Forecasting longer exchange-rate cycles

Extended technical report · Five forecast horizons on four- and five-currency graphs

Result. At the one-second target, all 22 graph-model fits improved on persistence on the original four-currency October test for both triangles and quadrilaterals. Every one of those fits fell behind persistence on CHF-containing triangles, quadrilaterals and pentagons. A linear ridge correction gave the lowest error on the longer CHF routes. The comparison tests numerical forecasts of recorded indicative quotes.

The source universe contains EUR, GBP, JPY and USD, augmented with CHF on the five-vertex graph. Models fit August K4 sessions and selected checkpoints on separate September dates. Both October evaluations use the same six session dates. Each graph uses direct bid and inverse-ask legs. Its target is the sum of log conversion rates at the future tenor, minus 0.1 assumed log basis point per leg.

What changes when CHF enters the graph?

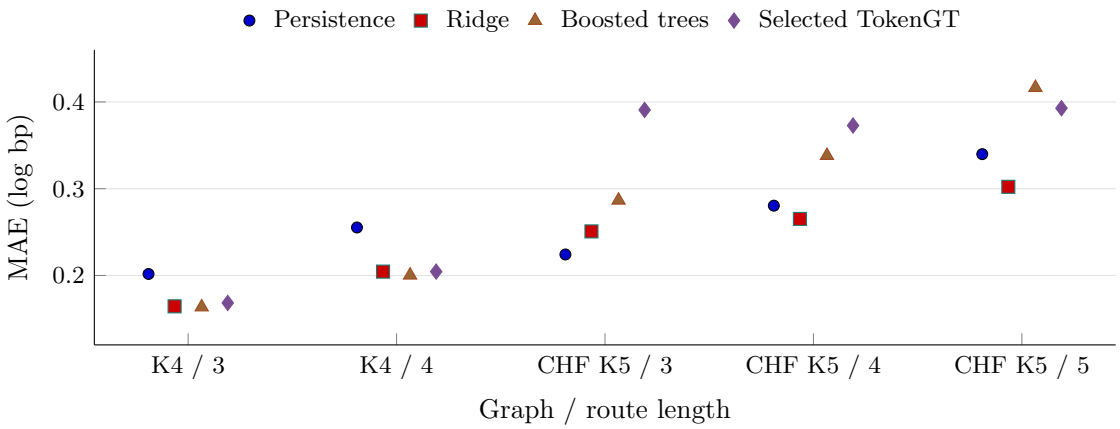


Figure 1. One-second target, one-second input cadence. Lower values indicate smaller mean absolute error. The displayed neural fit had the smallest September validation macro MAE among the 22 dense-history fits. It used classical route features and tropical TokenGT attention.

Graph	Legs	Labels	Persistence	Ridge	Boosted trees	TokenGT
K4	3	8,136	0.2017	0.1644	0.1636	0.1684
K4	4	4,830	0.2553	0.2045	0.2003	0.2045
CHF K5	3	9,880	0.2242	0.2508	0.2866	0.3908
CHF K5	4	13,584	0.2805	0.2652	0.3381	0.3728
CHF K5	5	9,810	0.3398	0.3022	0.4164	0.3928

Each label is a directed route, origin and tenor. A single origin can contribute several labels. The best two-seed neural mean on K4 quadrilaterals was 0.203728 for GAT with token attention. That same pair averaged 0.413984 on CHF pentagons. Boosted trees had the lowest one-second MAE across all 25 entries on both K4 lengths. The lowest K5 values shift to persistence on triangles and ridge on lengths four and five. The new graph changes currency, graph size and route support together. Five-leg routes were absent during training.

What the paired controls isolate

Same directed route, UTC origin and target horizon

The full history contains 16 one-second frames. The decimated history retains the four frames at $-15, -10, -5, 0$ seconds. Both end at the same origin and cover the same elapsed time. Holding the eligible keys fixed separately from the cadence change left 40,960 October K4 and 109,640 CHF K5 route–origin–tenor labels across the five target horizons. An additional sensitivity restricts origins to the five-second grid.

Circuit readout on the same message-passing backbone

The coherent arm evolves a four-qutrit, 81-amplitude state on a CUDA simulator. A phase depending on basis digits and adjacent-wire couplings precedes four local unitary mixers. Squared output amplitudes give probabilities for four expectation values. The dephased arm propagates probabilities using squared mixer entries. The max-plus arm accumulates maximum log-path scores and normalizes the terminal scores. The tropical attention control changes the TokenGT attention kernel separately. All measurements here come from software.

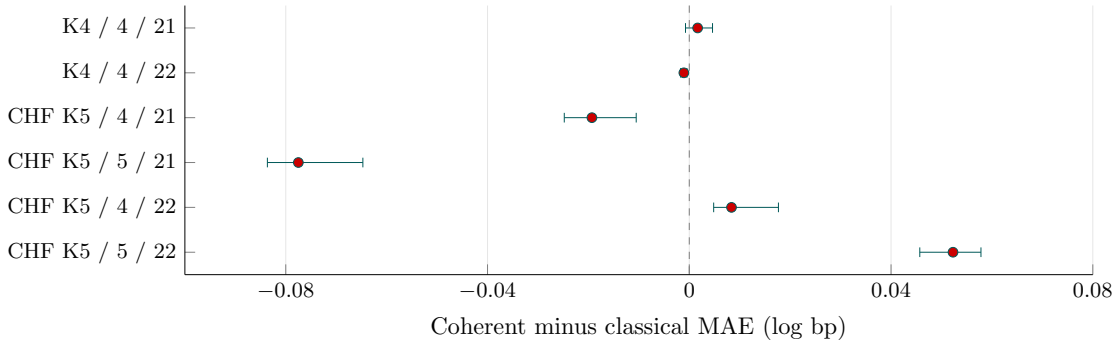


Figure 2. One-second cadence and target, common-mask labels. Rows identify graph, route length and the final two digits of the training seed. Each interval resamples six whole October dates 2,500 times. Negative differences favor coherent propagation.

The one-second CHF pentagon effect changes sign across seeds. The observation-weighted difference is -0.0775 log bp for seed 20261121 and $+0.0523$ for 20261122. For comparison, using five-second inputs gives differences of $+0.0166$ and $+0.0361$ on the same 6,944 pentagon labels. The fixed seed and cadence affect the relative result. On the native one-second CHF pentagon cohort both coherent fits still have higher error than persistence and ridge.

Interpretation

The target contains modeled quote availability and a fixed leg charge. The study has no observed executable size, queue priority, order arrivals or realized fills. Overlapping windows and routes share observations across all five horizons. The day intervals describe these six dates and the specified individual contrasts. The experiment does not establish a revenue estimate, a quantum hardware advantage or the mechanism of the transfer failure.

Construction of the supervised problem

Calendar boundaries, quote availability and outcome support

The frozen archive plan covers 30 HistData archive entries. The base graph contains the six direct pairs among EUR, GBP, JPY and USD. Adding CHF gives ten direct pairs. Each pair contributes two directions. A route starts at its lexicographically smallest currency, with both orientations retained. K4 has eight triangles and six four-leg routes. The CHF-only target filter on K5 leaves 12 directed triangles, 24 four-leg routes and 24 five-leg routes. Other K5 pairs remain available as input.

Fold	Session dates in 2026	Origins	4-leg labels, 1 s	5-s history subset
Train	Aug 4, 7, 12, 14, 18, 21, 25, 27	5,150	5,756	3,612
Validation	Sep 2, 4, 9, 11, 15	2,560	3,138	1,964
Stress	Sep 17, 18, 22, 24, 29	2,560	4,514	3,454
K4 test	Oct 6, 9, 14, 20, 23, 30	3,072	4,830	3,456
CHF K5	Oct 6, 9, 14, 20, 23, 30	3,072	13,584	10,018

Counts identify distinct units. Origins count graph windows, whereas labels also distinguish route and direction. Both right-hand columns use a one-second forecast horizon.

Information available at the origin

Each session is 12:00–13:00 UTC. The curator converts the provider's fixed EST convention, then adds one second of modeled release delay. If several rows have the same source timestamp it uses the minimum bid and maximum ask. A leg is available only when its released quote is no later than the frame and its measured age is at most 250 ms. Whole-second source stamps and query frames make the age of accepted quotes zero at this resolution. This clock cannot reconstruct intrasecond arrivals.

A window must contain a currently complete long route and a prior complete observation of that route. Curation samples qualifying origins by input masks with a fixed split/date seed, at most 768 per training session and 512 otherwise. The full 15-second lookback and future horizon up to 300 seconds fit inside the same session. A target additionally requires all its legs to have refreshed strictly after the origin. Changing the history cadence can change eligibility even on the same stored window.

Forecast target

For a route C of length L, let $r_e(t)$ be its oriented log quote in log bp. At horizon h, the target is $y(C,t,h) = \sum_{e\in C} r_e(t+h) - 0.1L$. The direct side of a pair uses bid, and the opposite side uses reciprocal ask. The target units are dimensionless log basis points. The baseline carries the admissible origin margin forward. The constant charge cancels in the time difference for a fixed route. It supplies no observed leg fee or depth.

Models, selection and audit coverage

Equal elapsed lookback with separately fitted cadence arms

Graph input summaries include edge half-spread, age and validity. Route summaries include average signed circulation and half-spread with inverse hyperbolic sine transforms, mean age, valid fraction, peak spread and length divided by five. Rotating the starting index preserves a route's cyclic ordered-pair average. Reversing a route changes those ordered pairs. The network combines spatial edge-token embeddings with a GRU over timed route summaries, a route patch and a shared horizon-conditioned residual head. An unavailable input route falls back to persistence.

The graph width is 48, with two graph layers and four attention heads where used. AdamW uses learning rate 0.001, zero weight decay, batches of 128 origins, smooth-L_1 loss with threshold 0.25 and gradient norm clipping at one. Each fit took 246 optimizer updates on an RTX 4090. Parameter counts range from 79,768 to 80,056. Training weights balance route-length and tenor strata present in August. Each fit runs six epochs. Only the September equal-length/equal-horizon macro MAE chooses that fit's checkpoint. The smallest validation score among the 22 one-second fits is 0.186430 at epoch 6. The aggregate below preserves every fit individually.

Backbone	Patch	Attention	Seed*	Epoch	Val. macro	K4 / 4	K5 / 4	K5 / 5
GAT	classical	softmax	21	6	0.1989	0.2158	0.4781	0.5814
GAT	classical	softmax	22	6	0.1889	0.2042	0.3866	0.4863
GAT + TOKENGT	classical	softmax	21	6	0.1913	0.2066	0.4290	0.4106
GAT + TOKENGT	classical	softmax	22	4	0.1868	0.2008	0.4161	0.4174
GCN	classical	softmax	21	6	0.1897	0.2034	0.4672	0.5304
GCN	classical	softmax	22	6	0.1881	0.2056	0.4112	0.4269
GCN + TOKENGT	classical	softmax	21	5	0.1888	0.2047	0.4924	0.5319
GCN + TOKENGT	classical	softmax	22	6	0.1877	0.2065	0.4008	0.4205
MPNN	classical	softmax	21	6	0.1897	0.2034	0.4656	0.5282
MPNN	classical	softmax	22	6	0.1881	0.2056	0.4106	0.4264
MPNN	coherent	softmax	21	5	0.1905	0.2042	0.4523	0.4553
MPNN	coherent	softmax	22	6	0.1882	0.2047	0.4221	0.4820
MPNN	dephased	softmax	21	5	0.1915	0.2049	0.4428	0.4437
MPNN	dephased	softmax	22	4	0.1917	0.2037	0.4132	0.4084
MPNN	maxplus	softmax	21	5	0.1908	0.2053	0.4372	0.4563
MPNN	maxplus	softmax	22	6	0.1930	0.2096	0.4396	0.4402
MPNN + TOKENGT	classical	softmax	21	5	0.1886	0.2044	0.4881	0.5275
MPNN + TOKENGT	classical	softmax	22	6	0.1877	0.2065	0.4013	0.4210
TOKENGT	classical	softmax	21	5	0.1890	0.2055	0.4860	0.5263
TOKENGT	classical	softmax	22	6	0.1877	0.2065	0.4043	0.4238
TOKENGT	classical	tropical	21	5	0.1905	0.2067	0.4888	0.5304
TOKENGT	classical	tropical	22	6	0.1864	0.2045	0.3728	0.3928

One-second input cadence. Test columns give one-second native-mask MAE. *Seed suffix 21 or 22 denotes 20261121 or 20261122. Hybrids apply a directed graph block followed by a token block.

The five-second cadence campaign fits classical and coherent MPNNs under both seeds, plus the same three standard controls. The checkpoint audit verified all 26 selected weight files and recorded 16 independent CPU arm–cadence–split replays for those MPNNs on October K4 and CHF K5. Across the audited metrics and interval endpoints, maximum discrepancy was 3.56×10^{-7} log bp against a 3×10^{-6} tolerance.

Absolute errors by forecast horizon

One-second history cadence, native eligibility

Each learned control uses the same frozen August split. Ridge fits one residual model per horizon ($\alpha = 30$). Histogram gradient boosting fits 100 iterations per horizon with 15 leaves, minimum leaf size 96, learning rate 0.06, and $L_2 = 1$. Inputs concatenate the last, mean, standard deviation, and last-minus-first values of the six route summaries and the transformed origin margin. Both methods add the fitted residual to persistence. Their inductive biases differ from the neural head's shared route and horizon representation.

Graph	Legs	Tenor (s)	Labels	Persistence	Ridge	Boosted trees	TokenGT
K4	3	1	8,136	0.2017	0.1644	0.1636	0.1684
K4	3	5	6,476	0.2065	0.1664	0.1699	0.1686
K4	3	15	6,552	0.2126	0.1778	0.1734	0.1773
K4	3	60	6,236	0.2109	0.1813	0.1723	0.1795
K4	3	300	6,064	0.2213	0.1952	0.2100	0.1953
K4	4	1	4,830	0.2553	0.2045	0.2003	0.2045
K4	4	5	3,822	0.2601	0.2040	0.2017	0.2045
K4	4	15	3,824	0.2630	0.2139	0.2074	0.2116
K4	4	60	3,636	0.2665	0.2178	0.2045	0.2119
K4	4	300	3,488	0.2771	0.2309	0.2379	0.2339
CHF K5	3	1	9,880	0.2242	0.2508	0.2866	0.3908
CHF K5	3	5	8,592	0.2345	0.2755	0.3336	0.4156
CHF K5	3	15	8,666	0.2309	0.2992	0.4051	0.4317
CHF K5	3	60	7,980	0.2325	0.3030	0.4413	0.4585
CHF K5	3	300	7,850	0.2481	0.3307	0.5959	0.4993
CHF K5	4	1	13,584	0.2805	0.2652	0.3381	0.3728
CHF K5	4	5	11,702	0.2975	0.2816	0.3571	0.3952
CHF K5	4	15	12,080	0.2809	0.2873	0.4271	0.4042
CHF K5	4	60	10,780	0.2842	0.2983	0.4536	0.4245
CHF K5	4	300	10,628	0.3047	0.3189	0.5673	0.4700
CHF K5	5	1	9,810	0.3398	0.3022	0.4164	0.3928
CHF K5	5	5	8,492	0.3644	0.3186	0.4341	0.4118
CHF K5	5	15	8,806	0.3349	0.3067	0.5288	0.4184
CHF K5	5	60	7,726	0.3426	0.3222	0.5241	0.4310
CHF K5	5	300	7,560	0.3629	0.3467	0.6018	0.4897

Every score uses the labels present for that length, tenor and split. Comparing rows across tenors changes both prediction distance and eligible observations. The complete CSV includes the September stress cohort, both cadences, origin- and route-macro MAE, RMSE and persistence differences.

The CHF ridge correction improves on persistence at all five pentagon tenors. On CHF four-leg routes it improves at one and five seconds, then falls behind at the three longer tenors. On the original K4 test, HGB's four-leg MAE changes from 0.2003 at one second to 0.2379 at 300 seconds. Pooling route lengths would hide both the training support boundary and these baseline reversals.

Paired comparisons on identical target keys

One-second target across the two October cohorts

Let $\Delta_i = |\hat{y}_{A,i} - y_i| - |\hat{y}_{B,i} - y_i|$ for the same directed route, UTC origin and horizon i. Tables give the observation-weighted mean and 95% day-block bootstrap interval, with 2,500 draws and fixed seed 20261121. Days retain all their constituent routes and horizons for the selected stratum. The machine-readable receipt also gives equal-day estimates and the sensitivity to retaining origins on a five-second grid.

Coherent circuit minus classical route patch

Graph	Legs	Cadence	Seed*	Labels	Δ	95% interval
K4	4	1 s	21	3,456	+0.0017	[-0.0007, +0.0046]
K4	4	1 s	22	3,456	-0.0011	[-0.0017, -0.0005]
K4	4	5 s	21	3,456	+0.0005	[-0.0043, +0.0032]
K4	4	5 s	22	3,456	-0.0020	[-0.0023, -0.0014]
CHF K5	4	1 s	21	10,018	-0.0193	[-0.0248, -0.0105]
CHF K5	5	1 s	21	6,944	-0.0775	[-0.0836, -0.0647]
CHF K5	4	1 s	22	10,018	+0.0084	[+0.0048, +0.0177]
CHF K5	5	1 s	22	6,944	+0.0523	[+0.0457, +0.0578]
CHF K5	4	5 s	21	10,018	+0.0458	[+0.0361, +0.0602]
CHF K5	5	5 s	21	6,944	+0.0166	[+0.0121, +0.0314]
CHF K5	4	5 s	22	10,018	+0.0051	[+0.0040, +0.0060]
CHF K5	5	5 s	22	6,944	+0.0360	[+0.0309, +0.0398]

Five-second history minus one-second history

Graph	Legs	Patch	Seed*	Labels	Δ	95% interval
K4	4	Classical	21	3,456	+0.0027	[-0.0007, +0.0097]
K4	4	Classical	22	3,456	+0.0025	[+0.0003, +0.0053]
K4	4	Coherent	21	3,456	+0.0016	[-0.0015, +0.0043]
K4	4	Coherent	22	3,456	+0.0016	[-0.0013, +0.0053]
CHF K5	4	Classical	21	10,018	-0.0403	[-0.0716, -0.0288]
CHF K5	5	Classical	21	6,944	-0.0944	[-0.1164, -0.0859]
CHF K5	4	Classical	22	10,018	-0.0054	[-0.0210, -0.0001]
CHF K5	5	Classical	22	6,944	-0.0176	[-0.0282, -0.0142]
CHF K5	4	Coherent	21	10,018	+0.0248	[-0.0003, +0.0356]
CHF K5	5	Coherent	21	6,944	-0.0004	[-0.0204, +0.0052]
CHF K5	4	Coherent	22	10,018	-0.0086	[-0.0344, -0.0000]
CHF K5	5	Coherent	22	6,944	-0.0339	[-0.0526, -0.0276]

*Seed suffix. Negative values favor the left arm. Zero-adjacent endpoints in these tables are rounded. The CSV retains full precision for all five tenors and all 120 paired strata.

The shared CHF five-leg cohort at one second contains 6,944 labels, compared with 9,810 on the native one-second history mask. A comparison that subtracts the two native-cadence headline MAEs changes both the predictors and the evaluated observations. Individual paired intervals have six October date blocks and do not provide a multiplicity-adjusted decision over the arm roster.

Cycle algebra and reproducibility

Finite checks alongside the forecasting experiment

A directed quote identity

For four distinct vertices, write $L_{ijk} = r_{ij} + r_{jk} + r_{ki}$ at one common complete timestamp and $S_{ij} = r_{ij} + r_{ji}$. Then

$$L_{123} - L_{012} - L_{023} + L_{013} = S_{13} - S_{02}.$$

The right side vanishes on antisymmetric midpoints. Directed bid/inverse-ask legs generally retain spread on each two-edge round trip. Equal fixed per-leg charges cancel in this four-triangle combination. The relation is an arithmetic consistency check on contemporaneous observations.

An exact ordered-transport companion

Take the unit quaternion group inside SU(2). Set the three star-tree links from 0 to identity and $U_{12} = i$, $U_{23} = j$, $U_{13} = +k$ or $-k$. Reverse links by inversion, and multiply new traversals on the left. Each connection yields zero normalized trace along the three independent star triangles. On triangle 123, $U_{13}^{-1}U_{23}U_{12}$ has normalized trace -1 for $+k$ and $+1$ for $-k$. An exact check exercised all 8^4 vertex gauges in both fixtures. A second implementation used 256 continuous vertex-gauge tuples with maximum discrepancy below 6.8×10^{-16}. The scalar market models and the four-qutrit simulator use their respective mechanisms above. No Wilson witness enters the training targets. Interactive K4 exhibit.

Reproduction

The archive acquisition and training instructions in `REPRODUCE.md` require authorized local quote archives. To rebuild the tables and make the same figures from the frozen aggregate receipts:

```
python3 build_tables.py --campaign "$CAMPAIGN" --audit "$AUDIT" --out .
latexmk -pdf transfer-figure.tex paired-figure.tex
latexmk -pdf compact.tex extended.tex
```

The source bundle includes the generated tables, the two plot-data files, editable figure sources, and the CSV results. Rendering from those generated aggregates needs only a TeX distribution containing PGFPlots. Recomputing neural weights or labels requires the original inputs and pinned project code.

Evidence and boundaries

The source bundle lists SHA-256 identities for the data manifest, runner, configuration, each fit-selection receipt and independent audit. The dataset covers sampled one-hour sessions and the matched experiment uses two training seeds. CHF evaluation shares October dates and several source pairs with the K4 test. Neither its asset change nor its graph-size change is separately identified. Performance at a different market, timestamp convention, sampling policy or trading cost would require new evidence.