominal_hour_utc	ISO 8601 rendering of the nominal hour
offset_ms	Integer timestamp minus nominal hour, in milliseconds
utc_date	UTC calendar date used for date-level resampling
period	Calibration before 2026-10-09; holdout thereafter
record_type	Semantic type of the public record
source_time_field	Public response field used for the timestamp
funding_rate	Public funding rate where returned; not used in offset estimation
funding_interval_min	Public interval where available; 240 for the observed edgeX contracts

The venue-hour file adds assets_observed, records, min_asset_offset_ms and max_asset_offset_ms. Its offset_ms is the median across available assets. Raw manifest entries store the public request, retrieval timestamp, status, elapsed time, response path, SHA-256 hash and any error. The manifest's elapsed time is provenance and is not used as a latency estimate.

D.2 Package map

The reproduction package contains the collector collect_public_history.py, offline analysis analyze.py, verifier verify_reproduction.py, raw and normalized data, documentation snapshots, data dictionary, endpoint notes, output tables and SVG/PDF figures. requirements.txt records Python package versions. outputs/verification.json records the final acceptance result.

Offline reproduction uses:

```
python3analyze.py
python3verify_reproduction.py
```

The expected verifier status is passed. A new online collection is optional and produces a new evidence snapshot; it is not required to audit the frozen numerical results.

D.3 Data-use boundaries

The package contains no credentials, user data, private endpoints, internal production fields, hidden thresholds or deployment details. Aster and ApeX failures remain explicit. Pilot observations are segregated from new estimates. The transport matrix is archival and single-host. Partial-identification counts are synthetic. All research tolerances are examples for public-data measurement and do not represent production settings.