

Online Resource 1

Historical-Memory Equity Selection: A Comparative Study of Retrieval, Neural Prediction, and Adaptive Mixing

Rohil Gujarathi Sangeeta Oswal Vaibhav Goyal Anjali Yeole

Digital Finance

Corresponding author: Rohil Gujarathi

Department of Artificial Intelligence and Data Science, Vivekanand Education Society's Institute of Technology, Mumbai 400074, Maharashtra, India

Rohil.wrk@gmail.com

S1 Scope and study chronology

This supplement documents the memory-centered equity-selection framework, the 2020–2025 walk-forward evaluation, fitting and execution specifications, the feature dictionary, and explicitly labeled legacy development branches. An admitted memory forecast is defined by explicit eligibility rules, selected historical record identifiers, distances, stored outcomes, and aggregation weights. When the query and neighbour records are retained, these components provide the case-level provenance needed to inspect and reconstruct signal formation. The final performance and inference results were reproduced by the complete 204-account production execution sealed on 18 September 2026. Component branches in Section S2 and the diagnostic in Section S5 remain legacy 2024 evidence; their point estimates do not describe the production-verified six-year comparison. The tracked analytical release is separate from this Online Resource and excludes the largest prediction, ledger, and coverage objects; Section S7 specifies the available artifacts and verification scope.

The earlier benchmark reconciliation and Batches 1–4 contained 636 simulation cells with a different score/architecture contract. The subsequent single-year memory study evaluated pooled-window predictors and their retrieval components, with 180 execution cells in its first matrix and 150 in its 2024 comparison. The production-verified final evaluation contains 204 continuous market/run accounts across 16 policies and six markets. These counts are execution records, not independent market histories; repeated reference arms and annual folds are not new independent markets.

Presentation labels map to final policy IDs as follows: memory only is MEM_SIM, random memory is MEM_RANDOM, historical prior is HIST_PRIOR, and plain nearest neighbors and annual ridge are KNN_PLAIN and RIDGE_ANNUAL. Neural arms use prefixes MLP_ or TRANS_ with suffixes BASE, GATE, MIX_MSE, and MIX_SR. Momentum, volatility-scaled momentum, and passive equal weight are MOMENTUM_21, VOL_MOMENTUM_21, and PASSIVE_EQUAL_WEIGHT. Earlier development scores do not define the multi-year policy ranking.

S2 Legacy component definitions and development evidence

Pooled input-window retrieval references use 252 sessions of 23 standardized features, pooled into 42 six-session means and flattened row-major to 966 components. The recovered MLP, annual ridge, and retrieval policies all consume this 966-dimensional representation; the recovered Transformer

Table S1: Chronology and interpretation of the final study.

Stage	Interpretation
Fold-specific history	For evaluation year y , training records and memory-bank origins end by 31 December of $y - 3$; validation uses $y - 2$, development uses $y - 1$, and all labels must mature by their applicable cutoff.
Development year $y - 1$	Fixed-mixture selection and gate fitting use only the immediately preceding development year.
2020–2025 evaluation	Six continuous annual folds with account state carried across boundaries; 1,568 union sessions in total.
18 September 2026	Full six-fold production execution, continuous portfolios, P1–P8 inference, and archived-return replay completed and verified.
Interpretation	The run is chronologically walk-forward but not an untouched holdout or point-in-time constituent study.

consumes the corresponding 42×23 matrix. Some legacy tables below used a narrower MLP representation and are labeled accordingly. The final memory predictor selects 25 eligible precedents and averages their 63-session returns. The value of a learned representation for forecasting is separate from its value as a distance metric.

S2.1 First memory matrix

M0 is the direct backbone; M1 uses constrained random-priority eligible retrieval; M2 uses pooled-input Euclidean distance; M3 uses frozen latent cosine similarity; and M4 applies a fixed softmax kernel with temperature 0.08 to the same M3 neighbors. Memory forecasts in this branch enter with fixed coefficient 0.25. M1 in this table is a hybrid, not the pure random-memory control in the final comparison.

Table S2: Legacy Round 1 results, before MLP input alignment. Returns and drawdowns are percentages.

Backbone	Mode	MSE	Rank IC	Return	Sharpe	Max DD
MLP	M0	0.027602	-0.0192	14.86	0.758	-15.03
MLP	M1	0.027116	-0.0194	12.93	0.785	-14.60
MLP	M2	0.026146	0.0029	10.43	0.659	-14.84
MLP	M3	0.027657	-0.0177	11.99	0.595	-15.24
MLP	M4	0.027668	-0.0176	11.82	0.601	-15.78
Transformer	M0	0.057303	0.0091	3.71	0.275	-18.47
Transformer	M1	0.044269	0.0091	1.91	0.202	-17.99
Transformer	M2	0.043011	0.0135	7.20	0.410	-18.42
Transformer	M3	0.053426	0.0098	1.02	0.100	-19.43
Transformer	M4	0.053427	0.0098	1.02	0.100	-19.43

The top Transformer neighbors differed in cosine similarity by about 0.004. At temperature 0.08, this range implies a largest-to-smallest weight ratio of $\exp(0.004/0.08) \approx 1.05$, so the M3 weights were nearly uniform. That pattern agrees with the small M3–M4 portfolio differences, although their forecast MSE values are not identical. After centering, the leading latent component explained approximately 64–75% of Transformer variance, and the diagnostic found no dead dimensions. Centering improved the portfolio point estimate for one seed only, so we treat the observation as a

prompt for further study rather than evidence for a geometric mechanism.

S2.2 Keys, aggregation, and value branches

K0 retains the M3 latent key. K1 substitutes a random 32-dimensional projection, whereas K2 learns a 32-dimensional key. All aggregation tests use K2 neighbors: A_{unif} assigns equal weights, A0 applies a fixed softmax, A1 adds a query-conditioned logit adjustment, and A_{ctrl} is the query-only capacity control. The value tests instead return to the M2 input-window neighbors. V0 stores the 63-session return, V1 also stores the 21-session return, and V2 adds drawdown. These branches answer separate questions and were not accumulated into one model.

Table S3: Legacy Rounds 2–4 portfolio results, before MLP input alignment. Earlier branch-specific significance labels are not treated as inference for the final comparison.

Branch	Mode	MLP Sharpe	Transformer Sharpe	Transformer return (%)
Keys	K0	0.595	0.100	1.02
Keys	K1	0.694	0.141	1.60
Keys	K2	0.507	0.253	3.38
Aggregation	A_{unif}	0.507	0.253	3.38
Aggregation	A0	0.508	0.171	1.57
Aggregation	A1	0.507	0.171	1.57
Aggregation	A_{ctrl}	0.560	0.154	1.84
Values	V0	0.659	0.410	7.20
Values	V1	0.659	0.410	7.20
Values	V2	0.652	0.397	6.82

None of these additions produced a consistent improvement in the development exercise, so we left them out of the final comparison. That decision applies only to the variants shown here. The experiment does not rule out other learned metrics, aggregation rules, or outcome definitions.

S2.3 Mixing and risk extension

In the early gate experiment, the reference forecast used a constant memory weight of 0.25. For the later walk-forward study, we reselected the MSE and Sharpe constants from the immediately preceding development year for each fold. We fitted the logistic gate on that same development period and froze it before evaluation. Its inputs are available at query time; future returns and the post-outcome accuracy diagnostic never enter the gate. The coefficients and risk results below belong to the 2024 development branch and were not reused across all six folds.

The secondary risk extension in `memory_study/run_round6_risk.py` computes a weighted mean over the lower q mass of neighbor returns. After sorting the returns as $r_{(1)} \leq \dots \leq r_{(k)}$, with normalized weights $w_{(j)}$, the code uses

$$\omega_j(q) = \max \left\{ 0, \min \left(w_{(j)}, q - \sum_{h < j} w_{(h)} \right) \right\}, \quad (1)$$

$$\text{TM}_q = \frac{1}{q} \sum_{j=1}^k \omega_j(q) r_{(j)}, \quad S_{i,t}^{\text{risk}} = \frac{\hat{r}_{i,t} - \beta |\text{TM}_{0.05,i,t}|}{v_{i,t}^{21} + 10^{-4}}. \quad (2)$$

With 25 equal weights and $q = 0.05$, this calculation gives $\text{TM}_{0.05} = 0.8r_{(1)} + 0.2r_{(2)}$; the second return supplies the remaining one-percent mass. The absolute value charges a penalty even when the lower-tail mean is positive, so this is not a negative-part loss. Returns and volatility are expressed as decimals, and β has no units. The runner examines up to 250 provisional candidates. If fewer than 25 pass the filters, it emits a zero memory forecast and sets the tail mean to -0.10 . Table S4 uses β for this penalty coefficient. Changing the score alone leaves the forecasts, MSE, and IC unchanged. The final comparison uses $\beta = 0$ and keeps the same stop and holding rules.

Table S4: Legacy risk-extension point estimates, before MLP input alignment; not results of the corrected gate.

Backbone	Forecast	β	Return (%)	Sharpe	Max DD (%)
MLP	Gate	0.00	15.37	0.803	-12.52
MLP	Gate	0.10	13.96	0.706	-13.80
MLP	Gate	0.20	16.17	0.799	-14.10
Transformer	Gate	0.00	12.08	0.686	-14.70
Transformer	Gate	0.10	13.28	0.712	-15.44
Transformer	Gate	0.20	11.75	0.615	-15.71
None	Pure memory	0.00	19.61	0.972	-11.16
None	Pure memory	0.20	13.42	0.664	-13.57

S3 Final fitting specification

The recovered release uses 36 fold-specific neural checkpoints: six evaluation folds, two architectures, and seeds 7, 17, and 37. For evaluation year y , training and memory-bank records end on 31 December of $y - 3$, validation uses year $y - 2$, and development selection uses year $y - 1$. The MLP is Linear(966, 64), LayerNorm(64), GELU, Linear(64, 128), LayerNorm(128), and Linear(128, 1). There is no second GELU. Its input is the row-major flattened 42×23 pooled annual representation. The Transformer takes the same 42 six-session means, projects 23 features to width 64, adds sinusoidal positions, and applies two post-norm encoder layers with four heads, feed-forward width 128, GELU, and dropout 0.1. Learned attention pooling precedes Linear(64, 128), LayerNorm(128), and the scalar head. Training, bank construction, validation, and evaluation use the same standardized-and-clipped representation builder.

The resolved neural contract uses equal-market MSE and AdamW with learning rate 10^{-3} , weight decay 10^{-4} , effective batch size 512, and gradient-norm clipping at 1.0. Training runs for at least 5 and at most 50 epochs. Early stopping has patience 5 and an absolute improvement threshold of 10^{-6} on equal-market validation MSE; the earliest checkpoint wins a tie. We did not retrain neural models for production. The runner authenticated and loaded the 36 saved checkpoints. The production receipt records their hashes together with architecture, fold, seed, and completion metadata. These records identify the models used in the run; they do not repeat the original fitting exercise.

For the evaluated gate, write $\tilde{z} = (|\hat{r}_{\text{base}}|, |\hat{r}_{\text{base}} - \hat{r}_{\text{mem}}|, |\hat{r}_{\text{mem}}|)^{\top}$. Standardization is $z_j = (\tilde{z}_j - m_{\text{dev},j}) / \max(s_{\text{dev},j}, 10^{-4})$, using development means and population standard deviations. The logistic gate $g = \text{sigmoid}(a^{\top}z + b)$ minimizes

$$\mathcal{L}(a, b) = \frac{\sum_{q \in \mathcal{D}_{\text{dev}}} w_q [(1 - g_q)\hat{r}_{\text{base},q} + g_q\hat{r}_{\text{mem},q} - r_q^{63}]^2}{\sum_{q \in \mathcal{D}_{\text{dev}}} w_q}, \quad w_q = \frac{N_{\text{dev}}}{6N_{m(q)}}. \quad (3)$$

Here $\mathcal{D}_{\text{dev}}$ contains preceding-year queries whose labels mature by the development cutoff, and $N_{m(q)}$

is the number of such queries in market $m(q)$. Initialization is $a = 0$ and $b = \log(1/3)$. Adam runs for 50 steps with learning rate 0.05 and weight decay 10^{-3} . We then fix the gate and its normalizer for evaluation. The gate receives three values: absolute base forecast, absolute disagreement, and absolute memory forecast. Asset-specific trailing volatility is applied later during allocation. The gate never receives a future outcome or post-outcome accuracy measure, and we do not interpret its output as calibrated epistemic uncertainty.

For each backbone, seed, market, and evaluation year, fixed-mixture selection evaluates

$$\hat{r}_{b,s,q}^\lambda = (1 - \lambda)\hat{r}_{b,s,q}^{\text{base}} + \lambda\hat{r}_q^{\text{mem}}, \quad \lambda \in \{0, 0.10, 0.25, 0.50, 1\}.$$

For MSE, the development objective is

$$J_{b,y}^{\text{MSE}}(\lambda) = \frac{1}{6} \sum_{m=1}^6 \frac{1}{3} \sum_{s \in \{7,17,37\}} \frac{1}{|\mathcal{D}_{m,y-1}|} \sum_{q \in \mathcal{D}_{m,y-1}} (\hat{r}_{b,s,q}^\lambda - r_q^{63})^2. \quad (4)$$

For Sharpe, the development objective is

$$J_{b,y}^{\text{SR}}(\lambda) = \frac{1}{6} \sum_{m=1}^6 \frac{1}{3} \sum_{s \in \{7,17,37\}} \text{SR}_{b,\lambda,m,s}^{\text{dev}}, \quad (5)$$

using the full development portfolio engine. If objective values differ by no more than 10^{-12} , we choose the smaller grid value. The selected coefficient remains fixed throughout the evaluation year. The compact public release does not include the numerical fold-by-objective coefficient table; reproducing the exact values requires the excluded fold-level selection artifacts. Section S7 identifies the pinned production release.

S3.1 Complete final policy contract

Unless a policy states otherwise, an admitted forecast $\hat{r}$ receives score $\hat{r}/(v^{21} + 10^{-4})$. A candidate must have a finite, positive score. We rank candidates by descending score and use the security identifier to break ties. A day with no eligible query produces no order. Retrieval policies fall back to the unconditional bank mean when fewer than 25 required neighbors are available.

Table S5: Complete contract for the 16 final policies.

Policy	Forecast or benchmark definition	Runs/ market
MEM_SIM	Mean 63-session target of the 25 nearest 966-vector records after same-ticker exclusion, maximum three per ticker, and 21-session within-ticker spacing.	1
KNN_PLAIN	Same input, distance, k , target mean, same-ticker exclusion, tie rule, and fallback as MEM_SIM; removes ticker cap and spacing.	1
MEM_RANDOM	Same eligible restrictions as MEM_SIM, but deterministic PCG64 random priority seeded by bank hash, fold, corrected query ID, and realization seed.	3

Policy	Forecast or benchmark definition	Runs/ market
HIST_PRIOR	Unconditional mean of all admitted bank targets, constant within a fold; no query-local or query-ticker conditioning.	1
RIDGE_ANNUAL	Float64 ridge on the 966-vector and 63-session target, equal market weight, penalty 0.001, unpenalized intercept, refit once per fold.	1
MLP_BASE	Direct Linear(966, 64, 128, 1) forecast from the flattened annual pooled input.	3
TRANS_BASE	Direct two-layer Transformer forecast from the 42×23 pooled input.	3
MLP_MIX_MSE	MLP-memory convex mixture with fold coefficient selected by J^{MSE} .	3
TRANS_MIX_MSE	Transformer-memory convex mixture with fold coefficient selected by J^{MSE} .	3
MLP_MIX_SR	MLP-memory convex mixture with fold coefficient selected by J^{SR} .	3
TRANS_MIX_SR	Transformer-memory convex mixture with fold coefficient selected by J^{SR} .	3
MLP_GATE	MLP-memory mixture with the fitted three-input logistic gate.	3
TRANS_GATE	Transformer-memory mixture with the fitted three-input logistic gate.	3
MOMENTUM_21	Direct score $C_t^{\text{model}}/C_{t-21}^{\text{model}} - 1$; it is not divided again by volatility.	1
VOL_MOMENTUM_21	Score $(C_t^{\text{model}}/C_{t-21}^{\text{model}} - 1)/(v_t^{21} + 10^{-4})$.	1
PASSIVE_EQUAL_WEIGHT	At the beginning of the continuous path, allocate 95% of capital equally over the market universe and hold until terminal liquidation.	1

S3.2 Execution and accounting

The execution adapter processes pending exits before entries. At close t , each entry receives budget $E_t/3$; at the next raw as-traded open $O_{i,t+1}^{\text{trade}}$,

$$q_i = \left\lfloor \frac{0.95E_t}{3O_{i,t+1}^{\text{trade}}} \right\rfloor, \quad \text{cost}_i = q_i O_{i,t+1}^{\text{trade}} (1 + 0.001).$$

An entry is cancelled when its price or quantity is non-positive or cash is insufficient. The engine queues no new entries during the final five sessions. Tickers awaiting exit do not occupy candidate slots and may be re-entered. Entry-day age is one, and the running peak includes the entry open and later closes. A stop or age 63 schedules an exit at the next open. The training label and traded payoff therefore cover different intervals. The label runs from the model-price close at signal origin t to the close at $t + 63$, whereas portfolio profit and loss begins at the next as-traded open and may end earlier through a stop or synthetic exit. MSE and gate fitting consequently optimize a forecast proxy rather than the exact path-dependent, cost-adjusted payoff.

For N sessions, close position value V_t , equity E_t , and ledger $\mathcal{L}$,

$$\text{Exposure} = \frac{1}{N} \sum_t \frac{V_t}{E_t}, \quad \text{Turnover} = \frac{252}{N} \frac{\sum_{j \in \mathcal{L}} q_j (p_j^{\text{entry}} + p_j^{\text{exit}})}{2\bar{E}}.$$

We record zero when a denominator is non-positive. Turnover uses two-decimal ledger prices and includes synthetic terminal records. Prices retain the cache’s numeric quote scale and are not converted through foreign-exchange rates. The adapter can process split and dividend actions, but the public production runner supplies an empty action map and a quote-unit multiplier of 1.0. The evaluated run therefore provides no test of nontrivial corporate-action or quote-unit handling. Its cache is a retrospectively assembled vendor dataset, not a contemporaneous point-in-time archive.

The 204-account workflow runs through `scripts/launch_a30_production.py`. Its immediate corrected-recovery predecessor is `scripts/recover_a30_targeted.py`; `memory_study/run_final_comparison.py` belongs only to the legacy single-year study. At the ranking stage, the final engine rejects a missing required forecast, excludes a nonpositive score, and breaks score ties by security identifier. `MOMENTUM_21` ranks on raw 21-session momentum. `VOL_MOMENTUM_21` divides that value by $v^{21} + 10^{-4}$, as do the other forecast arms with their respective forecasts. Model-price volatility uses 21 consecutive returns with $\text{ddof} = 1$. The ATR ratio and other technical inputs use model-price OHLC; fills and valuation use raw as-traded prices.

S4 Final inference and additional metrics

For the final estimand, we first average run-level metrics within each market and then give all six markets equal weight. This differs from computing Sharpe on one diversified cross-market return stream. The production analysis checks every original-sample contrast against the difference between its two displayed policy estimates at tolerance 10^{-10} .

For the bootstrap, we divided 1 January 2020 through 31 December 2025 into 314 calendar weeks with Monday boundaries. Each of the 204 accounts contributes four weekly statistics: $(\sum r, \sum r^2, \sum \log(1 + r), N)$. The resulting array has dimensions $204 \times 314 \times 4$. Every draw applies the same sampled week sequence to all accounts, preserving paired timing. We retain each account’s session count. Sharpe uses 252-session annualization and population standard deviation ($\text{ddof} = 0$), whereas CAGR uses the fixed six-year horizon. Each of the 10,000 draws concatenates non-wrapping four-week blocks until it contains 314 weeks. The default random generator uses seed 42.

For centered draws $e^{(b)} = \Delta^{(b)} - \hat{\Delta}$, the marginal interval is $\hat{\Delta} \pm Q_{0.95}(|e^{(b)}|)$ and the two-sided p-value is $(1 + \sum_b \mathbf{1}\{|e^{(b)}| \geq |\hat{\Delta}|\})/(10,001)$. We adjust P1–P8 together by the Holm procedure. The intervals are first-order symmetric estimates: they are not studentized or simultaneous, and they do not remove survivorship conditioning or convert an inspected period into an untouched holdout.

Every interval crosses zero. Memory’s point estimate is negative against random memory, the historical prior, plain nearest neighbors, annual ridge, direct MLP, and direct Transformer. P8, the Transformer gate comparison, is the only positive gate estimate, and its interval also crosses zero. The production execution reproduced the four-week primary analysis exactly. We retain the earlier single-year two- and eight-week sensitivity table only in the legacy release, not as evidence about six-year sensitivity.

Production confirmation. On 18 September 2026, the runner generated 4,066,060 policy-prediction records for 119,590 fold-specific evaluation queries. It completed all required forecasts without a non-finite prediction, reconstructed 204 continuous accounts over 1,568 union sessions, and reproduced the corrected P1–P8 estimates and probabilities. The archived-return replay also

Table S6: Primary differences in mean market Sharpe over 2020–2025. Marginal symmetric 95% intervals use four-week calendar blocks and 10,000 draws; Holm adjustment covers P1–P8 jointly.

ID	Candidate minus comparator	Δ SR	95% interval	Raw p	Holm p
P1	Memory – random	−0.1777	[−0.4156, 0.0602]	0.1416	0.9911
P2	Memory – prior	−0.1070	[−0.3901, 0.1761]	0.4508	1.0000
P3	Memory – KNN plain	−0.0497	[−0.3043, 0.2049]	0.7061	1.0000
P4	Memory – ridge annual	−0.1984	[−0.4318, 0.0350]	0.0936	0.7487
P5	Memory – MLP direct	−0.1380	[−0.3819, 0.1059]	0.2686	1.0000
P6	Memory – Transformer direct	−0.0290	[−0.2580, 0.2000]	0.7987	1.0000
P7	MLP gate – MLP mixture (Sharpe)	−0.0091	[−0.1940, 0.1759]	0.9200	1.0000
P8	Transformer gate – Transformer mixture (Sharpe)	0.0996	[−0.0689, 0.2681]	0.2415	1.0000

matched the released analysis at tolerance 10^{-10} . These checks confirm the production computation; they do not add independent market histories to the inferential sample.

S5 Legacy post-outcome gate diagnostic

We retain this pre-correction table only to document development history; it was not regenerated for the corrected MLP gate. Define $d = (r - \hat{r}_{\text{base}})^2 - (r - \hat{r}_{\text{mem}})^2$. A positive value means that memory has lower squared error than the base forecast for that query. The report gives a Spearman correlation of 0.2232 between gate weight and d , and Table S7 reproduces its aggregate bins. Because the diagnostic uses post-outcome errors and clustered observations, we do not use it for gate training or significance testing. The report also leaves the backbone and seed pooling of the 25,753 observations partly unspecified. We therefore do not treat them as independent replications across architectures. The first bin contains only seven observations.

Table S7: Legacy pre-correction gate bins; MSE is in squared decimal-return units.

Gate bin	Count	Mean gate	Mean d	Mem. win (%)	Base MSE	Mem. MSE	Gate MSE
[0, .2)	7	.1917	-.025287	42.9	.041203	.066490	.043320
[.2, .4)	5,081	.3422	-.001979	38.8	.018013	.019992	.018245
[.4, .6)	6,878	.5074	.000442	54.8	.015832	.015390	.015449
[.6, .8)	10,398	.6935	.000866	55.4	.030114	.029247	.028751
[.8, 1]	3,389	.8832	.022482	69.8	.060194	.037712	.038482

In this legacy aggregate only, high gate weight is associated with larger memory advantage. Gating need not outperform both component forecasts in every bin: for example, in the highest bin memory MSE is below gated MSE. Improved aggregate forecast accuracy also does not entail higher net portfolio Sharpe.

S6 Feature dictionary and universe

The following dictionary specifies the final `memory_study_v2` feature builder before pooled standardization. Let C^{model} , H^{model} , L^{model} , and V^{model} denote model-price OHLC and split-normalized volume; execution prices use the separate symbols O^{trade} and C^{trade} . The scaler is fitted once per fold on all admitted pooled training-market rows from 2013 through 31 December of $y - 3$, using population standard deviation floored at 10^{-4} and clipping to $[-5, 5]$. A feature row is valid only after 252 contiguous rows and when all 23 raw features are finite; unavailable values are not execution-time imputations. The universe manifest lists 108 requested identifiers, of which 103 are retained, but does not establish point-in-time historical membership.

Table S8: Final 23 technical inputs in encoder order. Here $r_t = C_t^{\text{model}}/C_{t-1}^{\text{model}} - 1$, MA_w is an arithmetic rolling mean, s_w is sample standard deviation with $\text{ddof} = 1$, and $\varepsilon = 10^{-9}$.

#	Implementation field	Formula and minimum local history
0	<code>return_1</code>	r_t ; 2 model-price rows.
1	<code>momentum_3</code>	$C_t^{\text{model}}/C_{t-3}^{\text{model}} - 1$; 4 rows.
2	<code>momentum_10</code>	$C_t^{\text{model}}/C_{t-10}^{\text{model}} - 1$; 11 rows.
3	<code>momentum_21</code>	$C_t^{\text{model}}/C_{t-21}^{\text{model}} - 1$; 22 rows.
4	<code>volatility_21</code>	$s_{21}(r)$ over returns $t - 20, \dots, t$; 22 rows.
5	<code>volume_mean_21</code>	$\text{MA}_{21}(V^{\text{model}})$; 21 rows.
6	<code>volume_ratio_21</code>	$V_t^{\text{model}}/(\text{MA}_{21}(V^{\text{model}}) + \varepsilon)$; 21 rows.
7	<code>volume_change_1</code>	$V_t^{\text{model}}/V_{t-1}^{\text{model}} - 1$; returns 0 when prior volume is 0; 2 rows.
8	<code>intraday_range</code>	$(H_t^{\text{model}} - L_t^{\text{model}})/(C_t^{\text{model}} + \varepsilon)$; 1 row.
9	<code>drawdown_252</code>	$(C_t^{\text{model}} - \max_{u=t-251}^t C_u^{\text{model}})/(\max C^{\text{model}} + \varepsilon)$; 252 rows.
10	<code>trend_slope_21</code>	OLS slope of 21 closes on $0, \dots, 20$, divided by $C_t^{\text{model}} + \varepsilon$; 21 rows.
11	<code>close_vs_mean_50</code>	$C_t^{\text{model}}/\text{MA}_{50}(C^{\text{model}}) - 1$; 50 rows.
12	<code>close_vs_mean_150</code>	$C_t^{\text{model}}/\text{MA}_{150}(C^{\text{model}}) - 1$; 150 rows.
13	<code>close_vs_mean_200</code>	$C_t^{\text{model}}/\text{MA}_{200}(C^{\text{model}}) - 1$; 200 rows.
14	<code>mean_200_trend_20</code>	$(\text{MA}_{200,t} - \text{MA}_{200,t-20})/(\text{MA}_{200,t-20} + \varepsilon)$; 220 rows.
15	<code>above_low_252</code>	$(C_t^{\text{model}} - \min_{u=t-251}^t C_u^{\text{model}})/(\min C^{\text{model}} + \varepsilon)$; 252 rows.
16	<code>from_high_252</code>	$(C_t^{\text{model}} - \max_{u=t-251}^t C_u^{\text{model}})/(\max C^{\text{model}} + \varepsilon)$; 252 rows.
17	<code>up_down_volume_50</code>	Sum of V^{model} on positive close changes divided by the sum on negative changes plus ε ; 51 rows.
18	<code>rsi_14</code>	Mean gain divided by mean gain plus mean loss plus ε ; flat windows return 0.5; 15 rows.
19	<code>atr_ratio_14</code>	Mean 14-row true range divided by $C_t^{\text{model}} + \varepsilon$; 14 rows.
20	<code>macd_signal_diff</code>	$(\text{EMA}_{12} - \text{EMA}_{26} - \text{Signal}_9)/(C_t^{\text{model}} + \varepsilon)$ with recursive EMAs; 34 rows.
21	<code>bollinger_width_20</code>	$4s_{20}(C^{\text{model}})/(\text{MA}_{20}(C^{\text{model}}) + \varepsilon)$; 20 rows.
22	<code>vol_ratio_63_21</code>	$s_{63}(r)/(s_{21}(r) + \varepsilon)$; 64 rows.

S7 Replication scope and artifact provenance

The final tracked production release is *Core-RL-Agent* commit `fe56a39960c90ac7da3dce9973035f9b4f5baa80`. The production receipt identifies execution code revision `16ca345fd40a8cce76c9ac83ea7aef110c7c79a2`. The historical-memory-equity-data v1.0.3 commit `0dba6f88504fffc802e915ff240177b515a2d8b9` is an upstream/legacy data reference, not the final 204-account production release.

Pinned production release and final scripts:

<https://github.com/Rohil72/Core-RL-Agent/tree/fe56a39960c90ac7da3dce9973035f9b4f5baa80>

Pinned data:

<https://github.com/Rohil72/historical-memory-equity-data/tree/0dba6f88504fffc802e915ff240177b515a2d8b9>

The receipt identifies a separate archive `historical-memory-equity-data-v1.0.3.zip` (4,963,342 bytes), with reported SHA-256 `43753023ecc69797d5e37098fa1cd5ad0f43406d57fda6679956d2bc9447e596`.

The release listing we inspected contains GitHub-generated source archives rather than this named asset, and the byte sequences are not interchangeable. For replication, use the pinned data-repository checkout. We have not independently verified the named archive’s availability, size, or digest.

Tracked production replay. The tracked release provides the 204-account daily-return series, P1–P8 contrasts, four-week bootstrap specification, replay report, runtime configuration, six fold summaries and prediction manifests, production receipt, pipeline-completion record, portfolio manifest, and final scripts. Its production entry point is `scripts/launch_a30_production.py`; `memory_study/run_final_comparison.py` is the older single-year runner. The replay status `REPLAY_VERIFIED` means that statistics recomputed from the supplied returns agree with the corrected recovery result at tolerance 10^{-10} . It does not rerun checkpoint training.

Coverage and input limits. Production summaries count 4,066,060 prediction records, with no missing required forecast or non-finite value. The public commit retains summaries or digests for the large execution objects but not the full coverage manifest, prediction arrays, account ledgers, decision logs, checkpoints, or fold-specific mixture coefficients. Forecast accounting is therefore complete for admitted queries. From the daily-return series alone, however, an extended zero-return period cannot always be separated into cash-only policy behavior versus a period without an admitted query. The public runner also supplies empty corporate-action maps and a unit quote multiplier. As a result, the tracked release does not demonstrate nontrivial split, dividend, or quote-unit handling. We make no point-in-time historical-universe claim.

Levels 2 and 3: deeper reconstruction. Rerunning forecasts and portfolios would also require the feature and price inputs, fitted fold transformations, authenticated neural checkpoints, random-memory seed derivation, and execution records. Those large objects are outside the tracked analytical commit. Code and manifests alone therefore do not provide an independently executable forecast-to-portfolio chain. Reacquiring vendor data may also fail to reproduce the evaluated vintage.

During production, the runner authenticated and loaded the 36 neural checkpoints, retained the corrected query identities and random-memory control, and reproduced the corrected recovery statistics. The receipt, fold manifests, pipeline-completion record, and replay report document this lineage. They do not test issuer isolation, raw vendor-data validity, or historical universe membership.

Online Resource 1 is a scientific supplement rather than the experimental archive. The code uses the Apache-2.0 license. The data repository applies CC-BY-4.0 to covered derived artifacts and lists its source-data exclusions; underlying vendor data remain subject to provider terms. We do not claim an archival DOI.

The main performance result is comparative. On the production-verified mean-market Sharpe point estimate, similarity memory trails random memory, the historical prior, plain nearest neighbors, ridge, MLP, and Transformer. The uncertainty intervals still allow modest benefits as well as harms, so non-rejection is not evidence of equivalence. Separately, the retrieval, aggregation, scoring, and execution contracts define a reproducible, audit-oriented decision architecture. When query records

are retained, the memory component preserves the provenance needed to inspect and reconstruct signal formation. The public-artifact limits stated above still apply.