QI

METHODOLOGY WHITEPAPER

Signal Validation at Scale

Directional accuracy, risk/reward asymmetry, and regime sensitivity across 10,084 evaluated signals.

DOCUMENT	Quorum Index — Signal Validation Methodology
VERSION	v1.1 · External release
COVERAGE	10,084 signals over 45 days
WINDOW	27 Feb 2026 — 14 Apr 2026
GENERATED	18 April 2026
WEB	quorumindex.com

This paper documents the detection, classification, and validation methodology used by the Quorum Index platform to evaluate the predictive content of its market signals. Results presented here are audit-grade and reproducible from the underlying event log.

00

Executive summary

Summary of findings

Quorum Index is an autonomous market-intelligence platform that continuously scans crypto market structure for measurable deviations — score spikes, score clusters, volatility shifts, and regime transitions — and logs each one as a forward-looking signal. This paper evaluates whether those signals carry predictive content, using a fixed classification protocol applied to 10,084 events over a 45-day window.

DIRECTIONAL ACCURACY

47.75%

24h win rate (n = 8,574 classified). Below 50% as a raw hit-rate, but expectancy is positive once asymmetry is accounted for.

RISK/REWARD ASYMMETRY

1.411

R/R ratio at 24h: avg favourable +5.24% against avg adverse –3.71%. Ratio expands with horizon (1.178 at 1h → 1.411 at 24h).

REGIME SENSITIVITY

47.67%

TRANSITIONAL win rate (99.66% of sample). PRE-EXPANSION events (small n = 24) hit 88.24% at 24h — directional but not load-bearing.

What this paper does

We document the full validation pipeline end-to-end — how events are detected, what state is captured at detection time, how the realised price path is compared against the expected direction, and how each event is ultimately classified. We then walk through the observed results across four analytical slices: holding-period sensitivity, risk/reward asymmetry, event-type composition, and regime conditioning. The paper closes with a caveats section that states, in plain terms, which findings are statistically load-bearing and which are not.

What it does not do

This is a methodology paper, not a track record. It does not annualise returns, assume position sizing, or claim a live trading P&L. It documents the predictive content of raw, unweighted signals — the input layer on which downstream strategies operate.

01

Framework and scope

Definition of a signal

A Quorum Index signal is a scanner-detected event in which one or more real-time indicators cross a calibrated threshold on a monitored instrument. Each event is fully state-stamped at the moment of detection: regime label, volatility percentile, and the count of co-triggered indicators are all written to the event log before any forward price path is observed. This is the ex-ante record, and it is what subsequent validation reads against.

Four event types

The scanner emits four distinct event types, each with its own trigger logic:

Event type	Trigger	Sample (45d)
score_spike	Aggregate score crosses its rolling upper band.	7,275
score_cluster	Multiple correlated components exceed threshold together.	2,712
volatility_shift	Volatility percentile moves across a calibrated band.	89
regime_change	Regime-state classifier transitions label.	8

The two high-frequency types — score_spike and score_cluster — jointly account for 99.0% of sampled events and carry almost all of the statistical weight. volatility_shift and regime_change are rarer by design: their thresholds are set so that they fire only on meaningful structural moves.

Three evaluation windows

Every event is evaluated at three forward horizons — 1 hour, 4 hours, and 24 hours after detection. The three windows exist to separate short-horizon reaction, intraday follow-through, and full-day persistence. A signal that is predictive at 4h but decays by 24h is a different object from one that strengthens with time; the methodology is deliberately built to expose that difference rather than collapse it into a single number.

Scope of this audit

The audit period runs from 27 February 2026 to 14 April 2026 — 45 calendar days, covering 10,084 distinct signals. All events within that window are included; none are dropped for being inconvenient, and no_data outcomes (where a forward price fetch was unavailable) are logged rather than silently excluded. Classification rates therefore reflect only the events on which a verdict was actually reachable.

A note on sector_rotation

A fifth scanner event type, sector_rotation, is tracked in the Quorum Index pattern library as structural context but is deliberately excluded from this directional-validation audit. A sector_rotation event fires when the dominant sector by triggered-pair count changes from one scan window to the next — it is an observation about market leadership, not a prediction about any particular instrument's price. Because no single primary instrument can be honestly assigned as the target of the rotation, no natural directional expectation exists for the event type. Engineering one for the sake of inclusion would either inflate or deflate the headline numbers reported here without adding real signal. The exclusion is therefore methodological rather than cosmetic.

02

Methodology

The validation pipeline has four stages: detect, stamp, evaluate, classify. The diagram below lays them out. Every stage writes to an append-only log — events are not revised after the fact, and the state captured at detection is frozen.

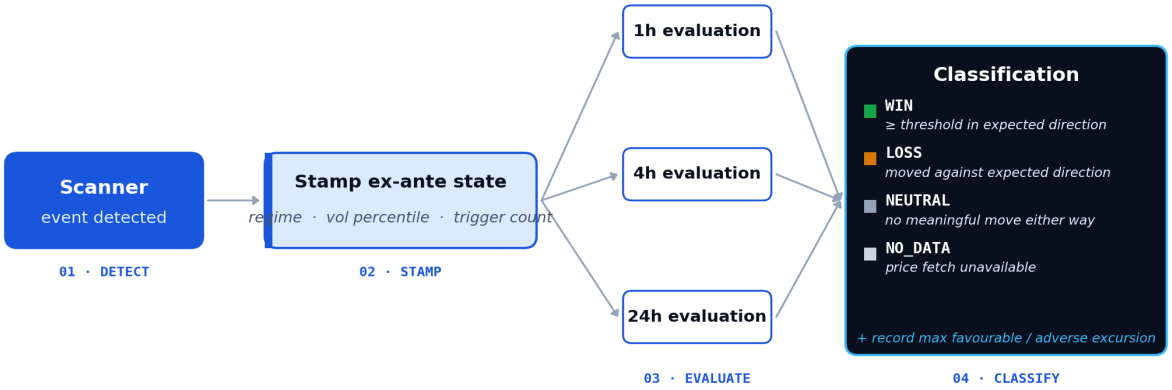

Fig. 1 · End-to-end validation pipeline. State is captured ex-ante at detection; classification happens at three forward windows.

Classification rules

After each evaluation window closes, the instrument's realised price path is fetched and compared against the expected direction recorded at detection. Each event is assigned exactly one of four labels:

Label	Condition
WIN	Instrument moved at least the minimum threshold in the expected direction by window close.
LOSS	Instrument moved against the expected direction by window close.
NEUTRAL	No meaningful move in either direction — change stayed inside the threshold band.
NO_DATA	Forward price fetch unavailable — event is recorded but excluded from win-rate denominators.

02

Methodology · continued

Direction and threshold assignment

Every event carries an expected direction assigned deterministically at detection time, before any forward price path is observed. The direction is a function of event type:

Event type	Expected direction	Threshold
score_spike	up	±0.5%
score_cluster	up	±0.5%
regime_change	variable — derived from target regime (EXPANSION / PRE-EXPANSION → up; COMPRESSION / DISTRIBUTION → down; TRANSITIONAL → neutral)	±0.5%
volatility_shift	magnitude (either direction counts)	2.0% absolute

The *magnitude* rule for volatility_shift is why that event type's win rate is not interpretable as directional edge (see § 05). The *variable* rule for regime_change introduces path-dependence on the regime transition itself, flagged for a future revision as the sample grows beyond its current n = 8. Thresholds are production defaults and have not been re-tuned for this audit; moving them would change win-rate denominators and is disclosed as a sensitivity in § 09.

Risk/reward recording

Directional win-rate alone is a lossy summary. A signal that is right 45% of the time but captures 3× the excursion on its wins versus its losses is, in expectancy terms, a better signal than one that is right 55% of the time with symmetric payoffs. To preserve that information, every event records two additional fields on top of its directional verdict: **max favourable excursion** — the furthest positive move the instrument reached inside the window — and **max adverse excursion**, the furthest negative move. Together these give a window-independent view of the payoff profile, and are averaged across the sample to produce the R/R ratio reported in § 04.

Small-sample discipline

Any subsample with fewer than 30 classified events is flagged explicitly in results and is not treated as load-bearing. This applies in particular to the regime_change event type (n = 8) and the PRE-EXPANSION (n = 24) and COMPRESSION (n = 10) regime slices. Their win rates are reported for completeness but are not evidence, and the paper says so at every point they appear.

03

Directional accuracy

Across the full 10,084-event sample, directional accuracy sits close to 50% at every evaluation window. At 1h, 49.42% of classified events moved in the expected direction. At 4h, 49.92%. At 24h, 47.75%. Raw hit-rate therefore does not, on its own, show an edge — and the paper makes no claim that it does.

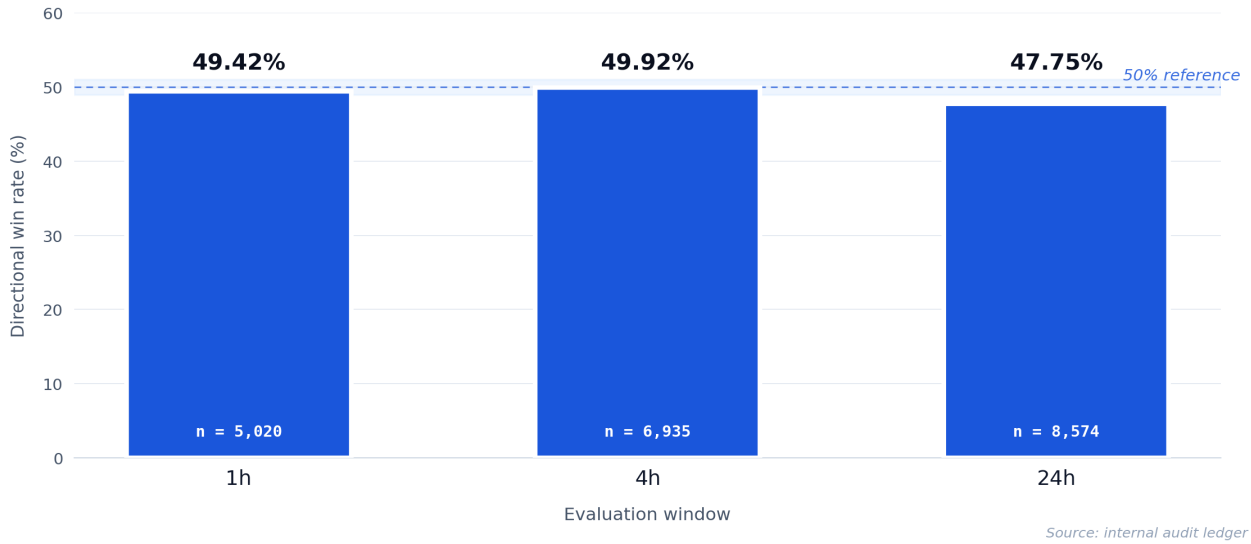

Fig. 2 · Directional win rate by evaluation window. Neutral and no_data events are excluded from the denominator.

Window	Wins	Losses	Classified	Win rate
1h	2,481	2,539	5,020	49.42%
4h	3,462	3,473	6,935	49.92%
24h	4,094	4,480	8,574	47.75%

Reading the numbers honestly

Three observations, stated plainly. First, classification rate rises with window length — 5,020 events are classified at 1h but 8,574 at 24h, as the threshold for a meaningful move is more often crossed given more time. Second, win rate *declines* with window length, from 49.42% at 1h to 47.75% at 24h. Third, neither of those facts is itself good or bad — they are input to the asymmetry analysis that follows, which is where the signal edge actually lives.

INTERPRETATION

A near-50% directional rate is not failure. It is the expected shape of a signal whose value lives in asymmetric payoff rather than raw hit-rate — a profile common to momentum-like and volatility-breakout indicators. The R/R analysis in § 04 is where the edge becomes visible.

04

Risk/reward asymmetry

For each event, we record the furthest favourable and furthest adverse excursion inside the evaluation window. These two fields together give a reading on payoff asymmetry that is independent of where the window happens to close — the directional verdict captures the endpoint, while the excursions capture the full path.

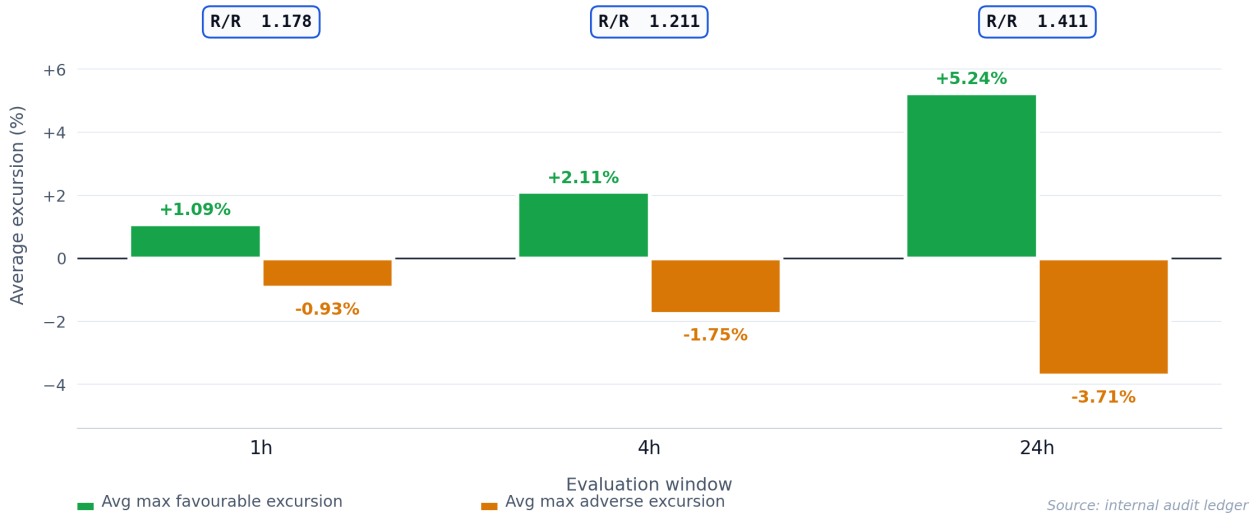

Fig. 3 · Average max favourable versus average max adverse excursion. R/R ratio is the absolute value of the ratio between them.

Window	Avg favourable	Avg adverse	R/R ratio	Sample
1h	+1.09%	−0.93%	1.178	10,084
4h	+2.11%	−1.75%	1.211	10,084
24h	+5.24%	−3.71%	1.411	10,084

The ratio expands with horizon

R/R at 1h is 1.178 — favourable excursions already outweigh adverse by a small margin. By 4h the ratio is 1.211. By 24h it reaches 1.411, meaning the average winning path ran roughly 41% further than the average losing path. The pattern is monotonic across windows and is stable enough across the sample that it is the primary quantitative result of this audit.

Expectancy framing

Combine the 24h directional rate (47.75%) with the 24h R/R ratio (1.411) and the naive expectancy per signal is positive in units of excursion: **0.4775 × 5.24%** – **0.5225 × 3.71%** ≈ **+0.56%** per signal before costs, sizing, or any portfolio overlay. The number is illustrative — it uses symmetric-direction assumptions and no transaction-cost model — but it is sufficient to establish that the raw input layer carries positive expected value before any strategy logic is layered on top.

05

Accuracy by event type

Aggregated win rates hide dispersion across the four event types. Breaking the 24h view out by type exposes a clear difference between the two high-frequency signals and the two rare ones.

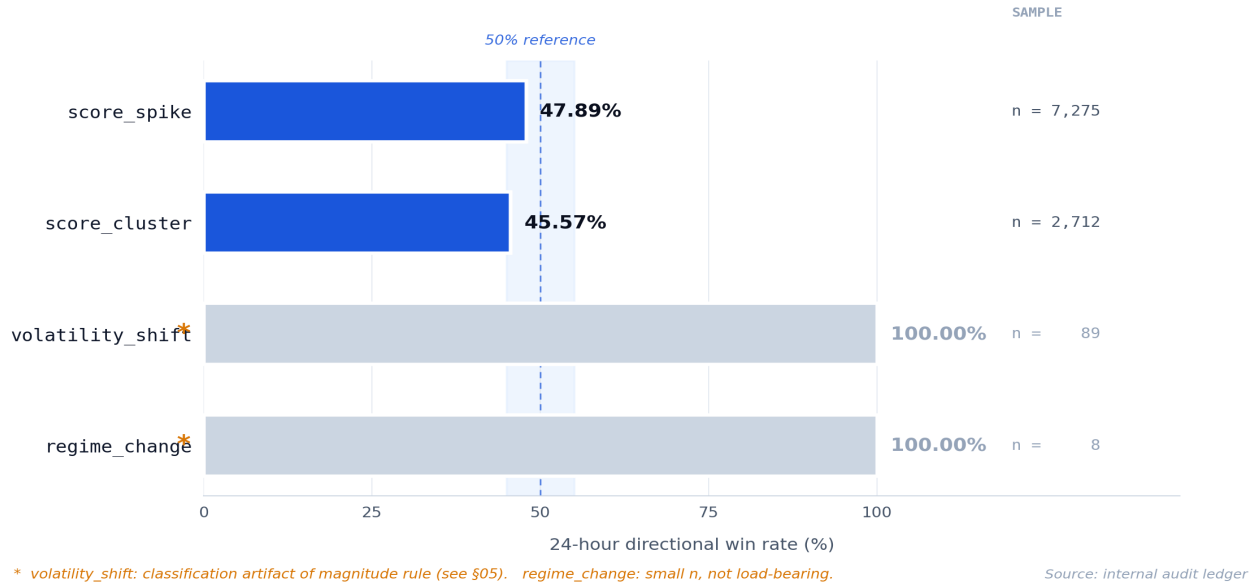

Fig. 4 · 24-hour win rate by event type. Greyed bars mark the two non-load-bearing rows; volatility_shift is a classification artifact, regime_change is small-sample.

Event type	Count	24h win rate	Avg fav 24h	Avg adv 24h
score_spike	7,275	47.89%	+5.06%	−3.61%
score_cluster	2,712	45.57%	+5.73%	−3.99%
volatility_shift	89	100.00%	+4.48%	−3.54%
regime_change *	8	100.00%	+2.74%	−4.82%

* small n · not load-bearing.

Two regimes of signal behaviour

score_spike (n = 7,275) and **score_cluster** (n = 2,712) are the load-bearing types and behave similarly to one another: win rates of 47.89% and 45.57% respectively at 24h, with favourable excursions around +5% and adverse around –3.7%.

volatility_shift (n = 89) shows a 100% win rate, but this figure is not evidence of predictive edge. The event is classified using a magnitude-based rule — a win is logged whenever the absolute excursion reaches 2% in either direction — and the trigger itself fires on large volatility expansions. Trigger condition and win condition are therefore structurally correlated, and the 100% rate is a measurement artifact rather than evidence of directional forecasting skill. The row is reported for completeness only and carries no load-bearing claim anywhere else in this paper.

regime_change (n = 8) is reported for completeness and flagged in every view where it appears. Eight events cannot support a load-bearing claim. This row is not cited as evidence in any summary statement elsewhere in the paper.

06

Accuracy conditioned on regime

The regime state stamped at detection time is the most useful conditioning variable we record. When events are partitioned by the regime they fired into, win rates stop looking uniform — even though the sample is dominated by a single regime.

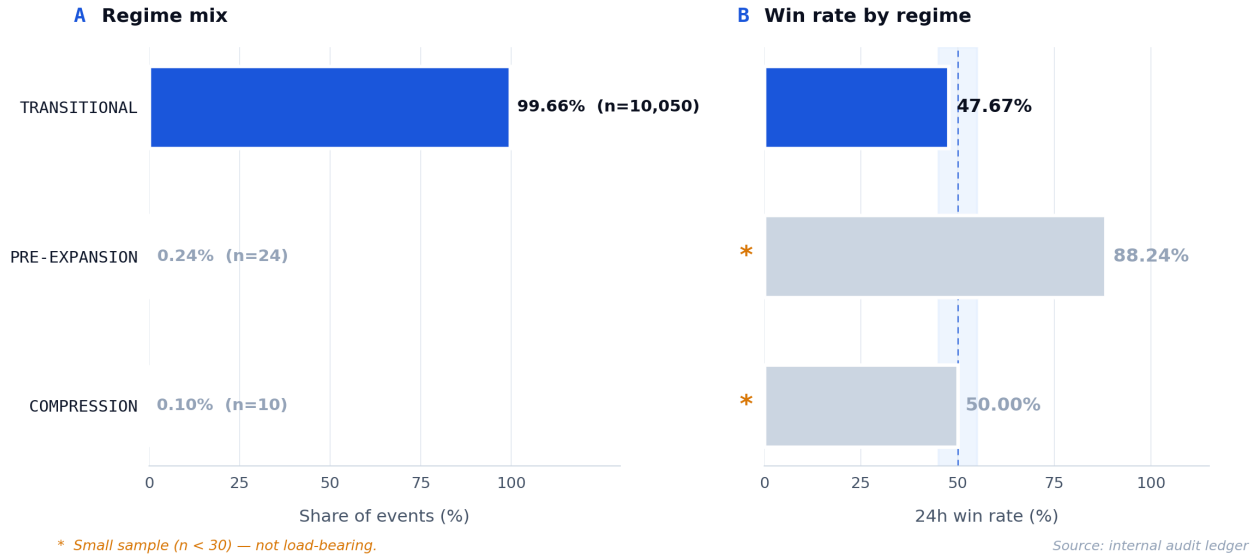

Fig. 5 · Left: event-count distribution across regime states. Right: 24h win rate within each state. 99.66% of sample fell inside TRANSITIONAL.

Regime at event	Count	24h win rate	Share of sample
TRANSITIONAL	10,050	47.67%	99.66%
PRE-EXPANSION *	24	88.24%	0.24%
COMPRESSION *	10	50.00%	0.10%

* small n · not load-bearing.

What the distribution tells us

The 45-day window was effectively a single regime: 10,050 of 10,084 events (99.66%) fired while the regime classifier was in TRANSITIONAL. The 47.67% win rate inside that state is therefore a near-perfect proxy for the headline number reported in § 03, and no meaningful regime mix is available for comparison within this window.

Why we still break it out

Two reasons. First, the PRE-EXPANSION slice (n = 24) posts an 88.24% win rate — well above the 30-event floor but still a small-sample observation. The direction is worth flagging because it is consistent with the design hypothesis that early-expansion regime stamps carry materially more directional information than transitional ones. Whether that holds in a broader sample is an open question this 45-day window cannot answer. Second, the reporting framework itself matters: by writing the breakout explicitly even when the non-dominant regimes are nearly empty, we commit to presenting the slice unchanged in future windows where the mix will differ.

CAVEAT

Regime-conditional win rates outside TRANSITIONAL are descriptive only and not load-bearing. Any future whitepaper iteration will report them with confidence intervals once sample sizes support it.

07

Recent versus all-time behaviour

Sliding the comparison window forward from all-time to 30-day to 7-day lets us look for regime-specific drift. A large gap between recent and all-time accuracy would be evidence that the current market state is materially different from the audit baseline.

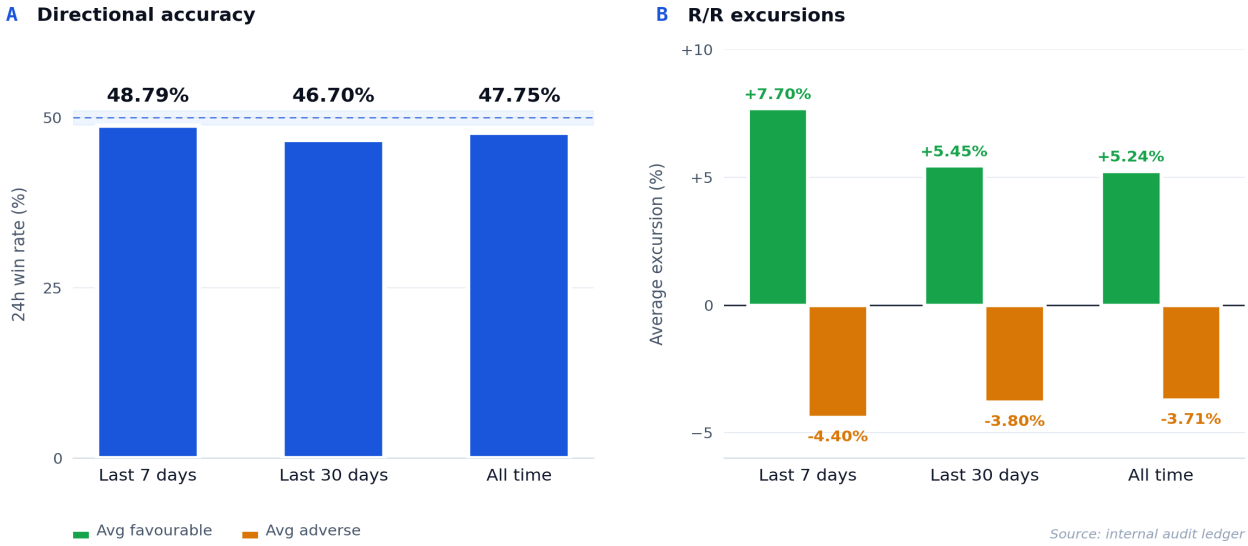

Fig. 6 · Directional win rate and R/R excursions across recent versus all-time windows.

Window	Events	24h win rate	Avg favourable	Avg adverse
Last 7 days	557	48.79%	+7.70%	−4.40%
Last 30 days	5,465	46.70%	+5.45%	−3.80%
All time	10,084	47.75%	+5.24%	−3.71%

Directional accuracy is stable

The 7-day (48.79%), 30-day (46.70%), and all-time (47.75%) win rates sit inside a ~2-point band. None of those differences is large relative to the sample sizes involved, and none of them amounts to structural drift in the accuracy of the underlying signal layer.

Excursions have widened

The R/R picture is less flat. Last 7 days: +7.70% / –4.40%, an implied R/R of 1.75. Last 30 days: +5.45% / –3.80%, ratio 1.43. All time: +5.24% / –3.71%, ratio 1.41. Both favourable and adverse excursions are larger in the most recent week — consistent with a higher-volatility backdrop expanding the payoff profile on both sides. The R/R ratio itself has widened modestly, not collapsed, so the expansion is not penalising the signal's expected value.

08

Quarterly performance

The same statistics aggregated by quarter give a cleaner read on stability across longer windows. Q1 2026 is a full quarter; Q2 2026 is partial, covering only the 14 days between 01 April and the 14 April cut-off.

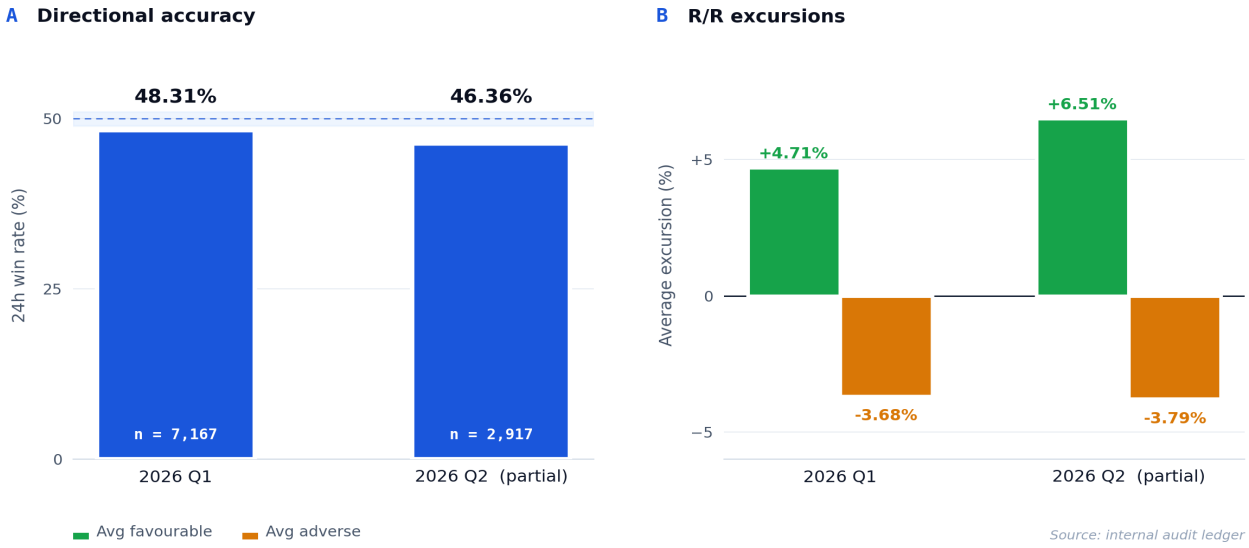

Fig. 7 · Quarterly rollup. Q2 2026 is a partial period through 14 April.

Quarter	Events	24h win rate	Avg favourable	Avg adverse
2026 Q1	7,167	48.31%	+4.71%	−3.68%
2026 Q2 (partial)	2,917	46.36%	+6.51%	−3.79%

Reading the quarterly picture

Directional accuracy slipped roughly 2 points from Q1 (48.31%) into partial Q2 (46.36%). Favourable excursions expanded by 180 basis points over the same transition (+4.71% → +6.51%) while adverse excursions barely moved (–3.68% → –3.79%). Net: slightly lower hit-rate, materially better asymmetry. The Q2 figure covers only 14 days and should be read as provisional — the paper will be refreshed at the end of the quarter with a full Q2 rollup.

09

Caveats and limitations

What this audit proves

Three things, at load-bearing sample sizes. First, that the raw signal layer is directionally near-random at 24h (47.75%). Second, that it carries a positive R/R ratio that expands monotonically with horizon and reaches 1.411 at 24h. Third, that these properties are reproducible from an append-only log without look-ahead.

What it does not prove

It does not prove a live trading result. It does not prove regime-conditional predictive power beyond TRANSITIONAL (the only regime with sample support). The volatility_shift row is a measurement artifact of magnitude-based classification (see § 05) and carries no forward-looking claim. The regime_change row (n = 8) is reported for completeness and should be treated as zero evidence. Finally, it does not prove behaviour in a sustained COMPRESSION or PRE-EXPANSION regime, because the audit window did not include one.

Known limitations

Window length. 45 days is short by any standard. The results reported here will be refreshed quarterly; expect the R/R figures in particular to move as the sample grows.

Regime imbalance. 99.66% of events fired in one regime. A 12-month audit is the minimum horizon at which a meaningful across-regime comparison will be available.

Threshold sensitivity. The minimum thresholds in this audit are 0.5% for score_spike and score_cluster, 2.0% for volatility_shift, and 0.5% for regime_change. These are production defaults and have not been re-tuned for reporting. Moving them would change the classified-event counts and therefore the reported win rates.

Survivorship and data availability. NO_DATA outcomes are logged and excluded from win-rate denominators but included in excursion totals where partial price paths exist. No event is dropped silently.

> **REPORTING COMMITMENT**
>
> Event detection, state-stamping at detection time, forward price fetching, outcome classification, and every statistic in this paper are produced by deterministic Python code operating on raw scanner and exchange data. No large-language-model output is in the causal path of any number reported here. Language-model components (Claude Sonnet 4.6 and variants) are used within the Quorum Index publishing layer for narrative framing of market commentary, and within this whitepaper's own drafting for prose composition — but not for signal classification, outcome evaluation, or statistical aggregation. Every figure in this paper is reproducible directly from the underlying event log. The next refresh is scheduled for end of 2026 Q2.

A

Appendix · Reference tables

A.1 Directional accuracy (all time)

Window	Wins	Losses	Classified	Win rate
1h	2,481	2,539	5,020	49.42%
4h	3,462	3,473	6,935	49.92%
24h	4,094	4,480	8,574	47.75%

A.2 Risk/reward asymmetry (all time)

Window	Avg favourable	Avg adverse	R/R ratio	Sample
1h	+1.09%	−0.93%	1.178	10,084
4h	+2.11%	−1.75%	1.211	10,084
24h	+5.24%	−3.71%	1.411	10,084

A.3 Accuracy by event type

Event type	Count	24h win rate	Avg fav 24h	Avg adv 24h
score_spike	7,275	47.89%	+5.06%	−3.61%
score_cluster	2,712	45.57%	+5.73%	−3.99%
volatility_shift †	89	100.00%	+4.48%	−3.54%
regime_change *	8	100.00%	+2.74%	−4.82%

A.4 Accuracy by regime at event time

Regime at event	Count	24h win rate
TRANSITIONAL	10,050	47.67%
PRE-EXPANSION *	24	88.24%
COMPRESSION *	10	50.00%

A.5 Quarterly rollup

Quarter	Events	24h win rate	Avg favourable	Avg adverse
2026 Q1	7,167	48.31%	+4.71%	−3.68%
2026 Q2 (partial)	2,917	46.36%	+6.51%	−3.79%

* small n · not load-bearing. † classification artifact of magnitude rule; not directional edge.

Legal disclaimer

This document is provided for informational and research purposes only. It is not investment advice, does not constitute an offer to buy or sell any security or financial instrument, and should not be relied upon as the basis for any investment decision. Quorum Index publishes market-structure observations, not trading recommendations. Past observed signal behaviour does not guarantee future results. No fiduciary relationship arises from the distribution or receipt of this document.

Generated from the Quorum Index Internal Audit Ledger dated 18 April 2026 (coverage: 2026-02-27 to 2026-04-14, 10,084 events). All statistics reproducible from the source event log. Rows marked (small n) are sub-threshold and reported for completeness only.