Exchange Operational Announcements and Public Transfer Status: A 180-Day Labeled Dataset and a Four-Venue Polling Study

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

humarb Research 07 · v1 · Methodological and empirical preprint · DOI pending Zenodo publication

Abstract

Operational announcements are often used as if they were synchronized records of market access. This paper separates announcement evidence from observed transfer status in two public, redacted datasets. Dataset A contains 996 labels across a 20-venue target map for 2026-04-23 to 2026-10-20, with explicit coverage gaps and separate publication and effective-time fields. Dataset B audits 502 polling rounds at four venues over a 169.1-hour log span. Among 929 auditable non-Gate key-transition records, 503 had no contemporaneous semantic announcement match; this record-level proportion is concentrated at LBank and is not an independent-event rate. For 81 matched announcement events, median detection-minus-publication lag is +5.81 minutes, with venue-specific signs and sensitivity bounds. The median of 517 continuously observed closed withdrawal intervals is 0.648 hours, but batch exclusions move it materially. The release supplies explicit estimands, gap-aware duration partitions, reproducible tables and figures, and evidence boundaries for research on operational information without claiming execution, profitability, or market-impact effects.

Keywords: exchange announcements; delisting; maintenance; deposit and withdrawal status; dataset documentation; observational design

JEL classification: G12; G13; G14; C58

1. Introduction

A trader, risk officer, or dataset consumer who sees an exchange announcement about a delisting, network upgrade, or transfer suspension faces three different questions. When was the announcement published? When does its text say that an operational change should occur? When did an independent public interface first show a corresponding state? These timestamps can differ by minutes, hours, or days, and they need not refer to the same object. Treating them as interchangeable turns an operational information problem into a silent measurement error. The problem is especially important for perpetual-funding research: a quoted spread may remain visible while the transfers needed to rebalance collateral are unavailable, a contract approaches delisting, or a venue changes the instruments supported on a network. Yet an announcement alone does not demonstrate that a switch flipped, and an interface change alone does not demonstrate that no announcement existed.

This paper asks a deliberately narrower question than whether operational events predict returns: **what can a carefully documented announcement label set and a short public-status polling study establish about the coverage, timing, and observability of exchange operational information?** Existing cryptocurrency market microstructure research establishes why cross-venue prices, contract design, and arbitrage frictions matter (Makarov and Schoar 2020; De Blasis and Webb 2022). Perpetual-futures research explains the funding mechanism and its pricing implications (He et al. 2022; Ackerer, Hugonnier, and Jermann 2026). Neither literature supplies a public record-level dataset that distinguishes announcement publication, text-stated effective time, and polled

transfer-state detection while preserving known observation gaps. That is the gap addressed here.

The study has two components. Dataset A is an announcement label set spanning 180 elapsed days and covering a target map of 20 centralized and decentralized trading venues from 2026-04-23 through 2026-10-20, inclusive, in coordinated universal time (UTC). It contains 996 rows from ten venues with nonzero events. Pacifica has an observable channel but zero in-window events; nine other venues have full-window or historical gaps. The unit is an announcement row after semantic event deduplication, not a token, chain, or state transition. The label schema stores publication time, event kind, market tags, instruments, chains, effective time when actually stated, a rescheduling link, source type, and a boolean showing whether the private annotation had a note. Public evidence excerpts and note text are not redistributed. Public URLs remain available for source checking.

Dataset B is a separate four-venue polling study of Bitget, KuCoin, Gate, and LBank. Its log spans 169.1 hours from 2026-10-21 13:02:10 UTC to 2026-10-28 14:05:59 UTC and contains 502 rounds. It is not seven days of continuous coverage: a known gap of approximately 13.2 hours, earlier lost temporary data, and long rounds limit observability. The analysis asks three questions. Q1 concerns a detected key-level state transition for which no semantically corresponding contemporaneous announcement was found in the saved, auditable snapshot range. Q2 concerns an announcement for which the observable key surface provides strict evidence that the announced state did not occur. Q3 measures the difference between a detection timestamp and announcement publication for matched events. A fourth descriptive object is the observed time between withdrawal-close and withdrawal-open detections when the interval does not cross a long observation gap.

Three design choices prevent the most consequential denominator errors. First, Q1 is measured in key-by-transition records, Q2 in announcement events, Q3 in unique matched announcement events, and duration in observed close–open intervals. The 81 Q3 events map to 420 key transitions; the 420 are not substituted for 81. Second, Gate is not assigned a zero or nonzero Q1/Q2 rate, because its announcement text coverage is unavailable. Third, continuous closure, gap-crossing endpoint pairs, right censoring, and recovery/opening records without an observed stop are separate partitions. These choices follow the documentation principle that data consumers need provenance and intended-use boundaries, not only rows (Gebru et al. 2021).

The announcement labels show substantial event flow but incomplete temporal and venue coverage. Maintenance accounts for 438 of 996 rows, delistings for 329, transfer suspensions for 188, resumptions for 30, and rescheduling notices for 11; there are no rows labeled retract. Binance, Bybit, and KuCoin contribute 816 rows, or 81.93% of the sample. Effective time is present at second resolution for 720 rows, at date-only resolution for 23, and null for 253. Null is not replaced with publication time. Among the 720 second-resolution pairs, median effective-minus-publication lead is 46.98 hours and the 90th percentile is 167.49 hours. Twenty-two values are negative and six are zero. These are text-stated scheduling distances, not measured execution delays.

The polling results show why record-level humility is necessary. On the three-venue auditable non-Gate surface, 503 of 929 key-transition records are Q1, a descriptive 54.14%. LBank supplies 398 of the 503 Q1 records, so the pooled figure is not a venue-invariant operational fact. Strictly confirmed Q2 is zero only on the currently decidable snapshots; five announcements lack an observable key, two are future notices, one is unresolved because a five-minute stop and resumption may fit between polls, and one key was already in the announced state. Gate remains undecidable. For Q3, median detection-minus-publication lag across 81 events is +5.81 minutes and the 90th percentile is +17.00 minutes. The sign is positive for 54 events, negative for 23, and uncertain after interval bounding for four. Bitget’s median is +8.50 minutes, KuCoin’s is 266.62 minutes on only four events, and LBank’s is 20.69 minutes on 20 events. A negative lag means detection before publication, not proof that the operational flip preceded an announcement that was publicly

invisible.

Observed durations are even more batch-sensitive. The main partition contains 517 continuously observed closed withdrawal intervals, 21 gap-crossing endpoint pairs, 72 right-censored records, and 98 openings or recoveries without an observed stop. The continuous-closure median is 0.648 hours, or 38.90 minutes. It remains unchanged when seven AMPL intervals are added or a 38-record GRAM batch is removed, but rises to 0.915 hours when the 295-record LBank ERC20 batch is removed. Removing all same-start batches of eight or more leaves 75 intervals and a 5.028-hour median. Consequently, the headline median is a property of a key-level record mix, not “the typical exchange suspension.”

The contributions are fourfold. First, the paper releases a structured announcement dataset with explicit source and coverage maps, separate timestamp granularities, and rescheduling links, while refusing to impute absent execution times. Second, it supplies an observation-first pairing framework for announcements and public transfer status, including counting units, semantic-matching requirements, and timing bounds. Third, it reports a gap-aware partition of status durations and demonstrates how chain-level batches and field oscillations change descriptive medians. Fourth, it converts these limitations into reproducible outputs: a public-schema validator, machine-readable statistics, table files, eight vector figures, and data-acquisition instructions that distinguish released derived data from unavailable raw logs.

Identification is descriptive rather than causal. There is no untreated control venue, no randomized announcement, and no counterfactual market outcome. The validity of Q1 depends on saved announcement coverage and semantic matching; Q2 depends on key coverage, due execution time, and polling resolution; Q3 depends on correct pairing and the fact that a poll detects rather than directly observes a flip. When these assumptions fail, the appropriate category is unknown, gap, or unresolved—not zero. No profit, return prediction, execution success, or price-impact claim is made. Sample-out prediction tests, forecast-loss comparisons, trading-cost economic significance, and multiple-model corrections are therefore inapplicable, rather than omitted. The related methodological issues are discussed explicitly in Section 7.

This study complements humarb Research 01 and 02. Liu (2026a) measures executable net carry after depth-dependent costs, and Liu (2026b) frames net carry using settled funding history, asynchronous settlement grids, and liquidation-aware accounting. Those papers address the measurement and presentation of net holding returns. The present paper adds the operational-event, announcement-time, and transfer-status observation layer. It does not validate their return estimates, and it does not turn a status record into a realized tradability measure.

Section 2 positions the work in the perpetual-futures, arbitrage-friction, dataset-documentation, and evaluation literatures. Section 3 summarizes the public funding mechanisms that motivate operational awareness without estimating funding returns. Section 4 documents both datasets. Section 5 defines estimands and observation rules. Section 6 reports results. Section 7 evaluates robustness and explains which finance-journal tests are not applicable. Section 8 states limitations and ethical boundaries. Section 9 concludes. Declarations, references, and appendices provide provenance, a symbol table, a data dictionary, and additional audit detail.

2. Related literature

2.1 Perpetual futures, funding, and operational frictions

Perpetual futures replace expiry with recurring transfers between long and short positions. He et al. (2022) derive frictionless no-arbitrage prices and bounds with trading costs, emphasizing that perpetuals do not guarantee convergence to spot. Ackerer, Hugonnier, and Jermann (2026) develop perpetual-futures pricing in a continuous-time framework. This literature makes the funding mechanism economically meaningful, but an implementable position also requires the contract to remain listed and the necessary collateral movements to remain operationally feasible. Our data do not enter a pricing kernel or estimate a funding premium. They record the operational announcements and public status observations that a pricing or execution dataset may otherwise silently treat as continuous availability.

The distinction matters because funding cash flows and operational access have different clocks. A funding rate can be computed from a premium index and interest component while a network upgrade is scheduled for a later time. A contract delisting can terminate future trading without changing the historical funding series. A transfer switch can affect collateral mobility without appearing in a candle or funding record. The paper therefore treats operational announcements as a complementary information layer, not as another price feature to be appended without temporal controls.

2.2 Cross-venue arbitrage, contract design, and market frictions

Makarov and Schoar (2020) document arbitrage and segmentation in cryptocurrency markets, while Shleifer and Vishny (1997) provide the broader limits-to-arbitrage framework. De Blasis and Webb (2022) connect arbitrage, contract design, and market structure in Bitcoin futures. Together, these studies motivate attention to capital mobility and implementation constraints. However, a status flag is not a direct measure of arbitrage capital, and a suspension announcement is not a measured transaction cost. The present evidence can help researchers avoid treating every venue and network as continuously fungible, but it cannot quantify the shadow cost of being unable to transfer a particular asset at a particular unobserved instant.

This distinction also separates our contribution from Liu (2026a, 2026b). Those studies measure or frame net carry using funding history, order-book walks, and liquidation-aware calculations. Our announcement data could be joined to such frameworks only through explicit venue, instrument, chain, and timestamp mappings. A naive join would create look-ahead bias when a later retrospective announcement is attached to an earlier state, or false availability when a future notice is treated as immediately effective. We publish the fields needed to discuss those joins but do not perform a return backtest.

2.3 Dataset documentation and annotation quality

Datasheets for Datasets argues that dataset creators should document motivation, composition, collection, preprocessing, uses, distribution, and maintenance (Gebru et al. 2021). We adopt that logic in a financial-operational setting. The release describes not only the 996 observed labels but also the nine venues with gaps and the distinction between an observable zero and an unobservable surface. It also distinguishes source type from access channel and preserves a public URL without republishing announcement text.

Annotation quality requires a separate evidentiary standard. Artstein and Poesio (2008) explain why inter-coder agreement requires multiple annotators and an appropriate chance-corrected statistic. This release has one annotator and an 88-item single-person review of an earlier 981-row version,

including two out-of-window rescheduling supplements. It therefore does not report Cohen's kappa, Krippendorff's alpha, or an "88-item error rate." Subsequent full-scan corrections cannot be reverse-engineered into the error rate of that earlier review. The release instead states what was checked, what changed, and why independent agreement remains unavailable.

2.4 Timing, censoring, and the limits of descriptive evidence

Polling produces discrete observations of a partially latent process. This differs from event-time data in which exact starts and ends are observed. Survival methods such as Kaplan and Meier (1958) and doubly censored generalizations such as Turnbull (1974) are relevant only after the censoring scheme and independence assumptions are justified. Here, long gaps, batch dependence, unknown starts, and field oscillations make an off-the-shelf survival estimate potentially misleading. We therefore report observed closed intervals and separate gap and censoring partitions rather than presenting a survival curve as identified.

Forecast evaluation has its own rigorous standards. Diebold and Mariano (1995) compare predictive accuracy, and Harvey, Liu, and Zhu (2016) address multiple testing in cross-sectional return research. Those tools are not decorative requirements for a dataset paper. There is no forecast target, no model search, and no return significance test in this study. We cite this literature to delimit what would be required for a future predictive or trading paper, not to imply that the current descriptive medians have passed such tests.

2.5 Positioning

The resulting contribution is closest to an empirically documented data infrastructure study. It does not compete with perpetual pricing models or arbitrage estimates. It supplies a measurement layer with enough provenance to prevent three common overstatements: that a publication timestamp is an execution timestamp, that no retrieved announcement means no announcement exists, and that adjacent polling observations prove continuous status. The empirical results are interesting precisely because they show where these shortcuts fail and how strongly record composition can affect a simple summary.

3. Institutional background

3.1 Why funding mechanics do not subsume operational status

Perpetual contracts use funding payments to anchor contract prices to an underlying index or oracle. Public venue documentation shares a conceptual core—an impact-price premium, an interest component, and a settlement rule—but differs in sampling, averaging, caps, and settlement grids. The formulas below are institutional background, not production parameters used by humarb and not inputs to any return estimate in this paper.

Let I_t be an index or oracle price, B_t an impact bid price, and A_t an impact ask price. A common premium form is

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

A stylized funding rate with interest component r, premium average bar P, and damper \delta is

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

The actual definitions of impact prices, sampling weights, averaging windows, interest rates, and caps are venue- and sometimes contract-specific. The operational layer enters outside equation (2):

an otherwise well-specified funding observation does not establish that a contract remains listed or that collateral can move on a selected network.

Binance's public documentation describes an impact-notional-based premium, a default eight-hour funding interval for many contracts, an interest-rate component, a ±0.05% damper in its standard formula, and contract-specific caps and interval adjustments (Binance, n.d., accessed 2026-11-03). Bybit documents an impact bid/ask premium, minute-level sampling, a time-weighted average, a similar clamp structure, and funding caps and intervals that can vary by contract (Bybit, n.d., accessed 2026-11-03). These public rules explain why funding histories require contract-specific metadata, but they say nothing by themselves about a particular deposit or withdrawal key at a particular polling instant.

Hyperliquid documents hourly funding payments based on an eight-hour-style rate divided across hourly intervals, five-second premium sampling, an oracle-price premium, and a stated hourly cap (Hyperliquid, n.d., accessed 2026-11-03). dYdX documents minute-level premium samples, an hourly funding tick, impact-notional prices, market-specific interest settings, and governance-controlled parameters (dYdX, n.d., accessed 2026-11-03). Thus even settlement frequency is not safely interchangeable across venues. Our announcement labels add delisting, maintenance, and transfer events to this institutional picture without pretending that the four documents exhaust the 20-venue target map.

3.2 Announcement publication, stated effectiveness, and actual visibility

An announcement has at least two meaningful times. Publication time is the timestamp attached to the public announcement record under the source-specific rule documented in the release. Effective time is the time stated in the text for trading termination, a maintenance-related transfer stop, or another scoped action. A later update can supersede an earlier notice. A retrospective notice can be published after the stated event. A completion-tense notice can omit a future effective time altogether. Assigning publication time to every null effective-time field would convert these different cases into fabricated precision.

Actual public visibility is a third concept. A source timestamp is evidence about a field carried by a public record; it is not always proof of the first moment at which every user could retrieve the record. Dataset B contains two visibility anomalies in which an announcement appeared in a later saved surface despite an earlier nominal time field. Those records are excluded from the main Q3 sample. We do not replace their nominal timestamp with an inferred visibility timestamp, because the required historical public-visibility evidence is unavailable.

3.3 Public transfer-status surfaces

The polling study uses public, no-key status responses. At Bitget, KuCoin, and LBank, a key represents a coin-by-network surface and deposit and withdrawal states can move independently. Gate additionally exposes coin-level and chain-level fields whose Boolean direction differs from the simple enabled convention. A missing key is not a disabled key. A reappearing key after an unobserved interval is not automatically a close–open episode. A coin-level field and a chain-level field are not interchangeable.

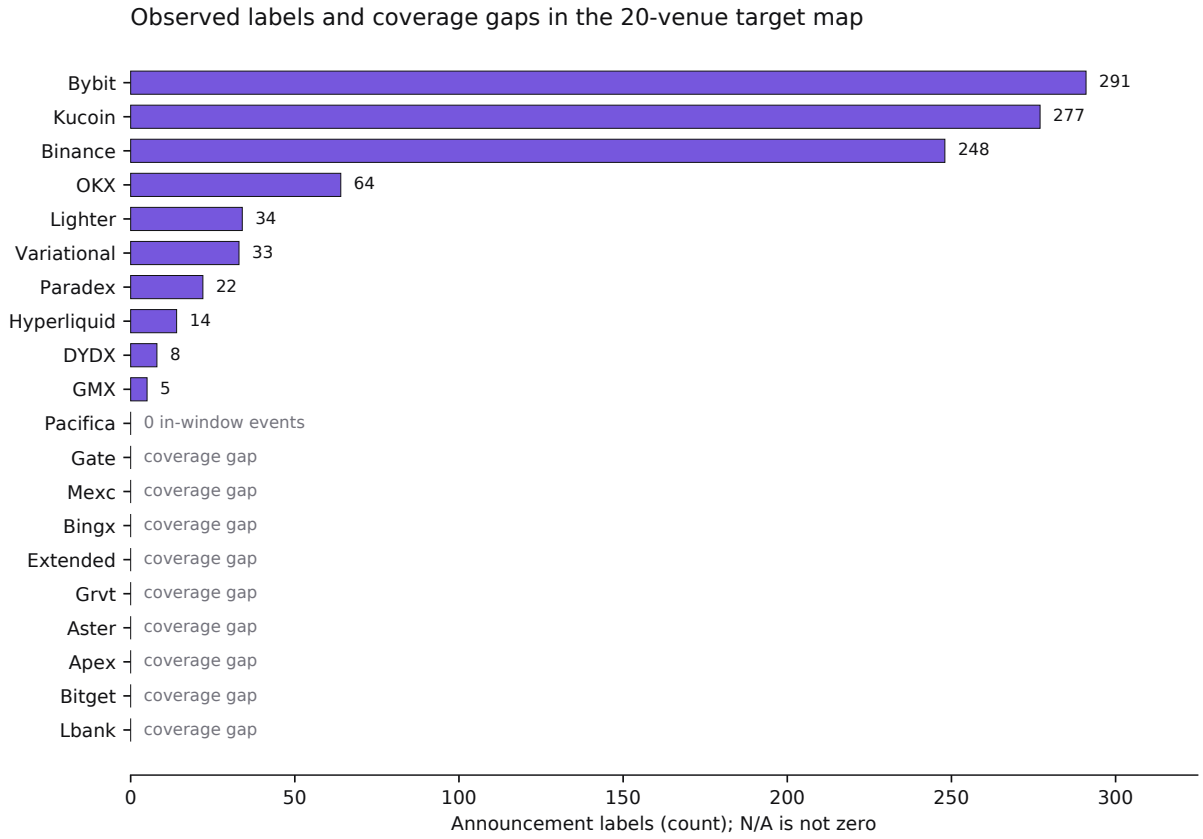
The observational unit is therefore a public interface key, not necessarily an asset in an economic sense. One chain-level announcement can correspond to many key transitions. Conversely, one economic asset can have several independent network keys. These many-to-many mappings explain why announcement counts, transition counts, and duration counts must be analyzed separately. They also explain why the paper reports exact key composition in the replication outputs rather than treating all rows as independent exchange decisions.

4. Data

4.1 Dataset A: target map, scope, and observation window

Dataset A covers the inclusive UTC window 2026-04-23 through 2026-10-20 and was collected and annotated on 2026-10-21 and 2026-10-22. The endpoints are 180 days apart; including both UTC dates yields 181 calendar dates. The target map contains 20 venues. Eleven have observable announcement sources under the study's access rules: Binance, Bybit, dYdX, GMX, Hyperliquid, KuCoin, Lighter, OKX, Pacifica, Paradex, and Variational. Pacifica contributes a genuine zero under its observed changelog and status-feed scope. The other ten observable venues contribute 996 labels. Gate, MEXC, BingX, Extended, GRVT, Aster, ApeX, Bitget, and LBank have full-window or historical gaps.

The distinction between zero and missing is fundamental. Pacifica's zero means that the scoped public channels were observable and no qualifying in-window event was found. A gap means that a complete historical enumeration could not be obtained under the study's access and compliance constraints. Gate's direct announcement access returned a restriction and no circumvention was attempted; MEXC and GRVT had access restrictions; BingX lacked a publicly enumerable history behind a device-signature wall; Extended's relevant channel was not publicly enumerable; Aster's historical enumeration was unavailable; ApeX had only a short current RSS history; and Bitget and LBank had current-page snapshots but incomplete earlier histories. These are observations about this collection surface, not universal claims that no public information exists anywhere.



Source: public redacted labels, 2026-04-23 to 2026-10-20 UTC. Hatched rows have historical or full-window gaps.

Figure 1. Announcement labels and coverage gaps across the 20-venue target map. Hatched zero-height bars denote gaps rather than zero events. Source: Dataset A, 2026-04-23 to 2026-10-20 UTC.

The inclusion rule covers delistings, maintenance and network upgrades, transfer suspensions and resumptions, and rescheduling notices within the documented market scope. Product or service closures are retained only as explicitly tagged boundary cases. Fee promotions, Earn-parameter

changes, purely promotional items, general product launches, software-development-kit updates without a scoped market event, compliance notices, account-only changes, and pure news are excluded. The same event reported on multiple channels is represented once, using the earliest suitable publication record and the most precise defensible time fields. Rescheduling links point from a later announcement to the earlier announcement it supersedes.

Table 1. Dataset composition and timestamp coverage. The announcement lead and Q3 detection lag are distinct estimands. Source: authors' calculations from the released public data.

Item	Value	Scope
Announcement labels	996	2026-04-23 to 2026-10-20 UTC
Target venues	20	20-venue target map
Observable venues	11	11; Pacifica has zero in-window events
Venues with nonzero labels	10	10
Venues with historical/full-window gaps	9	9
Second-resolution effectiveAt	720	Announcement labels
Date-only effectiveAt	23	Announcement labels
Null effectiveAt	253	Not imputed
Instrument assignments	3,762	683 non-empty rows
Chain assignments	294	274 non-empty rows
Supersedes links	8	Later announcement to earlier announcement
Top-three venue share	81.93%	Binance, Bybit, KuCoin
Median publication-to-effective lead (hours)	46.98	n=720 second-resolution pairs; not Q3 lag

4.2 Label schema and temporal granularity

Each public row contains venue, a globally unique identifier, a public URL, second-resolution UTC publication time, event kind, a multi-select market list, instrument and chain lists, effective time, an optional supersession identifier, source type, and a boolean note marker. The public schema intentionally excludes evidence excerpts and free-text confidence notes. The boolean marker records whether such a note existed in the private annotation; it does not permit its content to be reconstructed.

The six event kinds are mutually exclusive at the row level. Market tags are not: a row can reference spot and perpetual markets, or perpetual, contracts-for-difference, and margin scopes. Instrument strings preserve the named trading pair, contract, token, or asset surface when applicable; an empty list means no single instrument list was defensible, not that the event had no operational relevance. Chain names are included only when explicitly supported by the announcement. The release restores semantically equivalent single-letter percent encoding in public identifiers, URLs, and one chain name without changing their meaning.

Effective time has three states. A full ISO UTC timestamp is used only when the text supports second-resolution representation. A date-only value retains date granularity rather than appending midnight. Null means that no defensible effective timestamp was stated or that a completion-tense announcement already described the action as completed. Three genuine midnight values are retained as second-resolution timestamps because the source text actually states midnight; they are not date-only values promoted artificially.

Publication-to-effective lead is defined only for the 720 second-resolution effective-time rows. This selection is intentional and transparent. Date-only values are not included by choosing an arbitrary hour, and null values are not replaced by zero. Supersession links are not interpreted as fresh independent operational events without checking the linked record. Eight rows carry such links.

4.3 Venue and source composition

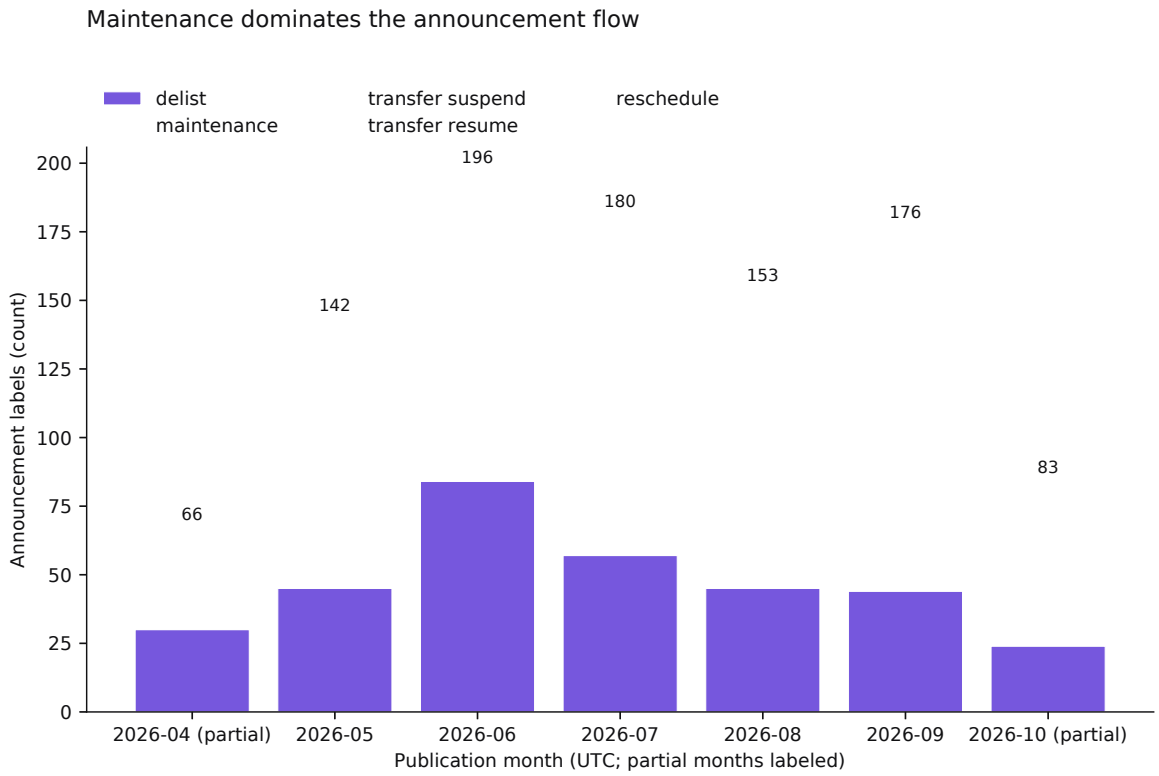
Bybit contributes 291 rows, KuCoin 277, Binance 248, OKX 64, Lighter 34, Variational 33, Paradex 22, Hyperliquid 14, dYdX eight, and GMX five. The three largest sources therefore account for 81.93% of all labels. Counts reflect both event activity and source architecture; they should not be interpreted as venue risk rankings. Venues with separate listing and technical channels, long maintenance histories, or granular transfer notices can produce more rows without necessarily experiencing more independent economic disruptions.

Official content-management systems supply 902 rows, public Telegram mirror records 76, status pages 11, and governance forums seven. Source type describes the evidence surface, not a claim about reliability. Governance proposals and status incidents may require cross-checks against execution surfaces. Public mirror messages can be edited or deleted without the mirror preserving a full revision history. OKX labels represent the en-ae public announcement surface; other regional surfaces are outside scope.

Table 2. Venue-level label counts and coverage status. N/A denotes unavailable historical coverage, not zero. Source: Dataset A coverage map and public labels.

Venue	Labels	Coverage
bybit	291	observable
kucoin	277	observable
binance	248	observable
okx	64	observable
lighter	34	observable
variational	33	observable
paradex	22	observable
hyperliquid	14	observable
dydx	8	observable
gmx	5	observable
pacifica	0	observable; zero in-window events
gate	N/A	historical/full-window gap
mexc	N/A	historical/full-window gap
bingx	N/A	historical/full-window gap
extended	N/A	historical/full-window gap
grvt	N/A	historical/full-window gap
aster	N/A	historical/full-window gap
apex	N/A	historical/full-window gap
bitget	N/A	historical/full-window gap
lbank	N/A	historical/full-window gap

Monthly counts range from 66 in partial April to 196 in June, with 83 in partial October. Maintenance is the largest category overall and in the aggregate monthly profile. Figure 2 preserves category composition and marks the partial months. Because venues enter the sample through different source surfaces and gaps are not random, the monthly series should not be used as a complete market-wide operational-event index.



Source: 996 public labels. Categories are mutually exclusive; market tags are not used in this decomposition.

Figure 2. Monthly announcement composition by event kind. April and October are partial months. Source: 996 labels in Dataset A.

4.4 Completeness checks and annotation review

Completeness evidence varies by venue. Binance, Bybit, KuCoin, OKX, and Paradex have enumerable public channels with pagination or cursor evidence extending beyond the window. OKX was cross-checked against an official announcement interface: 461 in-window records reduced to 68 under the three-category comparison, of which 46 matched this dataset one-to-one and 22 were outside scope; another 18 rows in this dataset came from other sections. The delisting and transfer categories matched completely within the stated en-ae scope. This is strong scope-specific reconciliation, not proof that every regional or deleted announcement is observed.

KuCoin's category reconciliation illustrates why category trees matter. Some in-scope delistings appear in latest or futures categories rather than only in the delistings category. Lighter's public technical channels and status page were deduplicated at the event level. dYdX governance proposals were checked against public governance execution status, while Hyperliquid delisting observations were cross-checked against structured market metadata. Paradex and Variational lack an equivalent second independent source in the release. These differences are documented rather than collapsed into a single completeness score.

The 88-item review was performed by one annotator on an earlier 981-row version and included two out-of-window rescheduling supplements. It cannot measure agreement across annotators. Later review rounds corrected fields across the dataset, but the number of corrected categories is not an error rate for the original 88-item review. The replication package validates the public schema, recomputes descriptive counts, and preserves this limitation. The earlier private 90-day baseline and the evidence excerpts needed to reproduce every value-level baseline difference are not included.

4.5 Dataset B: polling span and released evidence

Dataset B observes Bitget, KuCoin, Gate, and LBank through public transfer-status surfaces. The log spans 169.1 hours and 502 rounds. The first-to-last status-change span is 168.8 hours and is not substituted for the log span. A known approximately 13.2-hour gap occurred between a round beginning on 2026-10-24 at 03:48:54 UTC and a restart at 17:09:39 UTC. Earlier temporary data were lost. One additional round lasted 54.18 minutes. Common rounds were approximately 18–20 minutes; the ten-minute post-round sleep in the collector is not the full polling period.

All changes in a round share the round-start record in the historical data. Actual venue responses occur later and sequentially. The released bounds use a processing-upper-bound proxy and, where relevant, the next round. They are not exact response-time measurements. The historical collector did not store venue-specific response timestamps, separate key-reappearance events, or per-coin LBank success counts. Later software capabilities are not backfilled into the historical sample.

The release includes the audited key-transition classification, matched event-level timing bounds, and duration partitions. It does not include raw polling logs, raw state files, full announcement snapshots, or private process notes. Consequently, the replication package can independently recompute every reported statistic from the released derived JSON, but cannot reconstruct the raw log collection or prove the original announcement-search path from scratch. A new polling run would be a new sample, not a bitwise reproduction of this historical window.

4.6 Ethical and legal scope

Only structured labels and public URLs are released; announcement text, forum posts, and Telegram message bodies are not redistributed. Automated access was limited to public no-key surfaces, and restrictions such as access denials or device-signature requirements were not bypassed. The package does not provide instructions to evade them. Any new collection requires a fresh review of applicable terms, robots policies, redistribution rights, and low-frequency access practices. These boundaries are part of the data-generating process and therefore part of the dataset’s coverage limitation.

5. Methods

5.1 Two datasets and four counting units

Let announcement row r_i have publication time p_i, effective-time representation e_i, kind k_i, venue v_i, and instrument and chain sets. Dataset A is the ordered collection of such rows after event-level deduplication. Its descriptive unit is one announcement row. Dataset B contains several related but nonidentical units: a key-transition record, a matched announcement event, and an observed withdrawal close–open interval. A “key” is a public coin-by-network or coin-level status object under the venue-specific mapping.

This separation is not cosmetic. If one announcement maps to 326 ERC20 keys, counting keys estimates interface breadth while counting announcements estimates communication events. Neither count is inherently wrong, but substituting one denominator for the other changes the estimand. Similarly, two transitions on the same key can form one observed closed interval, while an opening without a prior observed stop cannot. All tables state their unit and all robustness transformations preserve it.

5.2 Announcement publication-to-effective lead

For rows with second-resolution e_i, define

$$L_i = e_i - p_i \tag{3}$$

where L_i is measured in hours. Positive values indicate a future stated effective time relative to publication, zero indicates coincidence at the recorded precision, and negative values indicate a retrospective stated time. The sample median and 90th percentile are conventional order-statistic summaries with linear interpolation. They describe the 720 eligible rows, not all 996 announcements. Date-only and null values remain outside the distribution.

The lead is not a forecast of execution. It does not establish that the event occurred at e_i, that the announcement was visible to every user at p_i, or that a market participant could trade or transfer throughout [p_i,e_i]. Its economic interpretation is limited to announced scheduling distance. The negative tail is retained because deleting retrospective notices would conceal an important announcement type.

5.3 Q1: transitions without a contemporaneous semantic match

Let z_j be an audited key-transition record and let M_j=1 when a semantically corresponding announcement is found in the saved contemporaneous snapshot range. Q1 records have M_j=0 after excluding key disappearances, fixed Gate oscillation records, records matched during the known gap, and visibility anomalies. On the three-venue non-Gate auditable surface, the descriptive proportion is

$$q_1 = \frac{\sum_j \mathbf{1}(M_j = 0)}{\sum_j \mathbf{1}(j \in \mathcal{A})} \tag{4}$$

where A is the 929-record auditable set. The numerator is 503. A semantic match requires the correct venue, asset or instrument, network where applicable, operational direction, and plausible time relation. A historical announcement about the same asset does not automatically explain a new transition. A listing notice does not prove a withdrawal resumption. A similarly named token or chain is not a valid match.

The identifying assumption for interpreting Q1 as “no matching announcement in the audited surface” is that the saved announcement snapshots and manual review cover the relevant contemporaneous surface. The assumption does not support “no announcement existed.” Incomplete historical pagination, deleted or edited posts, regional differences, later publication, and human matching error can all produce false Q1 classifications. Conversely, an excessively broad match can falsely remove a record from Q1. We therefore report the exact audit count and venue composition rather than a standard error that would pretend independent sampling.

Gate is excluded from equation (4). Its raw record count includes fixed field oscillations and key disappearances, and 545 remaining records have an announcement-coverage gap. Assigning them to Q1 would equate unavailable text with absent text; assigning zero would assert matches without evidence. The appropriate value is undecidable.

5.4 Q2: announced but strictly unobserved execution

A strict Q2 announcement requires four conditions: an observable corresponding key, an announced action already due, polling fine enough to detect the change, and evidence that the state did not change as announced. If any condition fails, the record is not strict Q2. We separately count no observable key, future notice, polling-unresolved events, and keys already in the announced state.

This rule prevents two opposing errors. It does not call every unpaired announcement an execution failure, and it does not interpret an unchanged poll as proof of compliance. A five-minute suspension

and resumption can fall entirely between rounds. A future maintenance notice should not show a change at publication. A key already disabled cannot supply a fresh disabling transition. Strict Q2 is therefore an observation-based classification, not a direct audit of venue conduct.

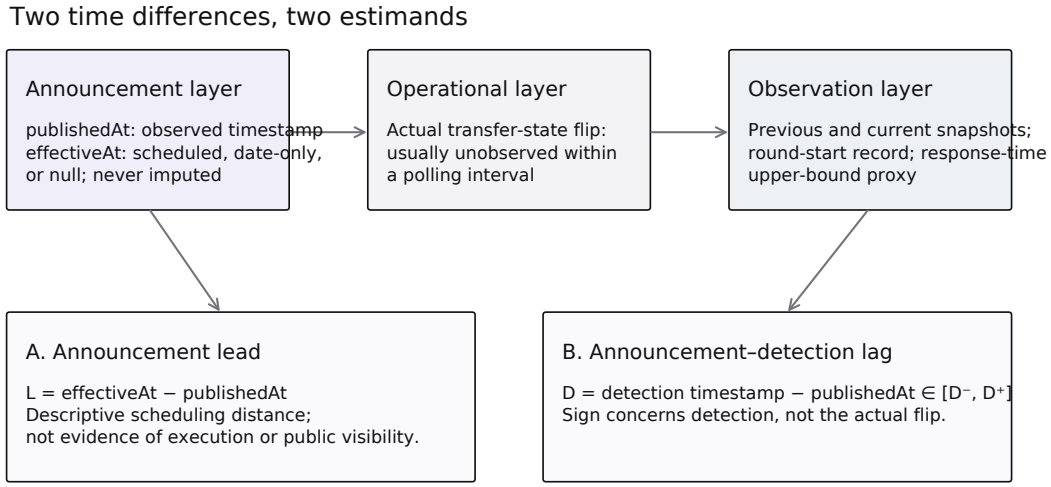
5.5 Q3: announcement–detection lag and bounds

For matched announcement event j, define the recorded lag in minutes as

$$D_j = t_j^{\text{det}} - p_j \tag{5}$$

Positive D_j means that publication precedes the recorded detection; negative D_j means that detection precedes publication. The main sample uses publication time, not an execution time stated in the body. It excludes three gap samples, 11 body-time comparisons, and two visibility anomalies. Each of the 81 rows is a unique announcement event; the 420 mapped key transitions are used only for Q1 accounting.

Because all detections in a round share a historical round-start timestamp, the true response detection time lies in an interval. Let D_j^- and D_j^+ be the released lower and upper lag bounds, incorporating the processing-upper-bound proxy and announcement minute precision where applicable. The sign is announcement-before-detection when D_j^->0, detection-before-announcement when D_j^+<0, and uncertain otherwise. We recompute medians and 90th percentiles once with all lower bounds and once with all upper bounds. This is a deterministic sensitivity envelope, not a joint confidence interval and not a probability distribution over response times.



The middle box is the unobserved operational event. Q3 intervals are sensitivity bounds, not confidence intervals.

Figure 4. Observation model for announcement lead and announcement–detection lag. The actual operational flip is usually latent between polls.

Even exact response timestamps would not identify the actual transfer flip. If one poll shows enabled and the next shows disabled, the flip occurred sometime between the relevant observations, but its exact time is unknown. A negative D_j can reflect retrospective communication; a positive D_j can reflect either an earlier actual flip or polling delay. Q3 therefore measures public-record alignment, not announcement effectiveness or operational causality.

5.6 Observed withdrawal close–open intervals

For a withdrawal-close detection c_k followed by a withdrawal-open detection o_k on the same key, define

$$U_k = o_k - c_k \tag{6}$$

Only pairs that do not cross a known gap longer than one hour enter the main continuous-closure distribution. Deposit-only changes are excluded. The other partitions are: gap-crossing endpoint pairs, for which elapsed time is descriptive but continuity is unproven; right-censored stops without a subsequent observed opening; and openings without an observed stop, which can include new listing enablement as well as recovery from an unobserved historical stop.

The phrase “continuous closure” is itself bounded: it means that the pair does not cross the known long gap, not that every intermediate poll proves uninterrupted disabled status. Short events can occur between ordinary rounds. A missing key truncates an observed segment rather than being interpreted as recovery. Recurrent fixed field oscillations are excluded from the main Gate duration sample, while AMPL and GRAM alternatives are reported separately. Same-start batches are retained in the main key-level distribution and removed in sensitivity analyses.

5.7 Descriptive inference and reproducibility

The paper does not assign classical standard errors to pooled medians or proportions. The released rows are a finite observed sample with venue-level clustering, same-start batches, deterministic interface coverage, and nonrandom gaps. A naive key-level bootstrap would manufacture independence. A venue-cluster interval with only three decidable venues would be unstable and would conceal, rather than solve, the measurement problem. We instead show exact counts, time splits, venue leave-one-out results, bound recomputations, and batch exclusions.

Every table and figure is generated by one public script from the released labels and derived JSON. A separate validator checks the redacted schema, unique identifiers, timestamp formats, category counts, Q1/Q3 units, and duration partition. Validation establishes internal reproducibility; it does not recreate unavailable raw logs, private evidence excerpts, or independent annotation.

6. Results

6.1 Announcement activity and timestamp completeness

The 996 labels are unevenly distributed by kind and venue. Maintenance contributes 44.0% of rows and delistings 33.0%; transfer suspensions contribute 18.9%, resumptions 3.0%, and rescheduling notices 1.1%. The absence of retract labels means no event was assigned that category under the annotation rules, not that no source ever edited an announcement. Market multi-select counts include 514 spot tags, 335 perpetual tags, 82 margin tags, 56 other tags, 29 CFD tags, and one delivery tag; they sum to more than 996 by design.

There are 683 rows with nonempty instrument lists, containing 3,762 assignments and 2,246 distinct instrument strings. Chain lists are nonempty for 274 rows and contain 294 assignments across 118 distinct strings. These counts measure annotation breadth, not unique economic assets: naming conventions can differ across venues and rows. Eight rows carry supersession links. The note marker is true for 937 rows and false for 59, indicating extensive granularity or interpretation annotations in the private source without releasing their text.

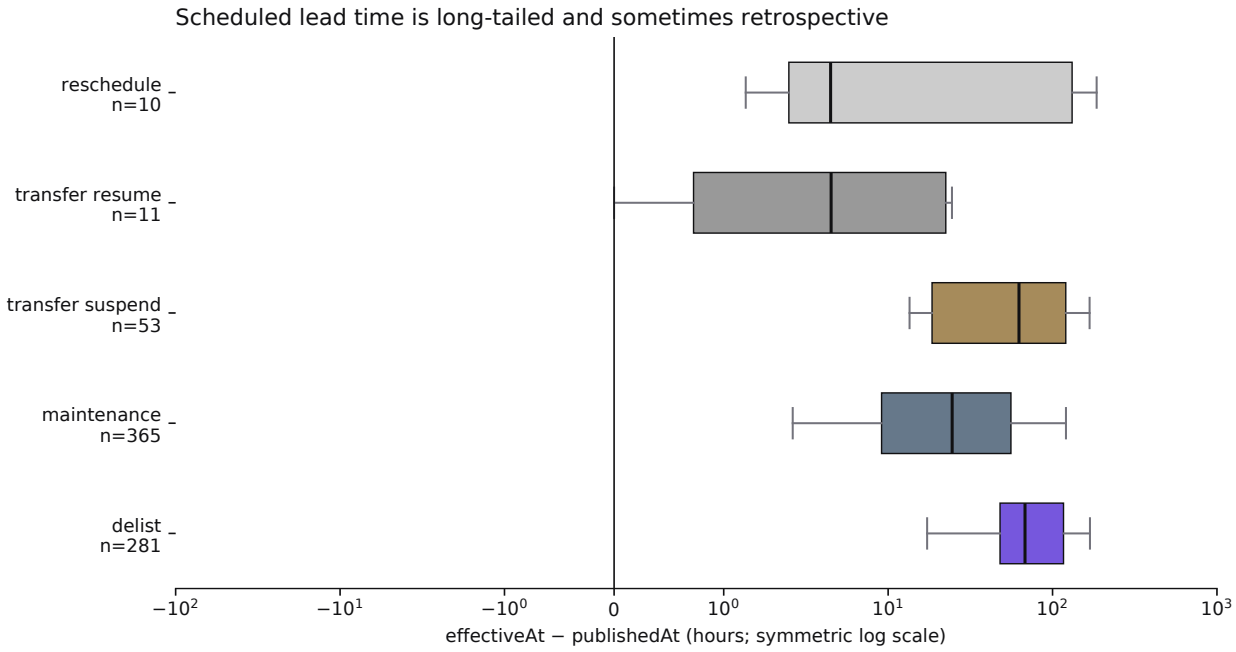
Effective-time completeness is the central descriptive result. Second-resolution values cover 72.3% of rows, date-only values 2.3%, and null values 25.4%. This composition varies by kind. Delistings have

281 second-resolution pairs among 329 rows; maintenance has 365 among 438; transfer suspensions only 53 among 188; and resumptions only 11 among 30. A pooled “average warning time” across all events would therefore combine different missingness mechanisms and should not be reported.

Table 3. Publication-to-effective lead by event kind, restricted to second-resolution effective times. Negative values are retrospective stated times. Source: Dataset A.

Kind	Labels	Second n	Other time	Median h	p90 h	Neg. / zero
delist	329	281	48	68.00	168.9	11 / 1
maintenance	438	365	73	24.51	122.7	7 / 2
transfer suspend	188	53	135	62.50	169.6	3 / 2
transfer resume	30	11	19	4.500	24.45	1 / 1
reschedule	11	10	1	4.469	246.2	0 / 0

Among the 720 eligible pairs, median lead is 46.98 hours, the interquartile range is 15.73 to 74.25 hours, and the 90th percentile is 167.49 hours. Delistings have the largest kind-level median at 68.00 hours, followed by transfer suspensions at 62.50 hours and maintenance at 24.51 hours. Resumptions and rescheduling notices have medians of 4.50 and 4.47 hours on much smaller samples. Twenty-two eligible rows have negative leads, six are zero, and 692 are positive. The longest retrospective value is approximately 1,184 hours; this tail is why a median and full distribution are more informative than a statement that announcements are “advance notice.”



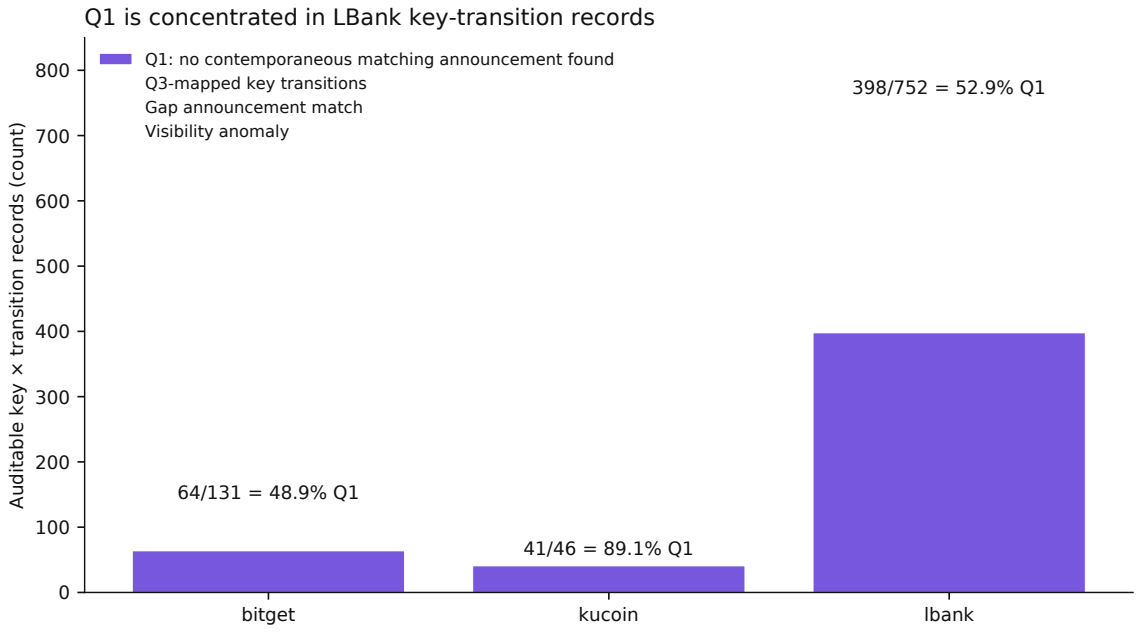
Boxes: quartiles; center: median; whiskers: p10–p90; outliers omitted for readability. Negative values are retained. n=720 second-resolution pairs.

Figure 3. Publication-to-effective lead by event kind on a symmetric-log scale. Boxes show quartiles and whiskers p10–p90; negative values are retained. Source: 720 second-resolution pairs in Dataset A.

The results support two practical conclusions. First, announced scheduling horizons are often long enough to matter for dataset timestamping and risk calendars. Second, the same field cannot be treated as universally prospective or precise. Researchers must retain null and date-only states and inspect retrospective notices before joining announcements to market data.

6.2 Q1: most auditable unmatched records come from one venue

The three-venue auditable set contains 929 key-transition records. Of these, 503 are Q1, 420 are mapped to Q3 publication events, three match announcements around the observation gap, and three are visibility anomalies. The pooled Q1 share is 54.14%. This is not an estimate that 54.14% of independent real operational events lack announcements: one chain event can generate many keys, repeated changes can occur on one key, and announcement coverage differs by venue.



Source: 929 non-Gate auditable transitions. Gate is excluded because announcement coverage is unobservable; records are not independent events.

Figure 5. Composition of 929 auditable non-Gate key-transition records. Percentages use each venue's auditable transitions as the denominator. Source: Dataset B announcement audit.

Bitget contributes 131 auditable records: 64 Q1, 61 Q3-mapped, three gap-announcement matches, and three visibility anomalies. Its Q1 share is 48.85%. KuCoin contributes 46 records: 41 Q1 and five Q3-mapped, giving 89.13%. LBank contributes 752 records: 398 Q1 and 354 Q3-mapped, giving 52.93%. LBank alone supplies 79.13% of all Q1 records and 80.95% of the auditable denominator. The pooled result is therefore principally an LBank-weighted interface description.

The interpretation of Q1 is asymmetric across directions and venues. Saved LBank snapshots contain repeated recovery-side absences and bundled or retrospective suspension notices. Bitget includes repeated short changes without contemporaneous matches as well as later-discovered retrospective announcements, the latter reclassified out of Q1. KuCoin includes changes for which a historical announcement exists but does not explain the current transition. These examples motivate semantic and contemporaneous matching rather than keyword matching over an asset's entire history.

Gate's 2,683 raw records do not enter the pooled ratio. Fixed oscillations account for 2,134, key disappearances for four, and announcement-coverage gaps for 545. The last group is unknown with respect to Q1, not evidence for or against announcement presence. Any pooled figure that silently includes Gate as either numerator or denominator would be invalid.

6.3 Q2: zero strict cases has a narrow meaning

No announcement on the currently decidable three-venue snapshot surface meets all strict Q2 conditions. This zero does not imply perfect execution. Five announcements have no corresponding key on the observed surface, two are future notices at the audit time, one five-minute Bitget KAT stop–resumption sequence may fall between polls, and one KuCoin ROSE key was already in the announced disabled state. Gate is undecidable because announcement coverage is missing.

The KAT example is particularly useful. Announcements describe a five-minute stop and resumption that leaves no polling trace, while a later stop is detected. Calling the first announcement a strict Q2 would confuse sampling resolution with nonexecution. Conversely, observing the later stop does not prove that the earlier five-minute sequence occurred exactly as written. The dataset preserves both announcement and observation limits.

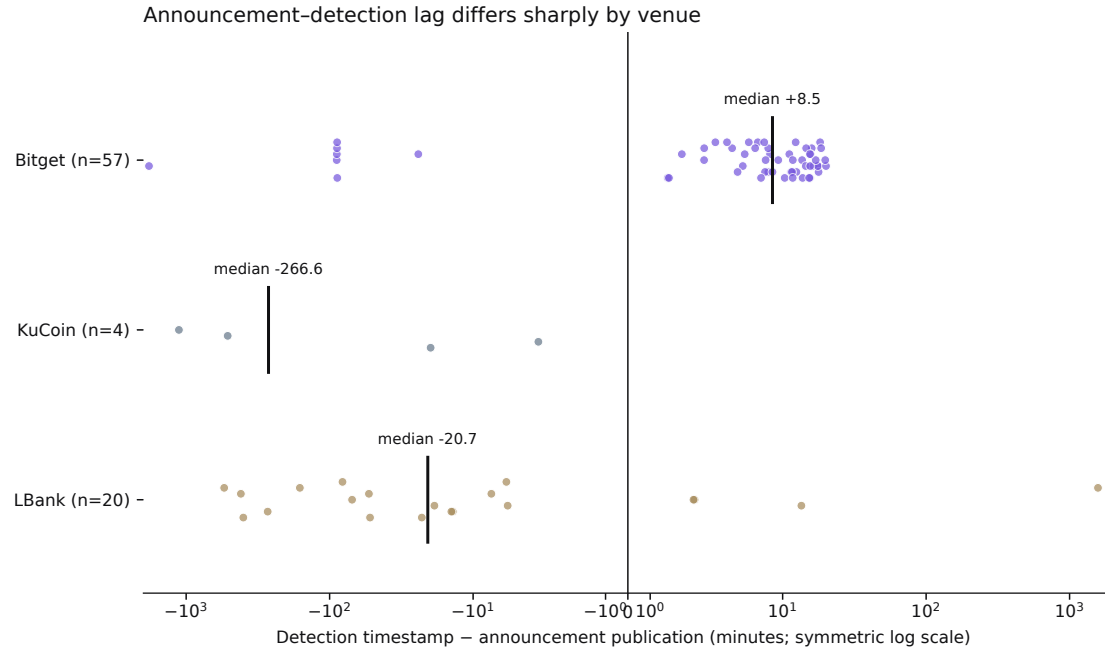
6.4 Q3: pooled median masks opposite venue signs

Across 81 matched publication events, median recorded lag is +5.81 minutes and the 90th percentile is +17.00 minutes. Recomputing with lower and upper detection bounds gives a median sensitivity envelope of +5.43 to +14.89 minutes and a p90 envelope of +17.00 to +25.42 minutes. All 81 events have finite bounds. Fifty-four have publication definitely before detection, 23 have detection definitely before publication, and four have uncertain sign.

Table 5. Announcement–detection lag by venue. Sensitivity envelopes recompute order statistics with lower and upper bounds; they are not confidence intervals. Source: 81 unique Q3 announcement events.

Venue	n	Median min	Median envelope	p90 min	p90 envelope	A / D / U
bitget	57	8.496	8.496 to 17.95	17.55	17.55 to 26.13	50 / 7 / 0
kucoin	4	-266.6	-267.6 to -257.6	-8.738	-9.738 to 0.572	0 / 3 / 1
lbank	20	-20.69	-21.69 to -11.05	4.043	3.897 to 13.86	4 / 13 / 3
overall	81	5.808	5.433 to 14.89	17.00	17.00 to 25.42	54 / 23 / 4

Bitget’s 57 events have median +8.50 minutes and p90 +17.55 minutes. Its interval directions are 50 publication-before-detection and seven detection-before-publication. KuCoin has only four events: median 266.62 minutes and p90 8.74 minutes, driven by retrospective notices including WARD and INIT. Three are definitely detection-before-publication and one is sign-uncertain. LBank’s 20 events have median 20.69 minutes and p90 +4.04 minutes; four are publication-before-detection, 13 detection-before-publication, and three sign-uncertain.



One point per announcement event (n=81), not per mapped key. Negative values mean detection preceded publication; neither sign orders the actual transfer flip.

Figure 6. Recorded Q3 lag by venue on a symmetric-log scale. One point per announcement event; vertical bars mark venue medians. Source: Dataset B timing bounds.

Two records illustrate why event counting and timestamp definitions matter. One LBank ERC20 chain-maintenance announcement maps to 326 keys but contributes one Q3 event; publication is about +1,577 minutes before the recorded chain-wide detection, while the separately stated execution time is roughly 30–39 minutes before detection under the body-time comparison. Bitget AL is a retrospective notice with recorded lag around 1,815 minutes. Mixing the chain announcement's 326 keys into the event distribution or replacing publication with body time would fundamentally alter the estimand.

The results do not rank venue communication quality. KuCoin's four-event sample is too small for a stable distributional claim, Bitget and LBank have different announcement and key mappings, and only detected matched events enter Q3. A venue with fewer observed pairs may have different coverage rather than different operational behavior. Q3 is best read as a set of documented examples showing why publication and detection clocks cannot be collapsed.

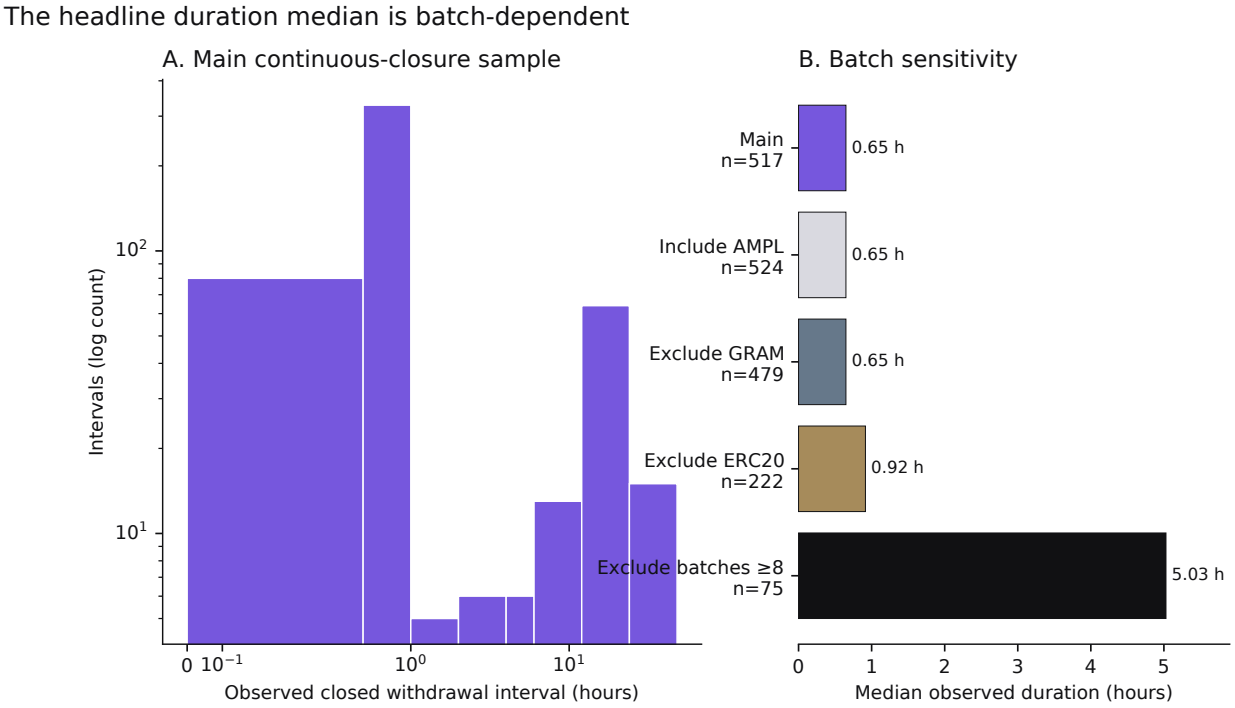
6.5 Duration: a precise median with fragile interpretation

The main continuous-closure sample has 517 intervals. Its median is 0.648 hours, minimum 0.300 hours, maximum 46.025 hours, and 90th percentile 14.413 hours. The buckets contain 408 intervals below one hour, 17 from one to six hours, 77 from six to 24 hours, and 15 at least 24 hours. Separate partitions contain 21 gap-crossing endpoint pairs, 72 right-censored stops, and 98 openings without an observed stop. The last group includes 38 GRAM openings in a token migration-like batch and possible new-listing enablement; it is not a set of 98 proven recovered suspensions.

Table 6. Continuous-closure duration under alternative batch and venue rules. Other duration partitions are not silently reclassified. Source: Dataset B duration data.

Scenario	n	Median h	Range h	Buckets: <1 / 1–6 / 6–24 / 24 h
main	517	0.648	0.300–46.02	408 / 17 / 77 / 15
including ampl	524	0.648	0.300–46.02	415 / 17 / 77 / 15
excluding gram batch	479	0.648	0.300–46.02	408 / 17 / 39 / 15
excluding lbank erc20 batch	222	0.915	0.300–46.02	113 / 17 / 77 / 15
excluding all start batches of 8 or more	75	5.028	0.301–46.02	21 / 17 / 30 / 7
without lbank	204	0.661	0.300–46.02	110 / 15 / 73 / 6
without gate	357	0.648	0.301–46.02	303 / 16 / 24 / 14

Venue composition explains much of the distribution. LBank contributes 313 of the 517 intervals, Gate 160, Bitget 37, and KuCoin seven. LBank's 295 same-start ERC20 closures all share approximately the same observed duration and strongly influence the pooled median. Gate contains short xStocks-related batches and the separate GRAM sequence. Bitget's median is 10.00 hours on 37 intervals and KuCoin's 5.03 hours on seven, but neither supports a general venue ranking because their event mix and censoring differ.



Continuous closure excludes known >1 h gaps but does not prove uninterrupted status between polls.
Excluding the 295-record LBank ERC20 batch raises the median to 0.92 h; excluding all start batches of eight or more raises it to 5.03 h.

Figure 7. Distribution and batch sensitivity of observed continuous-closure withdrawal durations. Exclusions change the key-level record mix, not the underlying raw records. Source: Dataset B.

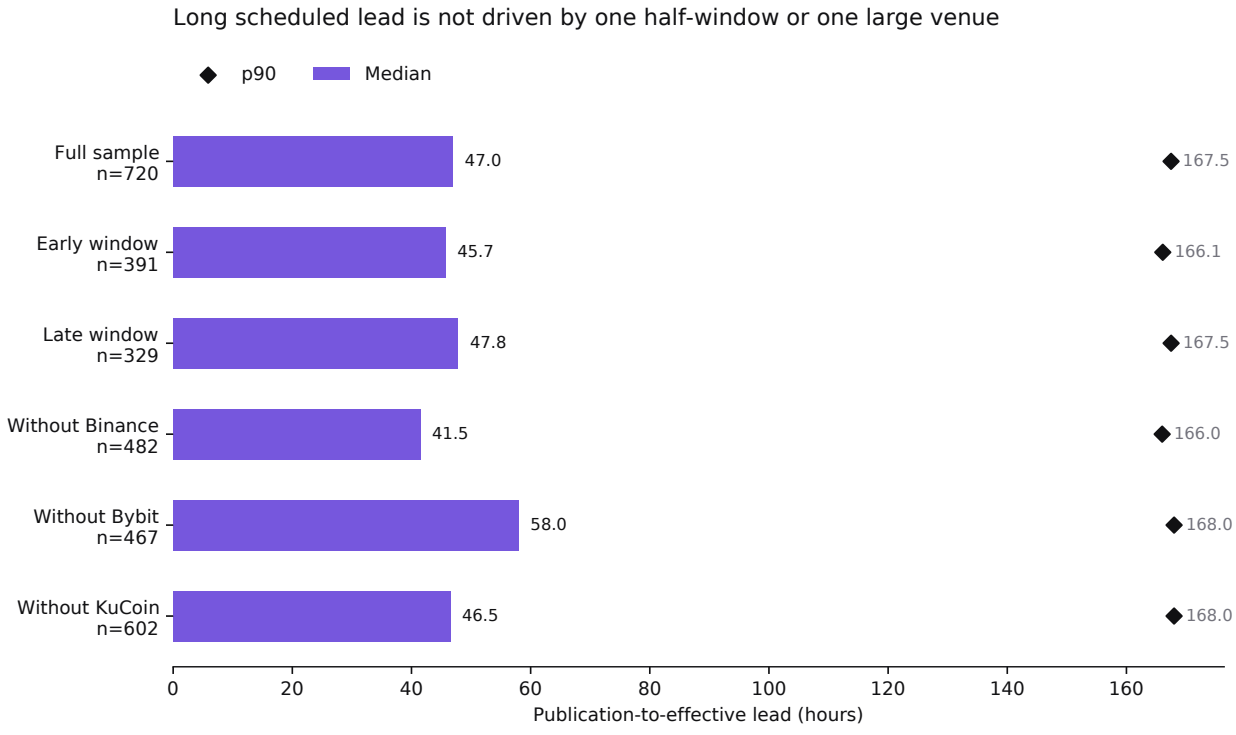
The median remains 0.648 hours when seven AMPL intervals are included or when the 38-record GRAM batch is removed. Excluding the LBank ERC20 batch leaves 222 intervals and raises the median to 0.915 hours. Excluding every same-start batch of eight or more leaves 75 intervals and raises it to 5.028 hours. The latter is a deliberately harsh composition diagnostic, not a proposed new population estimate. The correct conclusion is not that one of these medians is universally true, but that the key-level pooled median is batch-dependent and must always be reported with composition.

7. Robustness and evaluation boundaries

7.1 Announcement lead across time and venue composition

The announcement-lead result is first evaluated across fixed calendar subsamples rather than a data-driven median split. The early window from 2026-04-23 through 2026-07-21 contains 391 eligible second-resolution pairs and has median lead 45.70 hours with p90 166.10 hours. The late window from 2026-07-22 through 2026-10-20 contains 329 pairs and has median 47.83 hours with p90 167.50 hours. The similar central values support the narrow conclusion that multi-day stated lead is common in both parts of this particular sample. They do not establish time stability outside the 180-day window, and the late window contains the longest retrospective observation.

Leaving out one large contributing venue changes the median but not the order of magnitude. Excluding Binance leaves 482 pairs and lowers the median to 41.48 hours; excluding Bybit leaves 467 and raises it to 58.00 hours; excluding KuCoin leaves 602 and produces 46.50 hours. The corresponding p90 values remain between 166.00 and 168.00 hours. These shifts are substantively useful: they show that the pooled median is not a universal venue parameter, while also showing that the distinction between hours and days is not created by one venue alone. A descriptive restriction to nonnegative leads leaves 698 observations and median 47.94 hours. That alternative is reported only to show the limited effect of removing retrospective values; the main estimand retains them because retrospective communication is itself relevant.



Only second-resolution effectiveAt values enter; null and date-only values are not imputed. Diamonds show p90, not uncertainty intervals.

Figure 8. Publication-to-effective lead under fixed time subsamples and venue leave-one-out rules. Points show medians and whiskers p10–p90. Source: Dataset A, second-resolution pairs only.

7.2 Q1 across time and venue composition

Q1 is evaluated with a fixed UTC cutoff of 2026-10-25 00:00:00, avoiding a record-median cutoff that would itself be affected by large batches. Before the cutoff, 107 of 150 auditable records are Q1, or 71.33%. From the cutoff onward, 396 of 779 are Q1, or 50.83%. The decline warns against extrapolating the pooled share across even this short polling span. It may reflect changing event composition, announcement practices, key mappings, or saved snapshot coverage; the released evidence cannot distinguish these causes.

Venue leave-one-out results lead to the same measurement conclusion. Without Bitget, the share is 439/798 = 55.01%; without KuCoin, 462/883 = 52.32%; without LBank, 105/177 = 59.32%. Thus the qualitative presence of unmatched records is not exclusively an LBank phenomenon, even though LBank dominates the pooled numerator and denominator. At the same time, the leave-LBank denominator is only 177 records and retains Bitget's gap and visibility categories. The appropriate reading is a range of observed record mixes, not a confidence interval over a common venue effect. No variance-weighted average is estimated because the necessary independent sampling model is absent.

7.3 Q3 under alternative samples and timing bounds

Q3 has three distinct robustness dimensions: event composition, observation-time bounds, and timestamp definition. For composition, the first 40 ordered events have median +6.12 minutes and the remaining 41 have median +4.92 minutes. Excluding Bitget reverses the median sign to 21.27 minutes on 24 events. Excluding KuCoin leaves 77 events and median +6.71 minutes, while excluding LBank leaves 61 and median +7.95 minutes. These results demonstrate why the pooled +5.81 minutes cannot be described as a venue-invariant communications delay.

Two targeted exclusions probe influential retrospective and chain-level observations. Removing the two events whose absolute recorded lag exceeds 24 hours leaves 79 events and keeps the median at +5.81 minutes, while p90 changes to +16.53 minutes. Removing the single LBank ERC20 chain event leaves 80 events, median +5.62 minutes, and p90 +16.47 minutes. The median is therefore not driven solely by the most extreme records, but this stability does not rescue a universal interpretation: the venue leave-one-out sign change remains.

For observation-time uncertainty, all lower-bound substitutions produce median +5.43 minutes and p90 +17.00 minutes; all upper-bound substitutions produce median +14.89 minutes and p90 +25.42 minutes. The upper shift is material relative to the recorded median. It is a deterministic sensitivity envelope, not a 95% interval, because no probability model assigns mass within the detection intervals. Finally, the main analysis does not substitute body-stated execution time for publication time. The separate body-time comparisons illustrate that doing so would answer a different question and would mix timestamp semantics. This is an alternative-estimand boundary, not another estimate of the same lag.

Table 7. Robustness of announcement lead, Q1 share, and Q3 lag. Rows retain their original units and counting units; Q1 medians are proportions. Source: Datasets A and B.

Family	Scenario	n	Median	p90	Unit
lead time	full sample	720	46.98	167.5	hours
lead time	early 2026-04-23 to 2026-07-21	391	45.70	166.1	hours
lead time	late 2026-07-22 to 2026-10-20	329	47.83	167.5	hours
lead time	without binance	482	41.48	166.0	hours
lead time	without bybit	467	58.00	168.0	hours
lead time	without kucoin	602	46.50	168.0	hours
q3 lag	without bitget	24	-21.27	2.975	minutes
q3 lag	without kucoin	77	6.715	17.21	minutes
q3 lag	without lbank	61	7.954	17.52	minutes
q3 lag	first half	40	6.120	15.51	minutes
q3 lag	second half	41	4.916	18.24	minutes
q3 lag	exclude abs lag over 24h	79	5.808	16.53	minutes
q3 lag	exclude lbank erc20 chain event	80	5.624	16.47	minutes
q1 share	without bitget	798	55.01%		proportion
q1 share	without kucoin	883	52.32%		proportion
q1 share	without lbank	177	59.32%		proportion

7.4 Duration under batch, field, and venue rules

Duration robustness addresses record construction rather than tweaking a numerical estimator. Including seven AMPL intervals produces n = 524 and leaves the median at 0.648 hours. Removing the 38-record GRAM batch produces n = 479 and the same median, but p90 falls from 14.41 to 10.06 hours and the six-to-24-hour bucket falls from 77 to 39. Excluding the 295-record LBank ERC20 batch produces n = 222, median 0.915 hours, and p90 19.86 hours. Excluding all same-start batches of at least eight records produces n = 75, median 5.028 hours, and p90 22.26 hours. These alternatives do not identify a preferred sample; they expose how a key-level distribution weights operational batches.

Venue exclusions provide a complementary check. Removing LBank leaves 204 intervals and median 0.661 hours, whereas removing Gate leaves 357 and median 0.648 hours. Similar pooled medians can therefore coexist with very different venue distributions and upper tails. The fixed Gate field-flap union contains 266 intervals in the released sensitivity metadata, while the main construction excludes the documented recurrent oscillation pattern. AMPL, GRAM, chain-batch, and venue alternatives should be read together with the four duration partitions. Moving right-censored or no-observed-start records into the closed sample would be a change in evidence, not a robustness test.

7.5 Why forecast and trading-performance tests are not applicable

The study has no predicted outcome, forecast horizon, fitted return model, or held-out trading rule. Rolling or expanding sample-out evaluation would therefore require inventing a new research question. Mean absolute error, root mean squared error, forecast coverage tests, proper scoring rules, and the Diebold–Mariano test compare forecasts against realized targets; no such target–forecast pairs exist here. The calendar subsamples above are descriptive composition checks, not evidence of predictive generalization.

Likewise, there is no model or hyperparameter search whose selection could be controlled with a multiple-testing adjustment. The robustness scenarios are prespecified measurement diagnostics,

and all are reported rather than selected for significance. Harvey–Liu–Zhu adjustments, reality checks, and backtest-overfitting probabilities do not convert descriptive record counts into a validated return strategy. Classical standard errors are not added for the same reason: deterministic coverage, venue clustering, same-start batches, and nonrandom observation gaps violate a simple independent-sampling interpretation.

No return, executable profit, or economic-significance estimate is reported, so a per-$10,000 or annualized net-performance calculation would be incoherent. Fees, bid–ask spreads, depth-dependent market impact, transfer charges and time, borrowing costs, and liquidation costs would all be indispensable before any future trading claim. The present release supplies operational timestamps and observation boundaries that could inform such a later study; it does not supply trade decisions, fill prices, or realized P&L. This explicit inapplicability is preferable to importing finance-journal test names without their required estimands.

8. Limitations and ethics

8.1 Coverage, annotation, and historical evidence

Dataset A is a scoped public-source map, not a complete census of exchange communications. Nine venues have full-window or historical gaps, one observable venue has zero qualifying events, and source accessibility can differ by region and over time. Deleted or edited posts, category migration, regional pages, private channels, and sources outside the documented inclusion rule may be missing. The OKX reconciliation is limited to the stated en-ae scope and three categories; it cannot validate the other venues. Source counts also measure retrieved labels rather than the total communication effort of a venue.

The annotation was performed by one person. The 88-item single-person review concerned an earlier version and cannot supply inter-annotator agreement or a current error rate. Semantic matching in Dataset B has the same limitation. Instrument and chain strings are not normalized into a universal asset ontology, and eight supersession links do not prove that every edit chain was observed. Public rows contain URLs and structured labels but not announcement text or private evidence excerpts, so external users can inspect sources subject to their current availability but cannot reproduce every historical annotation decision from the release alone.

Dataset B is short, discontinuous, and derived from only four venues. The approximately 13.2-hour known gap, earlier lost temporary data, ordinary polling intervals, and one unusually long round constrain what can be detected. Raw logs and raw state files are not released. The supplied derived JSON supports exact recomputation of the reported tables, but not independent reconstruction of the collection path, venue response times, or every manual match. A future polling run will observe a new sample and should not be presented as bitwise reproduction of the historical one.

8.2 Measurement and external validity

The principal measurement risk is not random noise but mismatched meaning. Publication time, text-stated effective time, first public visibility, actual status transition, and polling detection are different. Even a correctly matched announcement and key do not identify the latent flip time. Q1 can be inflated by unobserved announcements or overly strict matching and reduced by overly broad matching. Strict Q2 zero is bounded by key availability and polling resolution. Q3 conditions on matched detected events and says nothing about unmatched announcements or undetected transitions. Durations omit or separately classify intervals whose start, end, or continuity is not established.

Dependence is substantial. A single chain announcement may map to hundreds of keys, repeated field oscillations may create artificial episodes, and venue event mix changes across the window. Time splits, venue exclusions, bound substitutions, and batch diagnostics reveal sensitivity but cannot create independence or repair missing observations. No venue ranking, causal effect, market-impact estimate, return forecast, or execution recommendation follows. External validity beyond the venues, access surfaces, and dates documented here remains untested.

8.3 Responsible release and reuse

The release redistributes structured labels and public URLs rather than announcement bodies, forum posts, Telegram text, private notes, user data, credentials, or nonpublic interfaces. Automated reproduction should remain low frequency, comply with current terms and robots policies, and stop at access restrictions. It must not bypass 403 responses, device-signature walls, or other controls. URLs may change or disappear; consumers should preserve their own lawful retrieval dates and avoid treating later content as the historical snapshot.

The public percentage-decoding normalization restores the single-letter forms %65 and %6F to their semantic characters and restores the obfuscated name fragment to Codex. These transformations affect public strings only and do not add evidence fields. The redacted schema deliberately retains has_note while excluding note content and evidence excerpts. Reusers must not infer or fabricate the missing text from other sources. This paper was reviewed against the confidentiality boundary before release; it describes public mechanisms, observational methods, and reproducible public-data results rather than production rules or deployment details.

9. Conclusion

This paper separates exchange operational announcements from public transfer-status observations and supplies two redacted datasets with explicit denominators, timestamp semantics, and coverage boundaries. The announcement sample documents substantial variation in stated scheduling lead and shows why null, date-only, and retrospective effective times must remain distinct. The polling study shows that unmatched key-transition records are concentrated by venue and batch, that announcement–detection lag can change sign across venues, and that a precise pooled duration median can be fragile to record composition.

The resulting contribution is a measurement framework and reproducible evidence base, not a trading system. Future work can extend independent annotation, preserve richer lawfully collected response-time data, and study how operational constraints affect execution or prices only after defining the corresponding outcome and identification design. Until then, unknown should remain unknown, a gap should remain a gap, and publication, effectiveness, visibility, transition, and detection should not be collapsed into one timestamp.

Declarations

Corresponding author. Hang Liu, humarb Research, London, United Kingdom. Research channel: https://humarb.ai/research/; institutional website: https://humarb.ai; ORCID: 0009-0008-1059-3926.

Competing interests. The author declares no competing financial or non-financial interests directly related to this descriptive dataset study.

Funding. This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Data availability. The redacted announcement labels cover 2026-04-23 to 2026-10-20 UTC and were annotated on 2026-10-21 and 2026-10-22. The derived polling audit covers a log span from 2026-10-21 13:02:10 UTC to 2026-10-28 14:05:59 UTC. The release includes public labels, public URLs, coverage documentation, and derived audit, lag-bound, and duration JSON. Raw polling logs, full announcement snapshots, private evidence excerpts, and note text are unavailable in the public package for the reasons stated in Sections 4 and 8.

Code availability. The accompanying reproduction package contains the public-schema validator and one script that recomputes all machine-readable results, seven CSV tables, and eight SVG figures. It also contains a data dictionary, acquisition guidance, environment details, and output-to-paper mapping. New collection is optional and is not required to reproduce the reported descriptive results.

AI assistance. Generative AI assisted with language editing, code drafting, and literature-discovery support. Numerical results were recomputed from the released files with the included scripts, and references were checked against scholarly metadata or official documentation. The author reviewed the work and is responsible for all content, interpretations, and remaining errors.

Series acknowledgment. This paper was completed by humarb Research; humarb.ai is its institutional website and https://humarb.ai/research/ is its research channel. It complements, but does not validate the return estimates of, humarb Research 01 and 02.

Not investment advice. This preprint is for research and dataset documentation only. It is not investment advice, an execution recommendation, a solicitation, or a claim of profitable or risk-free trading. It has not been peer reviewed.

References

Ackerer, D., Hugonnier, J., & Jermann, U. (2026). Perpetual futures pricing. *Mathematical Finance, 36*(3), 481–499. https://doi.org/10.1111/mafi.70018

Artstein, R., & Poesio, M. (2008). Inter-coder agreement for computational linguistics. *Computational Linguistics, 34*(4), 555–596. https://doi.org/10.1162/coli.07-034-R2

Binance. (n.d.). *Introduction to Binance Futures funding rates*. Retrieved November 3, 2026, from https://www.binance.com/en/support/faq/detail/360033525031

Bybit. (n.d.). *Introduction to funding rate*. Retrieved November 3, 2026, from https://www.bybit.com/en/help-center/article/Introduction-to-Funding-Rate

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

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

dYdX. (n.d.). *Funding*. Retrieved November 3, 2026, from https://docs.dydx.xyz/concepts/trading/funding

Gebru, T., Morgenstern, J., Vecchione, B., Vaughan, J. W., Wallach, H., Daumé III, H., & Crawford, K. (2021). Datasheets for datasets. *Communications of the ACM, 64*(12), 86–92. https://doi.org/10.1145/3458723

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

He, S., Manela, A., Ross, O., & von Wachter, V. (2022). Fundamentals of perpetual futures. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4301150

Hyperliquid. (n.d.). *Funding*. Retrieved November 3, 2026, from https://hyperliquid.gitbook.io/hyperliquid-docs/trading/funding

Kaplan, E. L., & Meier, P. (1958). Nonparametric estimation from incomplete observations. *Journal of the American Statistical Association, 53*(282), 457–481. https://doi.org/10.1080/01621459.1958.10501452

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

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

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

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

Turnbull, B. W. (1974). Nonparametric estimation of a survivorship function with doubly censored data. *Journal of the American Statistical Association, 69*(345), 169–173. https://doi.org/10.1080/01621459.1974.10480146

Appendix A. Symbols and estimands

Symbol or label	Definition	Unit and sample restriction
I_t, B_t, A_t	Index or oracle, impact bid, and impact ask prices	Institutional background only
P_t, bar P	Instantaneous premium and its venue-defined average	Equation (1); no estimation
r, \delta, F	Interest component, damper, and stylized funding rate	Equation (2); public mechanism illustration
p_i, e_i	Announcement publication and text-stated effective time	UTC; e_i can be second-level, date-only, or null
L_i = e_i-p_i	Publication-to-effective lead	Hours; 720 second-resolution Dataset A rows
M_j	Semantic contemporaneous announcement-match indicator	Audited key-transition record
q_1	Share with M_j=0 in auditable set A	503/929 non-Gate records; not an event rate
t_j^{det}	Recorded detection timestamp for matched event j	Round-based public-status observation
D_j = t_j^{det}-p_j	Detection-minus-publication lag	Minutes; 81 unique Q3 publication events
D_j^-, D_j^+	Released lower and upper Q3 lag bounds	Deterministic sensitivity envelope
c_k, o_k	Withdrawal-close and later withdrawal-open detections	Same public interface key
U_k = o_k-c_k	Observed close–open elapsed time	Hours; 517 non-gap-crossing closed pairs
Q2	Strict announced-but-unobserved execution class	Requires key, due time, resolution, and contrary state
Gap crossing	Endpoint pair crossing a known long observation gap	Kept outside the continuous-closure median
Right censored	Observed stop without a subsequent observed opening	Not assigned a completed duration
No observed start	Opening or recovery without an observed stop	Not treated as proven recovery duration

Appendix B. Public data dictionary

Dataset A is stored as UTF-8 JSON Lines with one announcement row per line. venue is the lowercase venue identifier; id is the unique public row identifier; url is the public HTTPS source URL; and publishedAt is the publication timestamp under the source-specific collection rule. kind takes delist, maintenance, transfer_suspend, transfer_resume, reschedule, or the schema-supported but unobserved retract. market is a multi-select list whose entries need not sum to one per row. instruments and chains are public string lists without a universal asset normalization. effectiveAt preserves a second-resolution timestamp, a date-only value, or null. supersedes links a rescheduling row where available and is otherwise null. source_type distinguishes cms, forum, status_page, and other. has_note records whether a private annotation note existed; neither that note nor evidence text is included.

announcement-audit.json records the Q1/Q2 audit and key-transition classifications. Its unit is not interchangeable with announcement events. lag-bounds.json contains unique matched publication events, recorded lags, and lower and upper timing bounds. duration.json contains continuous closed intervals and the separate gap-crossing, right-censored, and no-observed-start partitions, together with batch-sensitivity metadata. coverage.json records the target map, observable and gap venues, annotation dates, regional scope, polling timestamps, and known gaps. results.json is generated rather than manually edited and is the numerical source for the manuscript tables and figures.

The public validator checks 996 unique identifiers, the exact redacted key set, allowed category and market values, timestamp formats, HTTPS URLs, normalized public strings, and the central Q1, Q3, and duration counts. Its warnings about unavailable raw logs, private evidence, the earlier baseline, and independent annotation are intentional evidence boundaries. A PASS status means that the released public files pass these internal checks; it is not external certification of completeness or annotation accuracy.

Appendix C. Reproduction and audit map

Run python3 reproduction/scripts/reproduce.py from the delivery root to regenerate results.json, all seven tables, and all eight SVG figures. Run python3 reproduction/scripts/validate_public.py to validate the redacted public inputs. --skip-figures supports a faster numerical check. The scripts use the released local files and do not call an exchange, a private service, or an external API. The table files are named according to their order in the paper; Figures 1–8 correspond respectively to coverage, monthly composition, lead time, the observation model, Q1 composition, Q3 lag, duration sensitivity, and lead robustness.

The manuscript’s numeric claims were checked against the generated JSON rather than copied from intermediate console output. The main audit distinguishes three operations: schema validation, descriptive recomputation, and visual inspection. It does not claim a second annotator, a complete historical crawl, raw-log reconstruction, external plagiarism-database clearance, peer review, or a Zenodo identifier before one is assigned. The accompanying local text-repetition report is a screening aid, not an iThenticate or Turnitin result.