FOREIGN EXCHANGE / MATCHED SUPERVISION

Supervising five-leg foreign-exchange cycles

Can five-leg training outcomes improve a route predictor that already sees every five-currency quote graph? The matched experiment holds the graph, input histories, route head, initialization, batch order and validation criterion fixed. One fit learns from three- and four-leg routes. Its paired fit also learns from pentagons. This tests transfer across route lengths on a fixed asset universe.

Observed graph and target

Each graph contains CHF, EUR, GBP, JPY and USD. Ten quoted pairs supply 20 directed conversion rates. Choosing one starting asset per cyclic ordering gives $\binom{5}{k}(k-1)!$ directed cycles of length k, hence 20 triangles, 30 quadrilaterals and 24 pentagons. Both traversal directions are retained. All pentagons visit CHF.

For an oriented cycle c, the indicative route margin is

$$M_c(t) = 10^4 \sum_{i \to j \in c} \log r_{ij}(t) - 0.1\,|c| \quad \text{log bp.}$$

The canonical direction uses the bid and its reverse uses the reciprocal ask. Duplicate source ticks use the lowest bid and highest ask. Whole-second stamps are converted from fixed EST to UTC and shifted one second for modeled availability. A leg must be at most 250 ms old on that clock. Each future leg must also update strictly after the forecast origin. The five forecast tenors are 1, 5, 15, 60 and 300 seconds. The longest tenor stays inside its source session. These are indicative spread- and fee-adjusted quotes with an assumed 0.1 log bp per leg. No traded size or fill is observed.

Chronology and eligible denominators

Every session spans 12:00–13:00 UTC in 2026. Sampling uses only input availability, with deterministic date-specific caps and seed 20261124. Missing future observations have false masks. A single origin contributes multiple route labels.

	Sessions	Origins	Pentagon labels		Common history
Role			1 s tenor	5 s tenor	1 s tenor
August training	12	8,320	26,306	22,154	16,486
September validation	6	3,072	10,882	8,660	6,490
Late September test	5	2,560	8,920	7,026	5,432

Training ends on August 27. Validation uses September 1, 3, 8, 10, 14 and 16. Test uses September 21, 23, 25, 28 and 30. Prior project evaluations included non-CHF triangle outcomes on these September dates. The provenance audit found no previous CHF or pentagon outcomes on the selected September dates within the inventoried datasets. All five assets already occur in training.

Selection

The 1 s input history has 16 frames over 15 elapsed seconds. The 5 s spacing retains frames 0, 5, 10 and 15 at the same origins. Graph fits use width 48, three spatial layers, batch size 96, eight epochs and seeds 20261121 and 20261122. Select the first epoch minimizing validation MAE averaged across lengths three and four and the five tenors. Adding pentagons changes both label support and relative length weights. DLinear uses its separately registered all-length validation rule.

HELD-OUT EVALUATION

Later September pentagon forecasts

Table 1. Mean absolute error in log bp; smaller is better. Every column uses its forecast tenor's intersection of the two history masks. The shared sample has 5,432 route/origin labels at 1 s and 4,176 at 5 s. Stochastic rows average the two per-seed MAEs. Each method predicts the same observed routes at a given tenor.

	1 s forecast tenor		5 s forecast tenor	
Input frame spacing	1 s	5 s	1 s	5 s
Persistence	0.3129	0.3129	0.3247	0.3247
Ridge	0.2549	0.2664	0.2624	0.2662
Histogram gradient boosting	0.2424	0.2530	0.2638	0.2677
Hodge ridge	0.2507	0.2631	0.2621	0.2660
DLinear	0.2966	0.2903	0.3075	0.3001
Classical MPNN, lengths 3–4	0.2436	0.2551	0.2510	0.2614
Classical MPNN, lengths 3–5	0.2469	0.2488	0.2548	0.2543
Coherent MPNN, lengths 3–4	0.2442	0.2562	0.2529	0.2610
Coherent MPNN, lengths 3–5	0.2499	0.2551	0.2574	0.2607

The Hodge ridge control measures cycle-space residual energy of the antisymmetric edge field on the observed graph. Persistence carries the current margin forward. All classical and coherent MPNN rows use the matched six-wire, radix-three specification.

Effect of adding five-leg supervision

Table 2. Full minus short supervision at the 1 s forecast tenor. Negative values favor full supervision. Entries are the paired mean differences in absolute error with their 95% day-block bootstrap intervals. Each seed pair shares eligible route/origin keys.

Patch	Seed	1 s history spacing	5 s history spacing
Classical	20261121	$+0.0033$ $[+0.00095, +0.00532]$	-0.0118 $[-0.01427, -0.00983]$
Classical	20261122	$+0.0031$ $[-0.00004, +0.00680]$	-0.0007 $[-0.00551, +0.00996]$
Coherent	20261121	$+0.0057$ $[+0.00349, +0.00812]$	$+0.0007$ $[-0.00068, +0.00289]$
Coherent	20261122	$+0.0043$ $[-0.00195, +0.01294]$	-0.0031 $[-0.00846, +0.00473]$

These contrasts use native support for each cadence, 8,920 pentagon observations for 1 s histories and 5,432 for 5 s histories. Compare spacings only on shared support, as in Table 1. The interval resamples the five UTC day blocks 2,500 times, pooling route errors within each draw. They provide little information on transfer to other regimes. No per-route independence assumption is needed for the paired mean.

Reading the scores

At 1 s spacing and the 1 s tenor, gradient boosting leads the displayed rows on shared support (0.2424 log bp). At 5 s spacing the full classical MPNN leads them (0.2488). On native 1 s histories, adding pentagon supervision raises MAE in both seeds for both patches. On 5 s histories, the classical improvements differ substantially by seed, and the coherent signs differ. The effect depends on history and initialization.

The independent audit reproduced all 56,550 scalar comparisons within 1.67×10^{-16}. Eight selected MPNN CPU replays agreed with the sealed forecasts within 3.1×10^{-6} log bp. For 5 s gradient boosting, five of 947,200 CPU replay outputs exceeded 3×10^{-5} log bp, with maximum 0.04839. Aggregate scoring passed; the audit's all-checks flag is false because of this replay discrepancy.

METHOD AND REPRODUCTION

Exact single-wire cyclic-qudit readout

A six-feature route vector supplies angles a_i on w wires of radix q. Write $C_s = \cos(2\pi s/q)$ and $p_i(s) \propto \exp[\cos(a_i - 2\pi s/q)]$. The prepared product amplitudes are $\prod_i \sqrt{p_i(z_i)}$. A diagonal unitary adds phase

$$\Phi(z) = \sum_i a_i C_{z_i} + \frac{2\pi}{q} \sum_{i=1}^{w-1} \kappa_i z_i z_{i+1}.$$

To trace out every wire except i, fix its two indices s, s' and sum over a common environment e. The ket and bra phases contribute $e^{\mathrm{i}[\Phi(s,e)-\Phi(s',e)]}$. Terms involving only environment indices cancel. Product probabilities sum to one away from i and its neighbors. The exact remaining expression is

$$\rho_i(s,s') = \sqrt{p_i(s)p_i(s')}\, e^{\mathrm{i}a_i(C_s - C_{s'})} \prod_{j\sim i} \left[\sum_{t=0}^{q-1} p_j(t) e^{\mathrm{i}(2\pi/q)\kappa_{ij}(s-s')t} \right].$$

With $|+\rangle = q^{-1/2} \sum_s |s\rangle$, apply $U_i = I + (e^{\mathrm{i}\theta_i} - 1)|+\rangle\langle+|$. Mixers on every traced subsystem disappear under the partial trace. The measured scalar is $\mathrm{Tr}[\mathrm{diag}(C_s)\, U_i \rho_i U_i^\dagger]$. Thus the six-wire, five-level coherent patch has $5^6 = 15{,}625$ basis states but needs only $O(wq^2)$ single-wire storage and $O(wq^3)$ arithmetic at fixed degree. All measurements are classical GPU calculations. Joint measurements or a different circuit connectivity require a fresh contraction argument.

Circuit and graph controls

The product-phase control removes pair phases and keeps the mixers. The dephased control evolves local probabilities with $|U_i|^2$. The max-plus control retains the best incoming log mass per output and normalizes the resulting scores. Tropical outer attention changes the graph attention kernel separately. Dense-reference checks cover values and gradients in 16 settings at radices 2, 3 and 5, plus a legacy four-qutrit comparison and a twelve-wire finite-gradient check. A separate check accompanying this addendum explicitly evolves the full six-wire state at radices 3 and 5. Across the twelve wire readouts, its largest discrepancy from the marginal formula is 4.45×10^{-16} or less on the two fixed mathematical inputs.

Reproduce the comparison

Use the PhaseWeave research environment, verified licensed archives and a fresh destination. The wrapper fits and seals both cadences, then validates their selections and predictions before opening held-out targets. The registry includes seven spatial backbones at 1 s spacing and a focused four-arm MPNN comparison at 5 s.

```
manifest=4f6bc04f1ab9ebec085408c6fd1e9cb61c6c465233abadcd40e6490b08f4245a
./scripts/train_fx_cycle_k5.sh "$K5_DATASET" "$manifest" "$FRESH_CAMPAIGN"
```

SHA-256 of the aggregate score receipt

```
064fb473ba2fe5d7b11033ed560046c2652edc783e19b5374c56f3264a01b22c
```

PhaseWeave public lab. Local source notes are `docs/FX_K5_DATASET.md` and `docs/FX_K5_CAMPAIGN.md`.