Historical FX replication

Circuit phase across training seeds

A two-seed experiment suggested that a nearest-neighbor entangling phase improved a two-level clock-circuit feature for forecasting one-second FX triangle margins. We repeated the six matched circuit settings with twelve new initializations while preserving the model and training regime. The estimated two-level phase effect changed sign on both previously scored panels. Paired seed-by-day intervals spanned zero. At three levels per wire the phase increased mean error on both panels, with pointwise crossed intervals above zero. At five levels the crossed intervals included zero.

Design and estimand

Each width-200 TokenGT adds four circuit clock observables to the spatial representation. The wires have $q = 2, 3, 5$ levels. The phase arm adds the diagonal nearest-neighbor phase. Both arms retain the local phases and condition local mixers on neighboring input angles. All 72 models use 1,537,657 trainable parameters, the same 8,192 training windows, batch size 64, and 384 updates. Within each seed all six arms share initial parameters and training batches. All selected epoch three by validation MAE on the same 2,048 windows. This is a classical CUDA simulation.

For each seed and eligible row, we subtract the product-arm absolute error from the phase-arm error, averaging the two route orientations. We give equal weight to rows and to seeds. For the crossed interval, resample paired seeds and entire UTC-day groups independently, repeating all rows of a repeated day. The intervals use the 2.5th and 97.5th percentiles from 10,000 draws, with no multiplicity adjustment. Owen (2007)

Panel	q	Mean	Day only 95%	Crossed 95%	Improved
October	2	+0.001232	[+0.000774, +0.001681]	[−0.002798, +0.005052]	5
October	3	+0.005782	[+0.005074, +0.006473]	[+0.001683, +0.010024]	4
October	5	+0.004059	[+0.003472, +0.004628]	[−0.000027, +0.008025]	4
November	2	+0.002647	[+0.002272, +0.003043]	[−0.002673, +0.008116]	4
November	3	+0.006811	[+0.006165, +0.007434]	[+0.000039, +0.013951]	4
November	5	−0.000536	[−0.001637, +0.000541]	[−0.004940, +0.004231]	6

Values are log basis points. Negative effects favor the bond phase. Improved counts the seeds whose paired mean favors the phase, out of twelve.

Coverage. October has 2,048 windows on 26 UTC days. November has 2,042 eligible windows from 5,120 forecasts on 10 UTC days. Neither panel is a new set of market observations for the project. The forecast files were sealed before this evaluation opened the archived outcomes.

Interpretation. Day-only intervals are substantially narrower. The $q = 2$ effect is sensitive to the training initialization, even with identical architecture and data. All 72 new fits beat persistence in October, and 65 did so in November. The three-level result is specific to this matched control, graph, panels, and forecast horizon. The experiment evaluates prediction error on indicative quotes. Interactive results

Paired cluster calculation

For market row i with two directed labels y_{ij}, let $e_{si}^{A} = \frac{1}{2}\sum_{j=1}^{2} |\widehat{y}_{sij}^{A} - y_{ij}|$ for arm A. Define $L_{sd} = \sum_{i \in d}(e_{si}^{E} - e_{si}^{P})$ and $n_d = |\{i : i \in d\}|$ for the entangling and product arms E and P. The observed estimator and resampled estimator are

$$\widehat{\Delta} = \frac{\sum_{s,d} L_{sd}}{12 \sum_d n_d}, \qquad \widehat{\Delta}^{*} = \frac{\sum_{s,d} u_s v_d L_{sd}}{12 \sum_d v_d n_d}.$$

The seed multiplicities u_s come from 12 multinomial trials with equal probabilities on the 12 seeds. The day multiplicities v_d come from D trials with equal probabilities on D observed days. The fixed-day interval sets each $v_d = 1$. The fixed-seed interval sets each $u_s = 1$. Contrasts within a panel reuse weights. Both panels share seed draws and use separate day draws.

Panel	q	Seed SD	Seed only 95%	Crossed 95%	Signs match
October	2	0.007103	[−0.002734, +0.004936]	[−0.002798, +0.005052]	11
October	3	0.007633	[+0.001673, +0.009879]	[+0.001683, +0.010024]	8
October	5	0.007353	[−0.000013, +0.007892]	[−0.000027, +0.008025]	10
November	2	0.009943	[−0.002588, +0.008131]	[−0.002673, +0.008116]	11
November	3	0.012896	[+0.000027, +0.013903]	[+0.000039, +0.013951]	8
November	5	0.008226	[−0.004821, +0.004006]	[−0.004940, +0.004231]	10

Signs match counts the seeds with the same phase-effect sign across panels. These intervals retain within-day dependence. They do not model serial dependence across days. There are ten November blocks.

Persistence and the six circuit settings

The control carries each current margin forward by one second. Its October MAE was 0.206430 log basis points. Its November MAE was 0.180125. Each fitted setting averages twelve seeds.

q	Diagonal phase	Panel	MAE	Fits beating persistence
2	Entangling	October	0.187327	12 / 12
2	Product	October	0.186095	12 / 12
3	Entangling	October	0.193313	12 / 12
3	Product	October	0.187531	12 / 12
5	Entangling	October	0.196972	12 / 12
5	Product	October	0.192913	12 / 12
2	Entangling	November	0.167172	11 / 12
2	Product	November	0.164524	11 / 12
3	Entangling	November	0.173327	10 / 12
3	Product	November	0.166517	12 / 12
5	Entangling	November	0.174253	10 / 12
5	Product	November	0.174790	11 / 12

Reproducibility

Each seed is derived from the first four bytes of SHA-256 of a fixed domain string and an integer from zero through eleven. Source manifests and sealed aggregates record the fit, forecast, and scoring hashes. CUDA inference uses one float32 process and batch size 64 for all arms and both panels. November eligibility follows the prior one-second clock and quote-age gate. Aggregate score with all cells and intervals