COMPUTATIONAL FINANCE / 20 NOVEMBER 2026

Forecast controls and market evidence

A measurement addendum to PhaseWeave

Summary. Every width-200 forecaster improved on persistence in the October quote-margin test. Replacing a classical feature map with coherent qutrit simulation produced mixed results across seeds and backbones. An independent EUR/USD regime model improved on a constant probability, while a fixed ordinal rule scored better still. The remaining measurements establish computation coverage, depth-response counts and macro-release availability.

1. Matched interventions on the forecaster

From 16 consecutive FX graph frames, the model predicts the next-second quote margin for both triangle orientations. Mean absolute error (MAE) averages the two directions in log basis points. Each experimental cell fixes the backbone, seed, evaluation rows, starting parameter values, and training order while changing the patch algebra. This intervention supports a comparison of forecasting algorithms. It does not identify a causal effect on market prices.

The seven backbones are GCN, GAT, MPNN, TokenGT and the three local-to-TokenGT hybrids. Softmax or tropical attention is crossed with four patches: classical cosine, coherent complex propagation, dephased probability propagation, and max-plus propagation. All four use the same six directed rates and the same output dimension. Tropical aggregation is a separate choice inside the patch and the attention blocks. A local GAT retains softmax over neighbors. A conventional GPU simulated the qutrit circuit.

	Width 32 pilot	Width 200 registered
Fits; seeds	56; one	112; two
Unique train rows; rows per evaluation split	512; 256	8,192; 2,048
Passes; optimizer updates per fit	2; 64	3; 384
October persistence MAE	0.198725	0.206430
October model MAE range	0.199297–0.199622	0.176590–0.205274
Fits below persistence	0/56	112/112

The two cohorts change width, sample size, training exposure and seed roster together. The comparison across columns cannot isolate a width effect. At width 200, coherent propagation improved on its corresponding cosine model in 7/14 settings for seed 20261225 and 6/14 for seed 20261226. The individual 95% paired UTC-day bootstrap intervals lay entirely below zero in 2/14 and 6/14 settings, respectively. Both seeds reuse the same observations. The intervals have no multiple-comparison correction. All 112 fits selected epoch three by validation MAE. Each run saved its selected checkpoint and recorded all three epoch scores.

Evidence. Complete matched comparisons and per-seed scores. Immutable gated width-200 milestone, including fit receipts.

MODEL INSPECTION AND REGIME DYNAMICS

2. What the execution traces establish

Each of 231 validation-selected checkpoints processed the same released, 16-frame EUR/USD/JPY example. Across widths 32, 200, 448 and 832, the index records 8,734 operations and 42 checkpoint-specific coherent-circuit ledgers. The indexed runs include the 168 matched fits on page one. The wider roster also includes earlier sweeps, convergence runs and multi-horizon forecasters. Multiple traces of this single input expand operator coverage without adding independent market examples.

For each captured contraction or max-plus reduction, a public dossier records a selected output coordinate, a reconstruction check and up to 16 active terms. Embedding lookups record the selected entry. Each coherent ledger expands the four-qutrit circuit into 15 operations and records all $3^4 = 81$ complex amplitudes before and after each. Four preparation operators realize the native product state using an explicit unitary completion. Seven commuting diagonal factors reproduce the fused phase. Four trained-angle local mixers precede Born probabilities and four real-clock expectations. These checks expose arithmetic and gate order. A forecasting comparison still requires the held-out errors on page one. Operator workbench and indexed dossiers.

3. Does an active EUR/USD minute arrive before a quiet one?

HistData EUR/USD tick quotes give a one-minute range $10^4 \log(m_{\max}/m_{\min})$ of the observed bid–ask midpoint. An eligible minute contains at least two distinct quote timestamps. The August range quintiles define five fixed states. Let A be the lowest and B the highest. The committor q_i is the model's probability of reaching B before A. For the August one-minute transition matrix P, solve

$$q_A = 0, \qquad q_B = 1, \qquad q_i = \sum_j P_{ij} q_j \quad (i \notin \{A, B\}).$$

The 29,346 August transitions yield interior probabilities 0.402, 0.506 and 0.600. Paths stop at missing minutes, UTC midnight and split boundaries. Resolved interior starts supply first-hit labels. Starts whose paths end before either endpoint are censored.

Brier loss (lower is better)	Fitted	Reversible	Fixed rank	August constant
September 1–16	0.232763	0.232694	0.232615	0.249918
September 17–30	0.226223	0.225461	0.214272	0.247804
October	0.225649	0.223481	0.214132	0.244642

October contributed 15,198 resolved interior starts and 135 censored starts. The fitted-minus-constant Brier difference was −0.01899 (paired UTC-day bootstrap 95% interval $[-0.02062, -0.01746]$). Fixed rank assigns 0.25, 0.50 and 0.75 to interior states. Its lower error limits the case for the fitted model. The two-minute transition discrepancy rose from 0.10379 in August to 0.20519 in October, consistent with imperfect temporal aggregation or distribution shift.

Evidence and method. Frozen counts, scores and interval. Metzner, Schütte and Vanden-Eijnden (2009), transition path theory. Prinz et al. (2011), kinetic model validation.

MARKET MICROSTRUCTURE AND RELEASE TIMING

4. Depth imbalance on one BTC perpetual recording

A 41.427-second Gate BTC_USDT perpetual recording contains 18 ten-level books and 17 eligible subsequent responses. The midpoint comes from the book. The mark arrives on a separately clocked public ticker. An origin uses only a mark already received locally whose event and message times also pass the book's exchange-time checks. A later outcome requires a new eligible observation. Gaps between snapshots range from 2.111 to 3.494 seconds; intervening order events are unobserved.

For side quantities b_r, a_r at displayed rank r, weighted imbalance equals $(\sum_{r=1}^{10}(b_r - a_r)/r)/(\sum_{r=1}^{10}(b_r + a_r)/r)$. Its sign matched the next midpoint movement on 10/17 transitions. The midpoint fell 10 times, so the always-down control matched exactly as many. Under equal weighting of all ten levels, accuracy was 8/17. Against subsequent marks the inverse-rank signal matched 10/17 and always down matched 12/17. The last book has no future label. These are outcomes on one overlapping path. A next-snapshot response need not be the next tick or an attainable fill. Depth receipt; Gould and Bonart, queue imbalance study.

5. Original-vintage macro releases at decision time

The BLS event graph links one publisher, six releases and three indicators with 15 edges. The table links the original U.S. Department of Labor PDFs. Every event has an embargo of 12:30:00 UTC on its listed date and becomes eligible strictly *after* 12:31:00 UTC under the modeled 60-second allowance. Eligibility expires at 24 hours after embargo. Historical subscriber receipt times are unknown.

2026 release date	Reference month	Release	Change	Unemployment
Aug 7	July	Employment Situation	−23,000 jobs	4.1%
Aug 12	July	CPI-U	+0.1%	—
Sep 4	August	Employment Situation	+162,000 jobs	4.1%
Sep 11	August	CPI-U	+0.4%	—
Oct 2	September	Employment Situation	+29,000 jobs	4.2%
Oct 14	September	CPI-U	+0.6%	—

Changes are the original seasonally adjusted monthly CPI-U percentage change or the nonfarm payroll change. Within the eligible 24-hour window, the frozen join covers 6,595 August training anchors, 8,956 September validation anchors, zero September stress anchors, 7,307 October test anchors and 3,348 held-out CHF anchors. The last two folds share the same two October releases. Only 15 October test anchors and seven CHF anchors qualified within five minutes. These are availability counts. A test of economic response would need additional independently held-out releases and grouping by release date across markets.

Evidence. Immutable release graph, PDF checksums and clock definitions. Public aggregate graph and split coverage.