

Causal Online Validation of Transition-Conditioned Locality in Cryptocurrency Microstructure

Preprint — September 2026

Aleksandr Zhulin

Independent Researcher

xandermulder85@gmail.com

Abstract

This study asks whether a transition-conditioned Locality structure previously established retrospectively can be reformulated as an online-available object without future information and remain empirically observable under frozen validation contracts. The online construction was subjected to source-provenance, prefix-invariance, deterministic-equivalence, and information-availability audits. Three market experiments then served distinct evidentiary roles. A preregistered BTCUSDT continuity test directly linked the retrospective BTC result to the frozen causal construction (6,901 events, 8,832 controls, 744 UTC-hour clusters; incremental event coefficient 0.337, 95% CR1 CI 0.306–0.368). ETHUSDT supplied independent chronological validation (7,151 events, 8,843 controls; coefficient 0.306, 95% CI 0.282–0.331), while SOLUSDT supplied separate-market replication (6,908 events, 8,851 controls; coefficient 0.420, 95% CI 0.397–0.443). A separate forward experiment establishes finite, non-zero population-level persistence in ETH and SOL under simultaneous trajectory inference. The result is structural: it establishes online observability and continuity of the frozen Locality object under the tested contracts, not signed-price prediction, profitability, execution advantage, or economic causation.

Keywords: market microstructure; cryptocurrency; online validation; causal information set; event-conditioned analysis; locality; order flow; BTCUSDT; ETHUSDT; SOLUSDT

1. Introduction

Retrospective market structure and online market structure are not the same scientific object. A pattern may be stable, reproducible, and economically suggestive after an episode is complete while still depending on observations that were unavailable at the timestamp assigned to the episode. For transition-conditioned microstructure research, that distinction is not a software detail: it determines whether the object can legitimately be described as observable at signal time.

The preceding study in this series reported a localized transition-conditioned organization in BTCUSDT and ETHUSDT over 1 January–1 October 2025 (Zhulin, 2026). The completed field combined trailing-normalized price, activity, and order-flow information and differed from eligible random locations through a conjunction of near-transition excess, far-boundary attenuation, greater compactness, and lower spread. The same directions persisted through four chronological partitions in each market. That paper also imposed a strict boundary: confirmation and the symmetric completed representation used later observations, so the result was retrospective rather than online. A one-sided or progressively available construction was explicitly left as a new experiment rather than inferred from the historical result (Zhulin, 2026).

Paper II addresses that unresolved information-availability question. The target is not a new Locality formula and not a generic claim that online state detection is novel. The question is narrower: can the specific retrospective Locality object be reformulated under a frozen past-only information contract, pass implementation-level causality checks, survive independent chronological validation and separate-market replication, and preserve direct continuity in BTC, the market central to the retrospective evidence? A secondary question asks whether the resulting structural contrast persists beyond the detection instant for a finite interval.

This separation is important because contemporary microstructure research already contains strong precedents for as-of features, held-out evaluation, event conditioning, cross-symbol analysis, and persistence. Jeon (2026), for example, evaluates pre-event L2 and order-flow information in rolling held-out folds and explicitly measures incremental value relative to staged baselines. Karpov (2026) uses a preregistered as-of design with purging, embargoing, negative controls, and cross-symbol evaluation for a different predictive object. da Cruz et al. (2026) show, from a structural perspective, that observable liquidity deformation can be empirically meaningful while remaining largely decoupled from short-horizon price translation. These studies remove any basis for claiming novelty from online inputs, validation methodology, or the general distinction between structure and price movement.

The contribution tested here is therefore object-specific and empirical. First, the previously retrospective Locality object is converted into a past-only signal-time construction and frozen. Second, the frozen object is evaluated on a later fixed ETHUSDT chronology. Third, the same construction is applied to SOLUSDT without pooled inference or market-specific retuning. Fourth, a separately preregistered BTCUSDT continuity test asks whether the retrospective BTC lineage survives the final causal validation contract. Fifth, a forward experiment characterizes persistence with simultaneous uncertainty. Mechanical validity and scientific validity remain separate gates: deterministic execution cannot rescue a failed empirical test, and a strong empirical contrast cannot rescue a construction that uses unavailable information.

The evidence supports the online structural claim in all three tested markets, but for different reasons: BTC supplies retrospective-to-causal continuity, ETH supplies independent chronological validation, and SOL supplies separate-market replication. It does not establish signed-price direction, alpha, profitability, execution feasibility, or a causal market mechanism. Those distinctions remain binding throughout the paper.

Section 2 places the experiment relative to event-study, resiliency, predictive-state, and structural-liquidity work. Section 3 defines the data, chronology, and evidence provenance. Section 4 specifies the causal information contract and the distinct BTC, ETH, and SOL validation roles. Section 5 reports BTC continuity, ETH chronological validation, SOL cross-market replication, and forward persistence. Section 6 records robustness, falsification, and audit evidence. Sections 7–9 discuss interpretation, limitations, disclosure boundaries, and conclusions.

2. Related Literature

2.1 Event-conditioned microstructure and resiliency

Changes in order submission, spreads, depth, imbalance, and liquidity around events are established empirical phenomena. Erenburg and Lasser (2009) examine limit-order-book behavior around macroeconomic news; Rodrigues, Souza, and Stevenson (2012) study intraday microstructure before takeover announcements; and Gomber, Schweickert, and Theissen (2015) use an event-study framework to examine liquidity shocks and recovery. Resiliency work similarly studies how electronic books replenish or recover after disturbances (Large, 2007; Lo and Hall, 2015; Clapham, Haferkorn, and Zimmermann, 2020). These precedents eliminate broad novelty claims for event conditioning, event/control comparison, recovery, or persistence itself.

The present experiment uses event-conditioned populations but asks a different question. The object being tested is not whether a known disturbance changes liquidity, nor whether a generic shock decays. It is whether a specific Locality structure first established ex post remains observable after its information set is restricted to what exists by signal time. Persistence is secondary and object-specific.

2.2 Online/as-of validation and state prediction

Predictive microstructure research routinely separates training information from held-out outcomes. Jeon (2026) is especially relevant because the study uses BTCUSDT and ETHUSDT futures, staged baselines, rolling later-month evaluation, and event-clustered uncertainty to ask whether order flow adds information beyond a pre-event liquidity-state representation. Its dependent variable is a post-event liquidity state and its scientific task is prediction. The present study does not compete on that target. It uses a fixed structural object inherited from Paper I and asks whether that object remains distinguishable from its eligible background comparator after causal reformulation.

Karpov (2026) provides another methodological boundary. That study evaluates temporal-topology features for short-horizon tail-move prediction under an as-of, purged, embargoed, preregistered design and interprets a null primary contrast conservatively. The overlap with Paper II lies in information discipline and falsification-first evaluation, not in the observable, target, or mechanism. Past-only construction and held-out validation are therefore treated here as standard methodological requirements rather than contributions.

2.3 Structural observables versus price translation

A separate literature treats order-book and flow representations as structural objects rather than direct price forecasts. da Cruz et al. (2026) develop a geometric description of directional liquidity in BTC and ETH and report that liquidity deformation is largely decoupled from mid-price drift. Their observable and data representation differ from Locality, but the conceptual boundary is directly relevant: a reproducible structural state need not imply a stable directional-price signal. Paper II preserves that distinction by defining its central claim before any downstream directional interpretation.

The same caution is consistent with the broader order-flow literature. Order-flow imbalance can explain or predict short-horizon price changes under particular constructions (Cont, Kukanov, and Stoikov, 2014), but that does not imply that every structural flow-conditioned state carries incremental directional information. Accordingly, the regression used below tests distinguishability of the frozen Locality event population relative to its prespecified baseline and controls; it is not interpreted as a price-impact equation.

2.4 Position of the present study

The reviewed literature therefore removes novelty from every generic ingredient: event conditioning, as-of observables, held-out chronology, cross-market evaluation, regression with controls, and persistence. The surviving contribution is the empirical lineage of one object: retrospective transition-conditioned Locality in Paper I; frozen past-only reformulation; direct BTC retrospective-to-causal continuity; independent ETH chronological validation; separate-market SOL replication; and finite forward structural persistence. No “first”, “unique”, or equivalent priority claim is required.


3. Data and Evidence Provenance

3.1 Markets and experimental chronology

Paper I used BTCUSDT and ETHUSDT from 1 January through 1 October 2025 to establish the retrospective Locality result (Zhulin, 2026). The online branch was then separated into development, validation, replication, continuity, and forward-characterization stages. These stages have different evidentiary roles: the early development gate is not independent confirmation; ETH supplies independent chronological validation; SOL supplies separate-market replication; the later preregistered BTC experiment supplies direct retrospective-to-causal continuity under the final validation contract; and the forward experiment supplies only persistence characterization.

Stage	Market(s)	Chronology / role	Scientific status	Public role
Paper I foundation	BTCUSDT, ETHUSDT	1 Jan–1 Oct 2025	Retrospective result established	Prior result; online availability not established
Causal development gate	BTCUSDT, ETHUSDT	Early fixed development interval	Development only	Defines/falsifies implementation; not independent validation
Chronological validation	ETHUSDT	8 Jan–8 Feb 2025	PASS	Primary later-chronology evidence
Cross-market replication	SOLUSDT	Same frozen validation-era contract	PASS	Primary separate-market evidence
BTC continuity	BTCUSDT	8 Jan–8 Feb 2025	PASS	Direct retrospective-to-causal continuity
Forward persistence	ETHUSDT, SOLUSDT	1 Oct–1 Nov 2025	PASS	Secondary temporal characterization

Table 1. Experimental sequence. Development, later chronology, market replication, and persistence answer different questions and are not counted as interchangeable replications.



Three positive market results have distinct evidentiary roles; they are not pooled replications.

Figure 1. Evidence architecture from retrospective Paper I Locality to the frozen causal online object. BTC supplies retrospective-to-causal continuity, ETH independent chronological validation, and SOL separate-market replication; forward persistence is a secondary ETH/SOL characterization. The three market results are not pooled replications.

3.2 Evidence hierarchy and source binding

Numerical authority follows the same hierarchy used in Paper I: fixed source/data and primary execution artifacts govern over generated reports, registries, manuscript text, or model recollection. Before drafting the present Results, the BTC continuity CSV set, ETH validation CSV set, SOL replication CSV set, and both forward-lifetime CSV sets were rebound or reconciled against the frozen evidence package. Primary CSV values govern all numbers reported below.

This distinction matters because the research archive contains exploratory notes, development verdicts, deterministic reruns, and later summaries. Those objects are useful for provenance but are not independent evidence. A deterministic rerun establishes reproducibility of the same computation; it does not create a second scientific replication. Likewise, causal audit, prefix invariance, common support, chronological validation, market replication, and persistence answer different questions and are not added together as if they were independent tests of one null.

3.3 Public and private information layers

The manuscript follows a minimum-scientific-information disclosure rule. Markets, validation chronology, event/control counts, cluster counts, comparator meaning, inference family, primary coefficients, confidence intervals, causal information semantics, and the existence of finite persistence are public because they are required to evaluate the claim. Exact production detector formulas, tuned feature composition and weights, operational thresholds, reconstructive implementation sequence, exact operational lifetime, code, and trading-system integration remain private. This is the same disclosure logic used in Paper I: the scientific object must be auditable without materially shortening reconstruction of the production detector.

4. Methodology

4.1 From retrospective Locality to a signal-time object

Paper I defined Locality from a completed transition-conditioned representation. Its scientific limitation was explicit: the completed field used information from both sides of the nominal transition and confirmation completed after the historical anchor. Paper II does not relabel that field as causal. Instead, it constructs a distinct online object whose signal-time inputs are restricted to information physically available at or before the relevant time.

The online object preserves the scientific lineage of Locality—transition-conditioned organization relative to recent local market conditions—without publishing the detector recipe. The public description is intentionally functional rather than reconstructive. Price movement, trading activity, and signed-flow families remain the observable basis inherited from the prior programme; their online transforms use trailing information. Event selection is frozen before later validation. No full-sample normalization, future label, or backward assignment is permitted in the signal-time construction.

4.2 Information-set causality

Throughout this paper, “causal” means causal with respect to information availability, not causal in the economic or mechanistic sense. For an online quantity $Z(t)$, admissibility requires that every input needed to compute $Z(t)$ be available no later than t under the frozen construction. The audit separately records future dependence, full-sample dependence, and backward timestamp assignment. A quantity failing any of these requirements cannot support the signal-time claim.

The causal audit separates signal-time construction quantities from intentionally forward validation outcomes. In the ETH and SOL validation lineage, the audited signal-time construction objects show no future dependence, no full-sample dependence, and no backward assignment. The BTC continuity package independently records the same causal requirements, together with a prefix-invariance test in which appending future rows produced zero mismatches and zero numerical differences in required causal prefix quantities. Forward outcomes are not inputs to event selection and are allowed to use post-signal observations because their role is to evaluate what follows an already frozen event.

Information-availability audit: construction inputs versus validation outcomes

	Future dependency	Full-sample dependency	Backward timestamp
Signal-time construction 8 audited objects	NONE	NONE	NONE
Forward validation outcomes 3 audited objects	INTENTIONAL	NONE	NONE

Forward outcomes are evaluation targets, not inputs to event selection.

Figure 2. Information-availability audit. All eight signal-time construction objects are free of future, full-sample, and backward-timestamp dependence. Three forward outcome objects intentionally use later observations for validation but are not inputs to event selection.

4.3 Falsification architecture

Paper II retains the falsification architecture introduced in Paper I. The research funnel imposes successive gates: define and freeze the operational object; verify deterministic behavior and signal-time information availability; apply the relevant comparator and support prerequisites; then evaluate the frozen object under its designated chronological, continuity, replication, and forward-characterization tests. Failure at any stage removes or narrows the corresponding candidate claim rather than being rescued by a later positive result.

Mechanical validity and scientific validity remain separate. Correct execution cannot rescue a failed empirical test, and an empirical contrast cannot rescue a construction that uses unavailable information. Negative downstream results are retained when they materially narrow interpretation, but they do not falsify scientific claims that those tests were not designed to evaluate.

4.4 Frozen event and control populations

An event is an observation selected by the frozen causal event construction. A control is an eligible non-event/background observation processed under the same comparison design. Controls are not described as matched because no matching procedure is used. Event and control populations are kept conceptually separate from the baseline covariate: the former define the comparison groups, whereas the latter enters the prespecified regression as a continuous reference variable.

Eligibility and density requirements are frozen prerequisites. A common-support gate is evaluated before the incremental regression is admitted. This gate prevents a coefficient from being interpreted when the event population lies outside the usable baseline support of controls. The exact operational support thresholds are private; the public evidence reports whether the prespecified gate passes. It passes in the final BTC continuity test, ETH chronological validation, and SOL replication. The earlier BTC development run failed this prerequisite and remains recorded separately.

4.5 Primary regression and inference

For each market, the primary validation model has the form

$$Y_i = \alpha + \beta X_i + \gamma E_i + \varepsilon_i,$$

where E identifies the frozen event population, X is the prespecified causal baseline covariate, and Y is the prespecified forward structural validation outcome. The coefficient γ is the incremental EVENT association conditional on X . The equation is deliberately simple: the scientific question is whether the frozen online event population remains distinguishable after accounting for the baseline, not whether a flexible predictive model can be optimized on the validation interval.

Inference uses ordinary least squares with CR1 covariance clustered by UTC anchor hour. The reported 95% interval uses the cluster degrees of freedom. Clustering recognizes dependence among observations sharing an anchor hour without pretending that each event/control row is an independent experiment. The standardized effect G is reported as γ divided by the control-population standard deviation of Y ; it is a scale description, not a separate test.

4.6 Preregistered BTC continuity test

BTC is assigned a different evidentiary role from ETH and SOL. Paper I already contained retrospective BTC Locality evidence, while the early causal-development run failed its common-support prerequisite and therefore did not establish causal BTC validation. A later BTC continuity test was preregistered before candidate scientific output was inspected. It used 8 January through 8 February 2025, the first BTC Online Locality chronology following the consumed development interval and the same calendar interval used by the final ETH/SOL validation contract. The interval was unconsumed by prior BTC Online Locality testing, although it overlaps historical chronology used retrospectively in Paper I and is therefore not described as globally unseen market history.

The BTC wrapper was mechanically derived from the frozen validation implementation without BTC-specific changes to event construction, baseline, refractory logic, controls, forward outcome, common-support calculation, regression, covariance, or materiality gates. Before scientific reveal, contract equivalence, causal audit, prefix invariance, full production, an independent full rerun from zero, and deterministic/semantic equivalence were required to pass. The scientific question was continuity only: whether the retrospective BTC object survives the already frozen causal construction under the final validation contract.

4.7 Independent chronological validation

The ETH validation interval is fixed at 8 January 2025 through 8 February 2025. The validation wrapper reuses the authoritative frozen development implementation after a source-provenance lock; it changes the authorized chronology rather than reimplementing the scientific object. The final scientific verdict is admitted only after production execution, deterministic full rerun/equivalence, causal provenance audit, and prefix-invariance checks. The later chronology is therefore a test of the frozen object rather than a tuning interval.

4.8 Cross-market replication

SOLUSD is evaluated with the same frozen causal validation logic. The replication is market-separated: no pooled coefficient is estimated and no two-market average is used to manufacture significance. The claim is correspondingly narrow. A SOL pass shows that the frozen object survives in a second named market under the tested contract; it does not establish crypto-wide portability, equal effect size, or a common economic mechanism.

4.9 Forward persistence protocol

A separate forward-lifetime experiment asks whether the structural EVENT contrast exists only at detection or remains supported for a finite interval. At each prespecified forward evaluation, the same basic event/baseline regression is estimated on the lag-specific finite sample. Missingness is preserved by lag rather than converted into a complete-case panel. This matters because later evaluations can have different data availability and should not silently redefine the early population.

Family-wise uncertainty over the trajectory is controlled with a whole-cluster bootstrap. The bootstrap resamples UTC-hour anchor clusters, preserves duplicate sampled clusters as distinct bootstrap clusters, recenters the EVENT coefficient around its original estimate, studentizes with CR1 standard errors, and uses the maximum absolute t statistic across the trajectory. Both market runs requested 1,000 repetitions, obtained 1,000 valid repetitions, and recorded zero bootstrap failures. Lifetime is defined by a frozen sustained-crossing rule applied to the simultaneous band. The exact operational duration and timing parameters remain private; the public result is whether persistence is non-zero and finite.

4.10 Deterministic reproducibility and prefix invariance

Execution validity is a prerequisite, not evidence of market structure. The validation sequence therefore distinguishes source identity, deterministic equivalence, and prefix invariance from the empirical gate. Source provenance checks that the later wrapper calls the frozen object rather than a modified implementation. Deterministic equivalence checks that authorized repeated execution reproduces the same scientific outputs. Prefix invariance checks that extending the available history does not revise already computed signal-time quantities. Only after these mechanical conditions pass is the market-level coefficient interpreted.

5. Results

5.1 Causal construction verification

The causal implementation gate passes across the final evidence lineage. ETH and SOL share the frozen causal classification structure, while the preregistered BTC continuity package independently passes contract equivalence, causal provenance, prefix invariance, full-rerun equivalence, and final validation. In the BTC prefix audit, 9,000 prefix rows were compared before and after 3,000 future rows were appended; every required causal quantity had zero mismatches and zero maximum numerical difference. Mechanical validity is established before the market-level scientific verdicts are interpreted.

This result is necessary but not sufficient. It establishes that the tested event population can be generated online under the frozen information contract. It says nothing by itself about whether the population is structurally distinct from eligible controls. That question is answered by the market-level validation gates below.

5.2 BTCUSDT retrospective-to-causal continuity

The preregistered BTC continuity test passes the final frozen Online Locality validation contract. The sample contains 6,901 EVENT observations and 8,832 eligible controls across 744 UTC-hour clusters (15,733 regression observations). Density and common support pass. The prespecified forward-decay materiality criterion is satisfied. In the prespecified regression, the incremental EVENT coefficient is 0.336910 with CR1 standard error 0.015873 and 95% interval [0.305749, 0.368072]. The standardized magnitude is $G = 0.428$ and satisfies the prespecified materiality criterion.

The result closes a specific continuity gap rather than adding a third interchangeable replication. Paper I established retrospective BTC Locality; the early causal-development BTC run later failed common support; the preregistered final-contract test now shows that the BTC lineage survives the frozen past-only construction. The earlier development failure remains part of the record and is not retroactively converted into a pass.

5.3 ETHUSDT independent chronological validation

The frozen online object passes the later ETHUSDT chronology. The primary sample contains 7,151 EVENT observations and 8,843 eligible controls across 744 UTC-hour clusters. The density and common-support prerequisites both pass. The incremental EVENT coefficient is 0.306149 with CR1 standard error 0.012523; the 95% cluster-robust interval is [0.281564, 0.330734]. Relative to the control outcome dispersion, the standardized magnitude is $G = 0.452$.

The interpretation is deliberately conditional. The coefficient shows that, on the prespecified later chronology and conditional on the frozen baseline, the online Locality event population remains distinguishable from its eligible control population in the forward structural outcome. It is not a coefficient of market causation and is not a price-return forecast.

5.4 SOLUSDT separate-market replication

The same frozen validation logic also passes in SOLUSDT. The replication sample contains 6,908 EVENT observations and 8,851 eligible controls across 744 UTC-hour clusters. Density and common support again pass. The incremental EVENT coefficient is 0.419848 with CR1 standard error 0.011850; the 95% interval is [0.396585, 0.443110], and the standardized magnitude is $G = 0.604$.

The SOL coefficient is numerically larger than the ETH coefficient, but no equality or difference test was prespecified and no ranking is inferred. The replication claim concerns sign, support, and survival of the same frozen scientific test in a separate market, not equality of effect size.

Market	Role	EVENT	Control	UTC-hour clusters	γ EVENT	CR1 SE	95% CI	Std. G	Support
BTCUSDT	Continuity	6,901	8,832	744	0.336910	0.015873	[0.305749, 0.368072]	0.428	PASS
ETHUSDT	Later chronology	7,151	8,843	744	0.306149	0.012523	[0.281564, 0.330734]	0.452	PASS
SOLUSDT	Separate market	6,908	8,851	744	0.419848	0.011850	[0.396585, 0.443110]	0.604	PASS

Table 2. Primary Online Locality evidence under the final causal validation contract. BTC supplies retrospective-to-causal continuity; ETH supplies independent chronological validation; SOL supplies separate-market replication. γ is the incremental EVENT association conditional on the frozen baseline. Standardized G is descriptive; support PASS is a prerequisite, not an additional independent confirmation.




Figure 3. Market-separated causal Online Locality evidence under distinct roles. BTCUSDT is the preregistered retrospective-to-causal continuity test, ETHUSDT the independent chronological validation, and SOLUSDT the separate-market replication. Points show incremental EVENT coefficients; bars are 95% CR1 intervals clustered by UTC anchor hour. No pooled coefficient is estimated.

5.5 Forward structural persistence

The forward experiment contains 6,923 ETH EVENT and 8,846 control observations and 7,054 SOL EVENT and 8,845 controls; each market spans 744 UTC-hour clusters. At the earliest prespecified forward evaluation, ETH has $\gamma = 0.280335$ with simultaneous 95% interval [0.241337, 0.319332], while SOL has $\gamma = 0.264804$ with simultaneous interval [0.236557, 0.293050]. Both simultaneous lower bounds remain above zero.

The trajectory subsequently reaches the frozen sustained-crossing condition in both markets and is not right-censored at the experiment boundary. The resulting lifetime is therefore finite and non-zero in both ETH and SOL. Exact operational durations are withheld because they materially shorten reconstruction of the production timing logic and are not necessary to evaluate the central online-observability claim.




Figure 4. Public-safe persistence evidence. At the earliest prespecified forward evaluation, the EVENT contrast is above zero in ETHUSDT and SOLUSDT under simultaneous family-wise intervals. Both trajectories later satisfy the frozen sustained-crossing rule at a finite horizon; exact operational duration is withheld.

5.6 What the combined result establishes

The positive components answer different questions. BTC establishes direct continuity from the retrospective market central to Paper I into the frozen causal construction. ETH establishes that the frozen causal object survives independent chronological validation. SOL establishes that the result is not confined to BTC or ETH under the tested design. The lifetime experiment establishes that the structural contrast is not purely instantaneous at detection. Their conjunction supports a causal/online structural object across three tested markets, with finite persistence established separately in ETH and SOL.

The result is narrower than a forecasting claim. The dependent validation objects are structural, and no signed-price target enters the central test. A structural state may be stable, replicable, and temporally persistent without carrying incremental directional information. That boundary is consistent with the distinction between liquidity deformation and price translation emphasized by da Cruz et al. (2026).

6. Robustness, Falsification, and Audit Evidence

6.1 Development failure remains visible

The initial causal development gate did not produce a two-market causal validation claim. In BTCUSDT, the development population contained 1,515 events, 1,999 controls, and 168 clusters, but the common-support prerequisite failed. That result remains a development failure and is not rewritten. The later BTC continuity PASS belongs to a separate preregistered experiment under the final validation contract, selected before candidate BTC scientific output was inspected. The contrast between the two results is therefore part of the provenance rather than a post-hoc rescue.

6.2 Why mechanical audits are not replications

Source-provenance, deterministic-equivalence, and prefix-invariance checks strengthen confidence that the intended computation was executed without look-ahead or silent mutation. They do not create additional market observations. The manuscript therefore distinguishes four concepts: software-level deterministic equivalence; BTC retrospective-to-causal continuity; ETH independent chronological validation; and SOL separate-market replication. None is counted as an interchangeable replicate of the others.

6.3 Comparator and support discipline

The eligible control population represents background observations under the frozen comparison design. It is not a condition-matched counterfactual and cannot identify a mechanism. The common-support prerequisite addresses a narrower statistical concern: the baseline covariate must provide sufficient overlap for the conditional EVENT coefficient to be interpretable under the prespecified model. BTC, ETH, and SOL pass this prerequisite under the final validation contract, whereas the earlier BTC development branch did not.

6.4 Negative downstream evidence and claim boundary

Separate downstream experiments were conducted after the online structural result. A frozen absolute price-opportunity test did not establish the prespecified sustained displacement structure in ETH or SOL. A later signed-price Direction gate also failed its preregistered cross-market criterion. This test evaluated one prespecified directional formulation at a single forward horizon; it does not establish the absence of directional or predictive information under other, untested specifications. These experiments test different objects and therefore do not falsify Online Locality, but they materially constrain interpretation: the structural result must not be presented as evidence of directional price prediction or a deployable trading edge.

Because those downstream branches are not the scientific question of Paper II, their detailed coefficients, timing, and implementation are not reported here. Their role is the same as excluded branches in Paper I: negative evidence is retained when it narrows the permitted claim, but it is not allowed to turn a one-claim paper into a catalogue of later research.

Question	Falsification pressure	Observed outcome	Publication consequence
Is the signal-time object genuinely past-only?	Dependency, provenance, prefix-invariance audits	PASS	Online-available construction admitted
Does later ETH chronology preserve the frozen contrast?	Fixed later interval + support + clustered inference	PASS	Primary chronological validation
Is the result confined to ETH?	Same frozen gate in SOL	PASS	Separate-market replication
Is the state only instantaneous?	Forward trajectory + simultaneous bootstrap band	Finite non-zero persistence	Secondary persistence claim

Did the earlier BTC development branch satisfy all prerequisites?	Common-support prerequisite	FAIL	Development failure retained; no causal BTC claim from that run
Does retrospective BTC Locality survive the final causal contract?	Preregistered continuity test + frozen support/materiality gates	PASS	Direct BTC retrospective-to-causal continuity
Does structure imply downstream price opportunity/direction?	Separate frozen downstream gates	Not established under tested contracts	No price/trading inference

Table 3. Falsification and claim-boundary map. Each row targets a different failure mode; PASS outcomes are not counted as interchangeable independent tests.

6.5 Statistical and numerical reconciliation

All public validation coefficients, intervals, counts, and cluster totals were copied from or rechecked against frozen primary artifacts. The BTC continuity regression contains 15,733 observations and 744 clusters; ETH contains 15,994 observations and 744 clusters; SOL contains 15,759 observations and 744 clusters. All use the same three-parameter structural regression family with CR1 clustered inference under their frozen contracts. The BTC production run and independent rerun were semantically identical before final scientific reveal. Forward-lifetime summaries agree with their trajectory and bootstrap artifacts, and both bootstrap runs complete all 1,000 requested repetitions without failure.

The manuscript does not use remembered aggregate values to construct quantitative displays. The principal coefficient figure is based on the frozen BTC, ETH, and SOL regression evidence; the persistence display is based on the forward trajectory artifacts; and the information-availability display summarizes the dependency audit. This preserves the artifact-over-narrative rule used in Paper I.

7. Discussion

7.1 What changes relative to Paper I

Paper I established a retrospective structural fact: completed transition-conditioned fields in BTC and ETH had a localization pattern not reproduced by eligible random locations under the same representation. It explicitly stopped short of online availability. Paper II changes the status of the object rather than merely adding robustness. Under a new past-only construction, the event population can be generated without future information; BTC now provides direct retrospective-to-causal continuity, ETH supplies independent chronological validation, and SOL supplies separate-market replication.

That distinction is scientifically substantive because look-ahead can enter microstructure research through apparently innocuous operations: centered windows, full-sample standardization, future-completed events, or timestamps assigned earlier than the information needed to compute them. The present audit treats information availability as part of object definition. A timestamp alone does not make a quantity online; all dependencies must be available by the signal time.

7.2 Structural state is not directional edge

The strongest interpretation supported by the evidence is structural. The online Locality event population differs from eligible controls in the prespecified forward structural response, and the contrast persists for a finite interval. Nothing in that statement requires or implies a stable signed-price effect. This separation is consistent with da Cruz et al. (2026), whose liquidity-geometry analysis distinguishes internal order-book deformation from price translation, and with Jeon (2026), whose staged design requires each information layer to demonstrate incremental predictive value rather than inheriting it from a correlated state representation.

This boundary also clarifies why downstream negative tests do not invalidate the paper. A structural observable can be real and reproducible while failing a particular directional target, horizon, or economic criterion. Conversely, a profitable rule would require evidence beyond structural observability: signed outcomes, economic magnitude relative to costs, execution feasibility, unseen chronology, and prospective/live validation. Those layers are outside the present claim.

7.3 What finite persistence means

The persistence result is population-level. It describes how long the event/control structural contrast remains supported under the frozen lifetime protocol, not how long any individual event remains in a latent regime. This is not a reversal of Paper I's negative temporal state-trust branch: that branch tested a different degradation/trust

quantity against a random comparator under its own frozen criterion, whereas the present experiment measures persistence of the causally reformulated EVENT/control structural contrast under a separate forward protocol. It is also not a trading hold time. A position rule would require a directional target, an entry/exit contract, cost modeling, and execution evidence. None is supplied by the lifetime estimate itself.

The use of simultaneous trajectory inference matters here. Pointwise intervals at many forward evaluations would make the first apparent crossing sensitive to repeated inspection. The whole-cluster max- $|t|$ bootstrap instead controls the family of trajectory comparisons under the frozen protocol. The resulting finite crossing therefore has a different evidentiary status from visually inspecting a decaying curve and choosing a convenient endpoint.

7.4 Replication language

The evidence uses four distinct reproducibility concepts. Deterministic equivalence means the same computation reproduces the same authorized output. BTC continuity asks whether a market central to the retrospective result survives the frozen causal reformulation under a preregistered final contract. Independent chronological validation evaluates the frozen object on a later fixed chronology without redefining it. Cross-market replication applies the same frozen scientific gate to a separate market. Keeping these terms separate prevents software reproducibility, continuity, chronology, and market replication from being collapsed into a single replication count.

7.5 Relation to existing literature

The study sits between descriptive event-conditioned microstructure and predictive state modeling. Like event studies and resiliency research, it conditions on a defined population and examines a post-event structural response (Large, 2007; Gomber et al., 2015). Like contemporary predictive studies, it enforces as-of information and later evaluation (Jeon, 2026; Karpov, 2026). Unlike those predictive studies, however, the central target is not price or a post-event liquidity class. The scientific question is whether a previously retrospective structural object survives causalization, direct origin-market continuity testing, independent chronology, and separate-market replication.

For that reason, the paper makes no mathematical novelty claim. Regression, clustered covariance, bootstrap bands, event/control comparison, and persistence are established tools. The contribution is the empirical object lineage and the fact that the frozen object survives the specified tests.

8. Limitations and Disclosure Boundary

The final causal evidence covers three named markets with non-interchangeable roles. BTCUSDT supplies a preregistered retrospective-to-causal continuity test, ETHUSDT supplies independent chronological validation, and SOLUSDT supplies separate-market replication. This does not establish universality across cryptocurrencies, venues, market regimes, or asset classes. The BTC continuity chronology was unconsumed by prior BTC Online Locality testing but overlaps historical chronology used retrospectively in Paper I; it should therefore not be described as globally unseen market history.

“Causal” is limited to the construction and information set. The event definition is computable from information available by signal time, but the experiment is observational. The event/control contrast does not identify a treatment effect, and the eligible controls are not condition-matched counterfactuals. The results therefore do not establish that Locality causes subsequent structural or price behavior.

The primary regression is intentionally low-dimensional and prespecified. It establishes incremental association conditional on one frozen baseline under the tested population. It does not exhaust alternative baselines, nonlinear specifications, horizons, or state representations. Likewise, cross-market replication in SOL does not imply equal coefficients or a common mechanism.

The forward-lifetime estimate is protocol-dependent in the ordinary scientific sense: it depends on the frozen outcome, trajectory grid, simultaneous inference rule, and sustained-crossing definition. The paper discloses the population-level method and initial simultaneous evidence but withholds the exact operational duration and reconstructive timing parameters. This limits full independent reconstruction of the production timing logic while preserving the evidence needed to evaluate the claim that persistence is non-zero and finite.

Exact detector formulas, tuned feature composition and weights, operational thresholds, implementation sequence, and code are also withheld as non-public research IP. Their release would materially shorten reconstruction of the production detector. The manuscript instead discloses observable families, past-only semantics, event/control roles, baseline distinction, validation chronology, counts, clustering, primary coefficients and intervals, support verdicts, bootstrap family, and the boundaries of the claim.

Finally, the paper is not a trading-system study. No transaction-cost model, execution simulator, position policy, capital allocation rule, or prospective live performance is evaluated. The structural result should therefore be read as a validated market-state observation, not as an investment result.

9. Conclusion

A transition-conditioned Locality structure previously established retrospectively remains empirically observable after reformulation under a strictly past-only information contract. The frozen online object passes implementation-level causality and reproducibility audits. A preregistered BTCUSDt continuity test directly links the retrospective BTC lineage to the causal construction; the object also survives independent ETHUSDt chronological validation and separate-market SOLUSDt replication. A separate forward experiment shows finite, non-zero structural persistence in ETH and SOL under simultaneous trajectory inference.

The result closes the information-availability gap left by Paper I and, through BTC, closes the direct retrospective-to-causal continuity gap as well. It does not convert a structural observation into a stronger economic claim. Online Locality, as tested here, is an observable structural state across the tested contracts. No directional-price, profitability, execution, or economic-causation claim follows from these results. The separately failed Direction gate remains compatible with this conclusion: structural observability and directional edge are different empirical questions.

Appendix A. Public Evidence and Decision Framework

A.1 Evidence hierarchy

Material results were reconciled in the following order: fixed source/data and primary numerical artifacts; deterministic execution outputs and final reports; frozen executable logic and gate protocols; registries and project documentation; manuscript text; and model interpretation or recollection. A lower layer could locate evidence but could not override a contradiction in a higher layer. The primary-artifact reconciliation for Paper II passed before final numerical Results and figures were drafted.

A.2 Claim admission rule

A material manuscript statement is admitted only when it maps to a defined scientific object, executed experiment, market, chronology, comparator or baseline, inference object, information-availability status, primary evidence artifact, limitation, and disclosure class. Claims that fail this mapping are excluded or narrowed. Related audits and metrics derived from the same object are not counted as independent confirmations merely because they produce separate files.

A.3 Terminology

Term	Meaning in this paper	Excluded interpretation
Locality	Transition-conditioned structural object inherited from Paper I and causally reformulated here	Price opportunity, direction, alpha, mechanism
Causal	All signal-time inputs available at or before signal time	Economic or physical causation
Online-available	Computable under the frozen past-only information contract	Live deployment
Control	Eligible non-event/background observation	Matched counterfactual
Baseline	Prespecified regression covariate/reference model	Control population
BTC continuity	Preregistered test linking retrospective BTC Locality to the frozen causal construction	Third interchangeable replication or globally unseen chronology
Chronological validation	Frozen object evaluated on later fixed chronology	Independent market replication
Cross-market replication	Same frozen gate in a separate named market	Universal portability
Persistence	Population-level duration of supported structural contrast	Hold time, exit rule, directional persistence

A.4 Public disclosure rule

The public layer retains the minimum information needed to identify the empirical object, audit the past-only information claim, understand event/control and baseline roles, evaluate the validation chronology and market separation, reproduce the reported statistical interpretation, and see the principal magnitudes. Exact production

detector logic, tuned composition, thresholds, reconstructive timing, exact operational lifetime, code, and trading-system integration remain private.

A.5 Status of negative evidence

The early BTC development support failure, the downstream absolute price-opportunity failure, and the downstream signed-price Direction failure remain part of the bounded evidence record. The later preregistered BTC continuity result does not erase the early support failure; it answers a new frozen continuity question under the final validation contract. Publicly, the negative downstream results constrain interpretation: Online Locality is not presented as established price opportunity or directional edge.

Generative AI Disclosure

During preparation of this work, the author used multiple cloud-hosted and locally hosted generative language-model systems for implementation assistance, language and structure review, methodological criticism, causality and reproducibility audits, and adversarial checks. The systems served in differentiated advisory roles. The author defined the research questions, conceptual objects, hypotheses, scientific direction, acceptance criteria, interpretation, and final scientific decisions. AI-assisted outputs were reviewed and edited and, where applicable, checked against executable code and fixed research artifacts. The author takes full responsibility for the manuscript.

References

- Erenburg, G., & Lasser, D. (2009). Electronic limit order book and order submission choice around macroeconomic news. *Review of Financial Economics*, 18(4), 172–182. <https://doi.org/10.1016/j.rfe.2009.06.001>
- Large, J. (2007). Measuring the resiliency of an electronic limit order book. *Journal of Financial Markets*, 10(1), 1–25. <https://doi.org/10.1016/j.finmar.2006.09.001>
- Rodrigues, B. D., Souza, R. C., & Stevenson, M. J. (2012). An analysis of intraday market behaviour before takeover announcements. *International Review of Financial Analysis*, 21, 23–32. <https://doi.org/10.1016/j.irfa.2011.05.005>
- Gomber, P., Schweickert, U., & Theissen, E. (2015). Liquidity Dynamics in an Electronic Open Limit Order Book: An Event Study Approach. *European Financial Management*, 21(1), 52–78. <https://doi.org/10.1111/j.1468-036X.2013.12006.x>
- Lo, D. K., & Hall, A. D. (2015). Resiliency of the limit order book. *Journal of Economic Dynamics and Control*, 61, 222–244. <https://doi.org/10.1016/j.jedc.2015.09.012>
- Clapham, B., Haferkorn, M., & Zimmermann, K. (2020). Does Speed Matter? The Role of High-Frequency Trading for Order Book Resiliency. *Journal of Financial Research*, 43(4), 933–964. <https://doi.org/10.1111/jfir.12229>
- Cont, R., Kukanov, A., & Stoikov, S. (2014). The Price Impact of Order Book Events. *Journal of Financial Econometrics*, 12(1), 47–88.
- da Cruz, J. P., Granate, P., Costa, D., Teixeira, A., Brinca, P., & Duarte, J. B. (2026). Directional Liquidity and Topological Mechanics in Cryptocurrency Order Books. SSRN 6673705. <https://doi.org/10.2139/ssrn.6673705>
- Jeon, J. (2026). When Does Order Flow Matter? State-Dependent L2 Liquidity-State Transitions in Crypto Futures. arXiv:2607.09230. <https://arxiv.org/abs/2607.09230>
- Karpov, A. (2026). Does Temporal Topology Add Information Beyond Static Snapshots? A Pre-Registered Controlled Study of Zigzag Persistence for Short-Horizon Tail-Move Prediction on Hyperliquid Perpetuals. SSRN 7002981. <https://doi.org/10.2139/ssrn.7002981>
- Zhulin, A. (2026). Localized Transition-Conditioned Structure in BTC and ETH Microstructure. Preprint, August 2026.