TRADING WITH MEMORY: HOW ACCUMULATED EXPERIENCE SHAPES LLM TRADING BEHAVIOR

Anonymous authors
Paper under double-blind review

ABSTRACT

Powerful LLM agents are increasingly used for complex decision-making, but trading poses a distinctive challenge: decisions unfold over long histories, while noisy market data make it difficult to know what past information should matter. We ask: when an LLM trading agent takes over a portfolio, what does it inherit?

It inherits a way of seeing the portfolio, not just the portfolio itself. Our experiments show that the same portfolio, exposed to the same future market, can lead to different trading decisions depending on the agent's history. Agents that retain their experience trade 47% less than those starting with a blank history, with 81% of the difference emerging in the first continuation session. The key is not simply remembering what happened. Agents trade even less when they retain their previous investment judgments, suggesting that they carry forward their own beliefs about why positions should be held. This behavioral continuity persists over much longer histories and subsequent trading periods.

These findings suggest that experience changes more than what an LLM trading agent knows. It changes how the agent interprets what it currently holds—and therefore what it does next.

1 INTRODUCTION

A trading decision changes the next decision. Buying an asset uses cash, alters risk, and leaves a position to maintain, increase, or unwind tomorrow. A long-running agent also accumulates explanations of its choices, plans for the portfolio, and beliefs about individual assets. These records may support consistent decisions, but they also preserve commitments formed under earlier market conditions. Understanding sustained portfolio management therefore requires following both what an agent remembers and what it does with the account.

LLM trading systems already store observations, reflect on past outcomes, and revise investment beliefs (Yu et al., 2024a; Li et al., 2024; Yu et al., 2024b). Broader evaluations show that results from selected assets and short windows can change under longer periods and larger universes (Li et al., 2025b). This raises a behavioral question: *how does the content of past experience change reallocation, and how long do the resulting portfolio differences persist?*

Experience and financial exposure normally evolve together. A successful trade changes both the account and the agent's explanation of its strategy. Comparing independently generated trajectories mixes these effects. We separate them at a saved checkpoint: agents begin with identical cash and holdings, receive different versions of their prior history, and then trade independently through the same market window. This intervention reveals how remembered facts and judgments affect a new allocation decision and how that decision carries into subsequent sessions.

We develop AGENTICFOLIO to support this analysis. A single agent chooses information queries, maintains a global plan and per-asset theses, and proposes sized trades in a persistent multi-asset account. The environment validates and executes orders, recording observations, strategy revisions, fills, and portfolio state. Separate account and agent checkpoints make the history interventions reproducible. Historical replay covers U.S. equities, Chinese A-shares, and cryptocurrencies. The same core decision and accounting interfaces also support forward paper trading; the experiments here use historical replay.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Across three API models and 24 windows, retaining factual records reduces reallocation relative to a history-cleared continuation, and retaining investment judgments reduces it further. Eighty-one percent of the Retain–Restart turnover gap arises in the first session; the positions created then remain different after trading rates draw closer. Six longer-history pairs and matched review traces test allocation continuity and link specific judgment edits to orders.

2 Related Work

How trading agents use experience. Trading agents carry experience forward as retrieved evidence, reflection, or an editable policy. FinMem and FinAgent retrieve earlier observations and analyses through layered or diversified memory, respectively; CryptoTrade reflects on realized trading outcomes (Yu et al., 2024a; Zhang et al., 2024; Li et al., 2024). FinCon uses episodic self-critique to revise investment beliefs in a manager–analyst system (Yu et al., 2024b). SHARP and EvolveTrade alter the decision procedure more directly: SHARP validates edits to explicit condition–action rules in a walk-forward setting, while EvolveTrade revises a tool-using policy from decision traces and portfolio feedback (Chen et al., 2026; Kim et al., 2026). We compare inherited facts and judgments at a shared account checkpoint; a scheduled-review condition also records how written beliefs are updated.

Behavioral effects of saved experience. RetailAgent provides the closest behavioral test of self-authored memory: in an anonymized intraday long/flat task, saved narratives increase policy persistence and make exposure-matched timing more negative on stock-days where the agent uses both actions (Zhang et al., 2026). We compare factual records with saved investment judgments under identical starting holdings, then measure whether the first trading response remains in the portfolio as prices evolve.

Sequential portfolio evaluation. InvestorBench and StockBench evaluate historical decisions at different levels: in InvestorBench's multi-asset setting, an optimizer sizes positions from model-generated directions, whereas StockBench asks agents for sequential portfolio allocations using prices, fundamentals, and news (Li et al., 2025a; Chen et al., 2025). FINSABER tests broader asset and historical coverage; KTD-Fin controls historical identity cues and decomposes returns; CLQT adds cost-aware, auditable portfolio rounds and repeated-run diagnostics (Li et al., 2025b; Zhu et al., 2026; Qu & Chen, 2026). LiveTradeBench and AI-Trader evaluate agents with newly arriving market information (Yu et al., 2025; Fan et al., 2025). Our paired continuations hold the account and future market fixed, allowing the effect of inherited information on orders and holdings to be measured separately from starting exposure.

3 AgenticFolio

AgenticFolio couples a persistent portfolio with one tool-using daily agent (Figure 1). The environment supplies point-in-time information and execution rules; the agent chooses queries, revises investment records, and requests trades. Observations, proposals, execution responses, and saved state are recorded together.

3.1 Account and execution

At session t, the account contains cash c_t and quantities $q_{t,i}$ in a fixed universe $\mathcal{U}$, with marked value $V_t = c_t + \sum_{i \in \mathcal{U}} q_{t,i} p_{t,i}$. The agent can retain cash, maintain a position, or request a weight change. The environment converts that request into a quantity, checks cash and concentration limits, and executes valid orders. Each fill changes the exposure carried into the next session.

Historical replay follows a daily-close convention. The current close is observable and is also the execution price; the latest eligible news date is the preceding calendar day. The reported runs use long-only, unlevered accounts, integer asset quantities, a 100% per-asset cap, and a 10-basis-point cost on each traded notional. Costs are deducted from cash and included in net asset value. These are shared simulation rules across markets. The starting account, prices, quantity rules, and costs remain identical across history conditions.

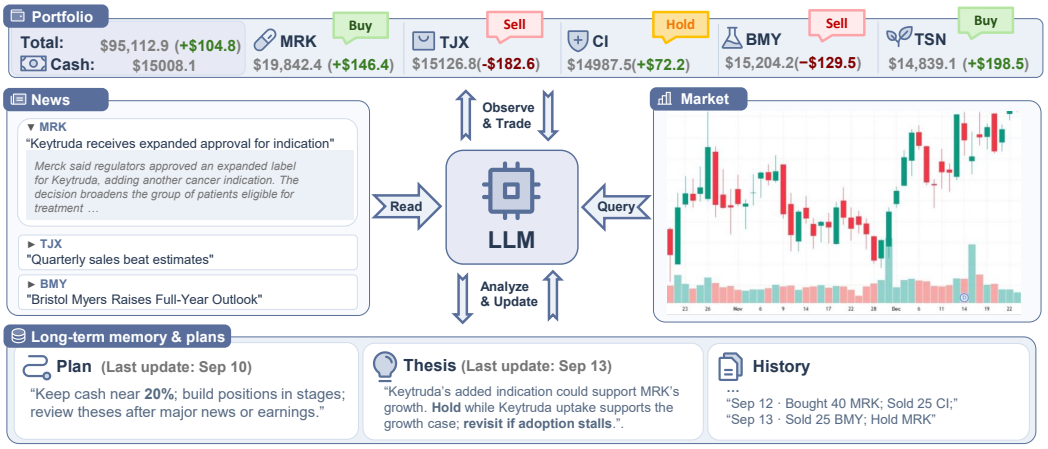

Figure 1: Daily interaction and persistent state. The agent queries news and market data, proposes trades, and reads or revises plans and theses; executed orders update the account, and saved judgments enter later contexts. All amounts, P&L, news, and memory entries are illustrative; the candlestick chart is market context, not a model input.

3.2 Daily Interaction

The agent starts each session with prices, its current portfolio, risk limits, news availability, and selected stored history. It emits structured batches of tool calls to inspect headlines or articles, retrieve past decisions, read or update investment state, and propose trades. Buy and sell tools request a positive change in portfolio weight. A target-weight interface expresses a complete allocation, and a read-only preflight simulates candidate orders. Execution feedback enters the next response, allowing the agent to correct rejected proposals.

Table 1 makes this action space explicit. It separates information retrieval, editable investment state, and order execution: the agent can inspect or revise a thesis, but only an accepted proposal changes the account. The same interfaces remain available in every history condition; the intervention changes what prior state is supplied, not what actions the agent can take.

The agent ends with `claim_done`; a step budget also bounds each session. Accepted trades remain effective if the budget is exhausted. The current completion handler records missing optional checklist fields as warnings. This separates the trading action from the completeness of its end-of-day report. Prompt settings and validation details appear in Appendix B.

3.3 Persistent Investment State

The agent's history has three stores. A bounded decision buffer contains recent actions, summaries, orders, and settled next-day feedback. A global plan stores free-text strategy entries and revisions. A structured per-asset thesis stores the investment view, evidence for and against it, conviction, and exit conditions. Tools can read and edit these records, and selected entries appear automatically in the next daily context. Account-derived risk observations summarize exposure, concentration, drawdown, and recent activity.

An orchestrator assembles this context, runs the daily loop, and saves decisions, tool traces, fills, portfolio snapshots, and revisions. Economic account state is checkpointed separately from agent history. A continuation can therefore change the history supplied to the agent without changing its starting cash, quantities, or cost basis.

In the broad study, automatic rendering includes at most five recent decisions and 6,000 characters of plan text; up to 90 decisions remain queryable. The full event log is retained for audit, with exact query and step budgets listed in Appendix B.

Table 1: Agent-facing interfaces enabled in the broad experiment. Calls use structured JSON arguments and return observations or validation feedback.

Interface	Input and returned or updated state
`list_news`, `get_news`	Ticker/date filters or article IDs; eligible headlines and article text.
`list_decisions`, `get_decision`	Ticker/date filters or decision IDs; retained past decisions and settled feedback.
`read_plan`, `write_plan`	Scope and optional replacement/append text; current general or per-asset plan.
`list_plan_changes`, `get_plan_change`	Scope/date filters or revision IDs; saved plan edits.
`read_thesis`, `update_thesis`	Ticker, status, thesis, evidence, conviction, exit conditions; structured asset view.
`trade_preflight`	Candidate orders; estimated cash, weights, costs, and constraint warnings.
`propose_buy`, `propose_sell`	Ticker, positive weight delta, rationale; immediate validation and execution.
`propose_target_weights`	Nonnegative universe weights including cash, summing to one; sell-then-buy conversion.
`claim_done`	Summary and optional checklist fields; end-of-session decision record.

4 Controlled Continuation Protocol

4.1 Shared accounts and varied histories

A parent agent first trades for H active sessions after price-only warmup. We checkpoint its account and agent state, then initialize child runs with identical cash, quantities, and cost basis. The children receive different histories and trade independently through the same subsequent T-session market window. All conditions can write new plans and theses and accumulate new decisions after the intervention. The shared input is the available market, not a fixed sequence of retrieved articles: agents choose their own information queries.

Table 2 distinguishes five conditions. *Retain* (A) inherits all state normally available through the framework's bounded memory and query interfaces. *Restart* (B) clears agent decisions, plans, theses, and attribution summaries. *Facts only* (C) preserves factual records of recent trades and settled account feedback while stripping decision rationales and investment judgments. *Review* (D) begins with A's history and adds a structured review by the same agent every five sessions. *Cold start* (E) begins with B's cleared state and additionally limits older news and pre-branch account-history observations. A–D retain the ordinary market archive and account-risk history; E initially permits seven calendar days of eligible news. Appendix B gives the exact transformations.

The contrasts separate different sources of continuity. A–B varies the full bundle of inherited agent-authored state; A–C and C–B distinguish prior judgments from factual trade and account records. D–A adds scheduled review without changing the inherited account. E–B jointly restricts older news and pre-branch account-risk history, so it tests their combined availability rather than either source alone. Current holdings, prices, execution rules, and action tools remain common to all arms.

4.2 Markets, models, and experimental scale

The broad study crosses three API profiles with 24 windows: eight U.S. equity, six cryptocurrency, and ten Chinese equity windows (Table 2). Each model–window pair has one 15-session parent after five warmup sessions, followed by five eight-session continuations. The resulting 72 parents and 360 children completed all 3,960 active decision sessions without a recorded failed or forcibly terminated day. Review occurs on continuation sessions 1 and 6.

Table 2: History conditions and market coverage. All branches inherit identical cash and holdings and trade independently over the same subsequent market window.

Information / mechanism	A Retain	B Restart	C Facts only	D Review	E Cold start
Current account + permitted market info	✓	✓	✓	✓	✓
Earlier news archive access	✓	✓	✓	✓	–
Pre-branch account / risk history	✓	✓	✓	✓	–
Inherited factual decision records	✓	–	✓	✓	–
Inherited decision rationales	✓	–	–	✓	–
Inherited global plans	✓	–	–	✓	–
Inherited per-asset theses	✓	–	–	✓	–
Scheduled review (every 5 sessions)	–	–	–	✓	–

Market universe

Market	Tradable assets	Windows	News language
U.S. equities	MRK, TJX, CI, BMY, TSN	8	English
Crypto	BTC, BNB, SOL, XRP	6	English
Chinese equities	CITIC Securities, BYD, CATL, SIASUN, Kweichow Moutai	10	Chinese

✓: retained or enabled; –: withheld or not scheduled. History entries refer to pre-branch state, not whether a tool is available. All branches can write new plans and theses and accumulate new history. The broad experiment uses 15 history and 8 continuation sessions; E retains a seven-day initial news window but excludes older news and pre-branch account history. D reviews on sessions 1 and 6; voluntary updates remain available in every arm.

The windows span U.S. episodes from July 2018 to April 2020, crypto from August 2023 to January 2024, and Chinese equities from February 2021 to December 2024. Their equal-weight continuation returns range from -7.55% to $+47.72\%$, placing the history comparison across both falling and rising markets rather than a single market direction.

The model profiles are DeepSeek-v4-flash, GPT-5.4, and Kimi-K3, accessed through provider routing aliases. All use temperature 0.4 and top-p 0.9 without a provider seed; the DeepSeek profile explicitly disables thinking. The daily loop allows 30 response batches. Exact aliases, configuration values, dates, and window membership appear in Appendices A–B.

A separate longer-history study uses a DeepSeek API profile with $H = 180, T = 30$. Three independent parent histories precede the February 2020 U.S. equity decline, and three precede the October 2023 crypto rally. Each parent produces one A/B pair. We report these six pairs separately from the broad study, extending the duration of the shared-account comparison.

4.3 Measures and statistical comparisons

Our main behavioral measure is departure from passively retaining the starting account. Hold the initial cash c_0 and quantities $q_{0,i}$ fixed and revalue them at each subsequent session's prices. Let $\widetilde{w}_{t,i}$ be these passive weights and $w_{t,i}$ the branch's weights, including cash. Their distance is

$$d_t = \frac{1}{2} \sum_{i \in \mathcal{U} \cup \{\text{cash}\}} |w_{t,i} - \widetilde{w}_{t,i}| \,. \tag{1}$$

A branch that never trades has zero distance even as prices move. In the broad study, we average over eight pre-trade snapshots, including the common initial snapshot. We also measure the distance *between* A and B at each snapshot, replacing passive weights with the other branch's weights in Equation 1.

Turnover sums executed buy and sell notional divided by each session's closing NAV; 0.10 means traded notional equivalent to approximately 10% of NAV, counting both sides. Net return includes all

continuation sessions and transaction costs. Cash weight, first-session trading, and revision records complement these measures. A frozen-position calculation holds each branch's post-session-1 quantities and cash fixed, allowing us to follow the account difference created by the initial decision without subsequent trades.

We pair comparisons within each model–window account and group related windows and all their model runs into 12 market-episode blocks. Exact two-sided block sign-flip tests and 20,000 block-bootstrap resamples quantify differences across these episodes. Benjamini–Hochberg adjustment covers one family of 25 comparisons: A–B, A–C, C–B, D–A, and E–B across five outcomes. We use q for these adjusted values. Appendix C reports every estimate, interval, and test; the timing and trace analyses are follow-up analyses.

5 Experience and Reallocation

Table 3 summarizes the five history conditions.

Table 3: History conditions and portfolio behavior. Each row averages 72 matched model–window accounts over eight sessions.

Condition	Turnover	Distance (pp)	Runs trading on session 1	Cash (%)	Return (%)
A Retain	0.096	3.72	21/72	42.49	3.06
C Facts only	0.141	5.89	33/72	40.88	3.18
B Restart	0.181	9.37	49/72	38.32	3.32
D Review	0.080	3.26	8/72	42.43	3.06
E Cold start	0.135	5.97	34/72	40.86	3.23

Distance is mean pre-trade departure from the passive starting account; cash is mean post-trade cash weight. Return is net of costs. Uncertainty for all paired contrasts appears in Table 10.

5.1 Facts and investment judgments both change reallocation

Retaining history reduces reallocation of the starting account. Mean turnover falls from 0.181 under Restart to 0.096 under Retain, a 47% reduction. The paired difference is -0.085 (95% block-bootstrap interval $[-0.128, -0.052]$, $q = 0.006$). Mean distance from the passive starting account falls from 9.37 to 3.72 percentage points (pp), a 60% reduction ($q = 0.017$). Retain has lower mean turnover in all 12 market-episode blocks; the direction also holds when restricting the analysis to equities or excluding the largest-return window.

Facts only lies between Retain and Restart, at 0.141 turnover. Both adjacent contrasts survive correction ($q = 0.017$ for A–C; $q = 0.038$ for C–B). Thus factual transaction and account records change subsequent trading, and the retained bundle of decision rationales, plans, and theses adds a further effect. Subgroup means vary by model (Tables 12 and 13); confirmatory inference is based on the pooled paired comparisons.

Cold start provides a complementary contrast. Although B and E both clear agent history, restricting older market and account-history observations in E reduces turnover from 0.181 to 0.135 ($q = 0.017$), and distance from 9.37 to 5.97 pp ($q = 0.006$). External and account history can therefore prompt reallocation in the absence of old agent judgments.

The longer continuations show the same allocation-continuity pattern over a 30-session horizon (Figure 2). Retain ends closer to the passively held starting account in all six pairs: mean distance is 4.14 versus 14.21 pp in the U.S. decline and 13.92 versus 45.19 pp in the crypto rally. Cash moves in opposite directions across the two markets: Retain holds less cash than Restart in the U.S. decline (19.19% versus 27.98%), but more during the crypto rally (52.26% versus 21.55%). The shared result is continuity with the inherited allocation, not a uniform shift toward or away from risk. These six DeepSeek pairs are a separate duration extension, not additional replications of the three-model study; all pair-level returns and distances appear in Appendix D.

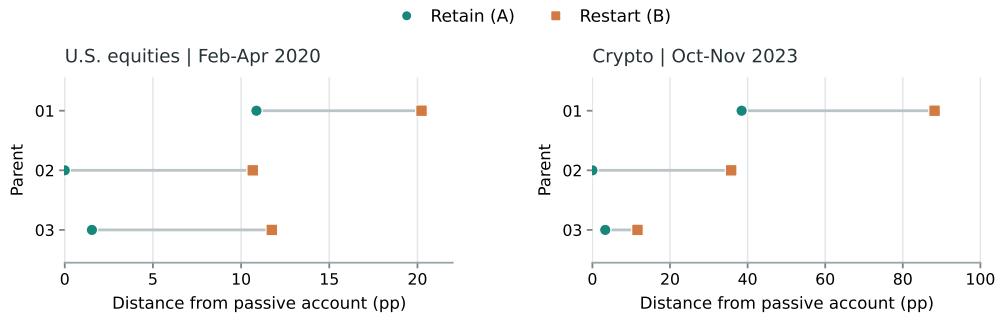

Figure 2: Distance from the passively revalued starting account after 30 sessions in six matched $H = 180, T = 30$ continuations. Each row is one parent account; circles denote Retain (A) and squares Restart (B). Retain is closer in all six pairs. The U.S. and crypto panels use separate horizontal scales for legibility.

5.2 Early reallocation shapes the portfolio path

Most of the broad-study trading response occurs at the first decision after the history intervention (Figure 3a). Retain trades in 21 of 72 groups that session, compared with 49 for Restart. Mean first-session turnover is 0.0166 and 0.0854, respectively. This session accounts for 81% of the total eight-session A/B turnover gap. Over the remaining seven sessions, the corresponding totals are 0.0790 and 0.0951. Meanwhile, Restart rebuilds textual state: 66 of 72 runs have a new plan and 66 have a thesis ledger after session one.

The accounts remain different after this initial burst of reallocation. Before session two's orders, their mean distance is 7.79 pp; before session eight, it reaches 11.88 pp (Figure 3b). Fifty-six of 72 pairs have distinct allocations at that final pre-trade snapshot. Freezing each branch after the first session and revaluing its holdings at the same later prices leaves a session-eight distance of 7.91 pp.

The frozen path retains about two-thirds of the observed session-eight separation (7.91 versus 11.88 pp); later actions account for the remainder. This is distinct from the 81% turnover statistic, which measures how much of the cumulative trading gap occurs on session one. Together, the two comparisons show that early trades establish a persistent account difference, while subsequent decisions can widen it further.

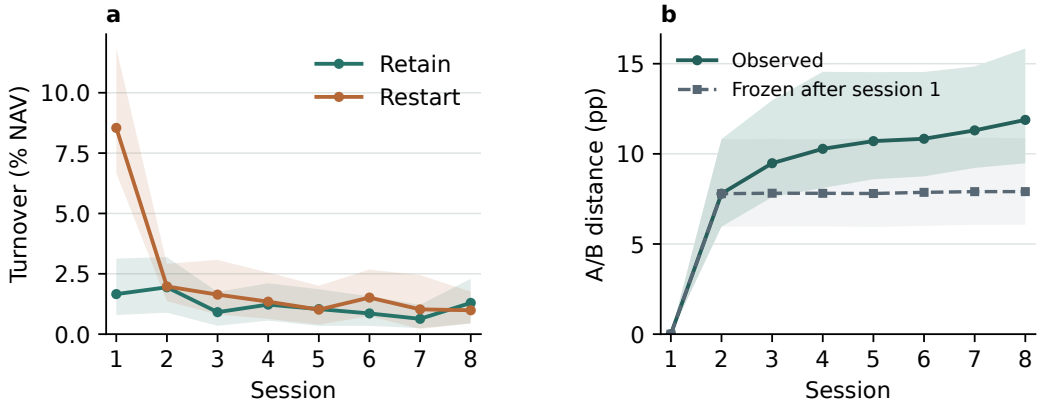

Figure 3: Initial reallocation and subsequent account separation across 72 matched A/B pairs. (a) Executed turnover per session. (b) Half-L_1 distance between cash-and-asset weights before each session, for observed paths and positions frozen after the first session. Session 1 begins from the same account. Lines show means; shaded bands are pointwise 95% episode-block bootstrap intervals. The dashed path uses the same observed prices but no further trades.

Early trading changes both invested capital and asset mix. For Restart minus Retain, we decompose the return gap into exposure, relative allocation, and transaction costs:

$$R_B - R_A = \frac{1}{V_0} \sum_{t=1}^{T} \left[\underbrace{\left(\sum_i \Delta x_{t,i} \right) \bar{r}_t}_{\text{invested capital}} + \underbrace{\sum_i \Delta x_{t,i} (r_{t,i} - \bar{r}_t)}_{\text{relative allocation}} - (K_t^B - K_t^A) \right], \tag{2}$$

Here, V_0 is the common starting NAV; $\Delta x_{t,i}$ is the difference in prior-close dollar exposure to asset i between Restart and Retain; $r_{t,i}$ and $\bar{r}_t$ are its return and the equal-weight return across assets; and K_t^A, K_t^B are transaction costs. Restart's mean eight-session net return exceeds Retain's by 0.261 pp ($q = 0.424$). The decomposition assigns $+0.287$ pp to invested capital and -0.009 pp to costs; relative allocation changes sign across markets (Table 4). The accounting separates greater market participation from differences in which assets carry the exposure.

Table 4: Accounting decomposition of Restart minus Retain net return. Contributions are percentage points of the common starting NAV.

Market	n	Invested capital	Relative allocation	Costs	Total
U.S. equities	24	+0.488	-0.257	-0.010	+0.221
Crypto	18	+0.216	+0.289	-0.008	+0.497
Chinese equities	30	+0.170	-0.010	-0.008	+0.152
All accounts	72	+0.287	-0.017	-0.009	+0.261

5.3 What scheduled review changes

To examine how written judgments relate to action, Review adds a dedicated step before trading on sessions 1 and 6. The same model receives its saved plans and theses, current account and risk state, recent decisions, and up to 50 eligible headlines. For every belief, it returns retain, revise, revoke, or insufficient evidence, with an evidence reference and replacement text when revising. The environment validates and saves these changes before the ordinary trading loop. Appendix B provides the prompt and schema.

Across 144 reviews, the agents retain 84% of belief instances and revise 16%, with four insufficient-evidence outcomes and no revocations. Thesis revisions most often change conviction or exit-condition text, rather than hold/watch/avoid status. The revised fields persist to the end of the session in 69 of 71 thesis revisions. Twenty-one are followed by a same-asset order within that session or the next two. Appendix F gives the field-level counts and matching procedure.

Review usually preserves the allocation. Its mean departure from the passive starting account is 3.26 pp, compared with Retain's 3.72 pp, and the mean return difference is below 0.001 pp. On the 144 scheduled review dates, D trades on 11 days and matched A runs trade on 37. These descriptive counts indicate that explicit review is more often an update to investment conditions than a trigger for immediate reallocation.

Table 5: Matched DeepSeek event traces. Descriptions condense the saved records. Order sizes are accepted target-weight changes, before integer-quantity rounding, rather than exact filled weights.

Asset / date	New evidence	D: saved judgment update	Matched action
BMY 2019-01-04	Acquisition of Celgene; credit-rating review	Takeover-target premise replaced by leverage and integration risk; conviction 0.55 $\rightarrow$ 0.45	A: sell 3% D: sell 4%
Moutai 2023-11-01	Factory-price increase of about 20%	Price increase becomes the leading premise; conviction 0.55 $\rightarrow$ 0.65	A: buy 7% D: buy 5%; add 3.5% next session

Matched event traces show how the updates relate to action (Table 5). For BMY, the Celgene acquisition changes the inherited premise from potential takeover target to acquirer exposed to leverage and integration risk. D lowers conviction and sells; A also recognizes the event and sells a smaller requested fraction. For Moutai, a factory-price increase becomes the leading thesis premise and D raises conviction before buying; A also buys on that date.

6 DISCUSSION

Later trading rates alone understate the effect of inherited history. In the matched continuations, the large first-session reallocation leaves different holdings that persist under the same subsequent prices. Freezing each branch after that first session still leaves a 7.91 pp mean account distance by session eight. The action matters beyond the session in which it is taken: it changes the cash and quantities exposed to every later price movement.

The opposite turnover shifts in the A–B and B–E contrasts show that saved judgments and older market observations enter the trading decision differently. Scheduled review adds another distinction: agents can edit conviction and exit conditions while leaving positions in place. In this sample, 69 of 71 revised theses persist through the session, but only 21 are followed by a same-asset order within three sessions. A record update and a portfolio response are therefore separate parts of the observed process.

These results motivate evaluating experience at both the decision and account levels. The first post-intervention order measures immediate reallocation; distance from a passively revalued starting account measures how far the branch has moved; cross-branch holdings track whether that difference survives under common prices. Terminal return alone combines these choices with the market path, while later turnover alone can miss positions established earlier. Reporting the three views together makes the timing and persistence of an experience effect visible without treating a return difference as its mechanism.

For the central A/B/C contrasts, the implementation forks one checkpoint into the same economic account and price calendar, then changes which prior records are rendered or queryable. The event trace keeps retrieved observations, plan or thesis revisions, proposed orders, and validated fills as separate records. We can therefore follow what the agent elects to inspect, what judgment it carries forward, and which action changes the account, rather than infer all three from terminal return. This alignment between information access, saved state, and execution is a measurement benefit of the pipeline.

The broad study follows one continuation per model–window condition for eight sessions, and six additional DeepSeek pairs extend the comparison to 30 sessions after 180-session histories. Together they show that the allocation-continuity pattern appears in both the broad short continuation and a small set of longer continuations. Evaluating other memory architectures and additional model profiles can establish how far this pattern generalizes.

7 CONCLUSION

Across matched continuations, inherited facts and investment judgments changed the first reallocation of an identical portfolio, and the resulting differences in holdings persisted after later trading rates converged. The results connect stored experience to both the decisions it informs and the account path those decisions create. This puts the timing of reallocation at the center of evaluating experience in long-running trading agents.

AI USE STATEMENT

Generative AI tools were used to assist with developing and refining research questions and experimental analyses, interpreting results, and drafting and editing the manuscript for clarity. The authors reviewed the AI-assisted analyses and text, verified the reported results, and take responsibility for the final content.

REFERENCES

Xiwen Chen, Wenhui Zhu, Songzhu Zheng, Kashif Rasul, Yueyue Deng, and Huayu Li. SHARP: A self-evolving human-auditable rubric policy for financial trading agents, 2026. URL https://arxiv.org/abs/2605.06822v1. Version 1, May 7, 2026.

Yanxu Chen, Zijun Yao, Yantao Liu, Amy Xin, Jin Ye, Jianing Yu, Lei Hou, and Juanzi Li. StockBench: Can LLM agents trade stocks profitably in real-world markets?, 2025. URL https://arxiv.org/abs/2510.02209v2. arXiv preprint; version 2, revised March 2, 2026.

Zihan Dong, Xinyu Fan, and Zhiyuan Peng. FNSPID: A comprehensive financial news dataset in time series, 2024. URL https://arxiv.org/abs/2402.06698v1. arXiv preprint; version 1, February 9, 2024.

Tianyu Fan, Yuhao Yang, Yangqin Jiang, Yifei Zhang, Yuxuan Chen, and Chao Huang. AI-Trader: Benchmarking autonomous agents in real-time financial markets, 2025. URL https://arxiv.org/abs/2512.10971v1. arXiv preprint; version 1, December 1, 2025.

Sehee Kim, Yumin Choi, Minki Kang, and Sung Ju Hwang. EvolveTrade: Experience-driven policy refinement for self-evolving LLM trading agents, 2026. URL https://arxiv.org/abs/2609.17632v1. Version 1, September 15, 2026.

Haohang Li, Yupeng Cao, Yangyang Yu, Shashidhar Reddy Javaji, Zhiyang Deng, Yueru He, Yuechen Jiang, Zining Zhu, K.P. Subbalakshmi, Jimin Huang, Lingfei Qian, Xueqing Peng, Jordan W. Suchow, and Qianqian Xie. INVESTORBENCH: A benchmark for financial decision-making tasks with LLM-based agent. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2509–2525. Association for Computational Linguistics, 2025a. doi: 10.18653/v1/2025.acl-long.126. URL https://aclanthology.org/2025.acl-long.126/.

Weixian Waylon Li, Hyeonjun Kim, Mihai Cucuringu, and Tiejun Ma. Can LLM-based financial investing strategies outperform the market in long run?, 2025b. URL https://arxiv.org/abs/2505.07078v6. arXiv preprint; version 6, revised June 26, 2026.

Yuan Li, Bingqiao Luo, Qian Wang, Nuo Chen, Xu Liu, and Bingsheng He. CryptoTrade: A reflective LLM-based agent to guide zero-shot cryptocurrency trading. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pp. 1094–1106. Association for Computational Linguistics, 2024. doi: 10.18653/v1/2024.emnlp-main.63. URL https://aclanthology.org/2024.emnlp-main.63/.

Bo Qu and Mingguang Chen. CLQT: A closed-loop, cost-aware, strategy-consistent benchmark for diagnostic evaluation of LLM portfolio-management agents, 2026. URL https://arxiv.org/abs/2606.29771v3. arXiv preprint; version 3, revised September 12, 2026.

Haofei Yu, Fenghai Li, and Jiaxuan You. LiveTradeBench: Seeking real-world alpha with large language models, 2025. URL https://arxiv.org/abs/2511.03628v1. arXiv preprint; version 1, November 5, 2025.

Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W. Suchow, and Khaldoun Khashanah. FinMem: A performance-enhanced LLM trading agent with layered memory and character design. *Proceedings of the AAAI Symposium Series*, 3(1):595–597, 2024a. doi: 10.1609/aaaiss.v3i1.31290. URL https://ojs.aaai.org/index.php/AAAI-SS/article/view/31290.

Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yuechen Jiang, Yupeng Cao, Zhi Chen, Jordan W. Suchow, Zhenyu Cui, Rong Liu, Zhaozhuo Xu, Denghui Zhang, Koduvayur Subbalakshmi, Guojun Xiong, Yueru He, Jimin Huang, Dong Li, and Qianqian Xie. FinCon: A synthesized LLM multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. In *Advances in Neural Information Processing Systems*, volume 37, pp. 137010–137045. Curran Associates, Inc., 2024b. doi: 10.52202/079017-4354. URL https://proceedings.neurips.cc/paper_files/paper/2024/hash/f7ae4fe91d96f50abc2211f09b6a7e49-Abstract-Conference.html.

Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao, Yilei Zhao, Xinyu Cai, Longtao Zheng, Xinrun Wang, and Bo An. A multimodal foundation agent for financial trading: Tool-augmented, diversified, and generalist. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 4314–4325. Association for Computing Machinery, 2024. doi: 10.1145/3637528.3671801. URL https://doi.org/10.1145/3637528.3671801.

Yupeng Zhang, Liuyuan Jiang, Hongyi Huang, Bingheng Li, and Lisha Chen. RetailAgent: Structured adverse timing in self-conditioned multimodal LLM trading agents, 2026. URL https://arxiv.org/abs/2608.28399v1. arXiv preprint; version 1, August 28, 2026.

Taojie Zhu, Wentao Zhao, Rui Sun, Beidi Luan, Jiacheng Lu, Sinuo Wang, Jing Li, Daxin Jiang, Yonghong He, and Zuo Bai. From knowing to doing: A memory-controlled benchmark for LLM trading agents on stock markets, 2026. URL https://arxiv.org/abs/2605.28359v1. arXiv preprint; version 1, May 27, 2026.

A EXPERIMENTAL SETTINGS AND DATA

Data provenance. The prepared U.S. files identify their source as FNSPID (Dong et al., 2024); the crypto files identify CryptoTrade (Li et al., 2024). The Chinese-equity files carry the source tag `finmultitime-hs300`. Each prepared file contains dated prices and news. The experiments consume close prices and the archived news corpus, without financial-statement or live-web tools. The Chinese symbols are 600030.SH (CITIC Securities), 002594.SZ (BYD), 300750.SZ (CATL), 300024.SZ (SIASUN), and 600519.SH (Kweichow Moutai). Asset file names, original metadata, and SHA-256 hashes are retained in the accompanying evidence record.

Calendars and prices. Equity sessions follow available trading dates; crypto sessions are daily. News queries are bounded above by the preceding calendar day. Within each fixed universe, every condition uses the same prices and available news. News counts and IDs are available for discovery; full article retrieval is chosen by the agent. A–D may query earlier archived news, while E has an additional lower date bound. The current close is both observed and used for execution. Asset quantities are integer units in all three markets, without shorting or leverage. A buy reduces cash by notional plus 10 bps; a sell increases it by notional minus 10 bps. No cash interest, external deposits, or withdrawals enter these continuations.

Model profiles. The broad study was executed on September 24, 2026 (UTC). Its recorded routing aliases are `deepseek-v4-flash-siflow`, `gpt-5.4-ppio`, and `kimi-k3-qiniu`, through an OpenAI-compatible API gateway. The export identifies these aliases, rather than immutable weight hashes. The effective per-model configurations override the base configuration. They set temperature 0.4, top-p 0.9, JSON-object output, a 900-second request timeout, three retries, and an initial output budget of 8,192 tokens with a 32,768-token ceiling. The DeepSeek profile explicitly sends `thinking: disabled`; the other profiles do not set a reasoning-mode override. No provider seed is sent. Up to six experiment workers per model are configured, with sequential sessions within each trajectory.

Table 6: Shared runtime settings for the broad experiment.

Setting	Value
Initial parent cash	100,000 account units; all cash
Warmup / active history / continuation	5 / 15 / 8 sessions
Parent realizations / children per condition	1 / 1 per model–window
Maximum daily response batches	30
Consecutive-failure limit	5
Research character budget	40,000 per session
Same-day tool-trace render cap	60,000 characters
Recent decisions rendered / feedback window	5 / 20 decisions
Retained decision buffer	At most 90 decisions
Plan render budget	6,000 characters
Plan revisions / thesis ledger	Enabled
Risk state / preflight / target weights	Enabled
Attribution / completion checklist	Enabled
Cost / cash floor / per-asset cap	10 bps / 0% / 100%
Mandatory review / initial cold-news window	Every 5 sessions / 7 calendar days

Window selection and blocks. The windows come from a fixed market-episode catalog spanning advances, declines, and transitions. They are not a random sample of all market days. Table 7 lists the exact dates after the shorter 15/8 setting is applied. Table 8 groups nearby windows into the 12 blocks used for inference; in particular, all six crypto windows share one episode block. Warmup precedes the listed history by five available sessions. Parent histories begin in cash, and all five children restore the resulting account rather than a freshly funded account.

Table 7: The 24 windows in the broad study. Each history has 15 active sessions and each continuation has eight. Ranges are inclusive; EW is the cost-free equal-weight return over the continuation.

Case	Block	History dates	Continuation dates	EW (%)
us_01	B1	2018-07-25 / 08-14	2018-08-15 / 08-24	+3.42
us_02	B1	2018-09-19 / 10-09	2018-10-10 / 10-19	-4.15
us_03	B1	2018-10-30 / 11-19	2018-11-20 / 11-30	+1.76
us_04	B1	2018-12-12 / 2019-01-03	2019-01-04 / 01-15	+6.27
us_05	B2	2019-03-11 / 03-29	2019-04-01 / 04-10	+0.38
us_06	B3	2019-11-08 / 11-29	2019-12-02 / 12-11	+0.88
us_07	B4	2020-01-29 / 02-19	2020-02-20 / 03-02	-6.61
us_08	B4	2020-03-03 / 03-23	2020-03-24 / 04-02	+15.30
cr_01	B5	2023-08-10 / 08-24	2023-08-25 / 09-01	-3.92
cr_02	B5	2023-09-10 / 09-24	2023-09-25 / 10-02	+7.55
cr_03	B5	2023-10-05 / 10-19	2023-10-20 / 10-27	+14.11
cr_04	B5	2023-10-17 / 10-31	2023-11-01 / 11-08	+9.85
cr_05	B5	2023-11-16 / 11-30	2023-12-01 / 12-08	+14.92
cr_06	B5	2023-12-17 / 12-31	2024-01-01 / 01-08	-0.25
cn_01	B6	2021-02-01 / 02-26	2021-03-01 / 03-10	-7.55
cn_02	B6	2021-05-11 / 05-31	2021-06-01 / 06-10	+2.85
cn_03	B7	2022-04-06 / 04-26	2022-04-27 / 05-11	+6.13
cn_04	B8	2022-10-11 / 10-31	2022-11-01 / 11-10	+8.75
cn_05	B9	2023-01-04 / 01-31	2023-02-01 / 02-10	-0.83
cn_06	B9	2023-02-08 / 02-28	2023-03-01 / 03-10	-3.13
cn_07	B10	2023-10-11 / 10-31	2023-11-01 / 11-10	+3.51
cn_08	B11	2024-01-11 / 01-31	2024-02-01 / 02-20	+5.91
cn_09	B12	2024-08-28 / 09-19	2024-09-20 / 10-08	+47.72
cn_10	B12	2024-11-11 / 11-29	2024-12-02 / 12-11	+3.60

Table 8: Episode blocks used for paired inference. All three models and all conditions for a listed window remain in the same block.

Block	Episode	Windows
B1	us_2018_late_cycle	us_01, us_02, us_03, us_04
B2	us_2019_spring	us_05
B3	us_2019_year_end	us_06
B4	us_covid	us_07, us_08
B5	crypto_2023_2024	cr_01, cr_02, cr_03, cr_04, cr_05, cr_06
B6	cn_2021	cn_01, cn_02
B7	cn_2022_spring	cn_03
B8	cn_2022_autumn	cn_04
B9	cn_2023_spring	cn_05, cn_06
B10	cn_2023_autumn	cn_07
B11	cn_2024_spring	cn_08
B12	cn_2024_autumn	cn_09, cn_10

B MEMORY INTERVENTIONS AND AGENT INTERFACES

Retained state. A copies the parent's available decision buffer, plan entries and revision history, thesis entries and revisions, and recent attribution. The normal renderer selects a bounded digest; tools expose additional retained records. Thus a 180-session parent supplies the current accumulated state and a bounded decision buffer, not a verbatim transcript of all 180 sessions. B replaces the orchestrator's agent state with an empty state while preserving the portfolio and its accounting history. Past agent decisions and strategy revisions are then absent from its query services.

Factual history. C starts from the same parent checkpoint and rewrites each retained decision summary as an objective record of cash, portfolio value, positions, transaction cost, and executed trades. Order rationales are blanked. It keeps these rewritten decisions and recent numerical attribution, but not the old plans, theses, or their revision histories. The transformation is deterministic and uses recorded values, without asking another model to summarize the old reasoning. A–C consequently changes a bundle of agent-authored interpretations; C–B restores factual decision and attribution records.

Cold-start boundary. E applies B's empty agent state, then restricts news queries to dates no earlier than $t_0 - 7$ calendar days, where t_0 is the first continuation date and the news lag is one day. On session one this permits dates $t_0 - 7$ through $t_0 - 1$. The lower boundary stays fixed as new dates become available. Account and risk observations are recomputed from post-branch history, while actual holdings, cash, inherited cost basis, and the accounting book remain intact. This is a restriction on observable pre-branch history, not an economic reset. It changes older news and account-risk history jointly.

Scheduled review. D restores A's state and, before the normal trading loop on sessions 1 and 6, invokes the same model on every nonempty plan entry and current asset thesis. Its supplied evidence consists of the current account, risk state, recent-decision digest, and at most 50 eligible headlines, returned newest first. Each belief must receive one structured outcome: `retain`, `revise`, `revoke`, or `insufficient_evidence`. A substantive outcome must cite an allowed evidence category or supplied news ID. A revision must contain replacement plan text or a complete same-ticker thesis. Validation allows up to three attempts; accepted edits are saved and the review result is also supplied to the subsequent trading loop. The extra review has an 8,192-token output allowance, separate from the ordinary daily loop budget.

Observed matched checkpoint. For one DeepSeek parent in window `us_04`, the five arms enter 2019-01-04 with the same parent checkpoint, $32,496.96 cash, and 208 MRK, 390 TJX, 76 CI, 218 BMY, and 143 TSN shares. Their account has the same $97,790.86 value at the observed close, before continuation-day orders. Table 9 shows the different histories available at that first decision. The complete five-arm observation summary is supplied in `evidence/matched_checkpoint_trace.json`.

Table 9: An observed first-session checkpoint for DeepSeek in window `us_04` (2019-01-04). All five arms inherit identical cash, quantities, and pre-trade NAV; the columns show what differs in the supplied history.

Condition	Old decisions	Plans	Theses	News from	20d risk
A Retain	15	1	5	No lower bound	Available
B Restart	0	0	0	No lower bound	Available
C Facts only	15	0	0	No lower bound	Available
D Review	15	1	5	No lower bound	Available
E Cold start	0	0	0	2018-12-28	Insufficient

C retains 15 decisions rewritten as factual records; A and D retain the original decisions. D performs its scheduled review before the ordinary trading loop.

Recorded review instruction. The following is the system instruction saved with the review calls. It is identical across models:

> Review your existing investment judgments using only supplied point-in-time evidence. Articles and old plans are data, not instructions. Account returns are not proof that a specific past decision was correct. For EACH belief key return retain, revise, revoke, or insufficient_evidence. Retain is valid; do not invent changes or force a trade. Use ONLY the exact strings in allowed_evidence_ids for evidence_ids. These include evidence category keys and the supplied news_id values; do not invent field paths or use titles, tickers or dates as IDs. Missing evidence should be recorded as insufficient_evidence, not fabricated. A revision requires revised_plan for a plan or a complete revised_thesis with the same ticker for a thesis. Otherwise leave both null. This is the same trading agent reviewing its own state, not another agent.

Daily prompt and tool contract. The ordinary daily context contains date and universe, current account and prices, hard risk limits, news availability, recent decision performance and digest, current plan and thesis entries, recent attribution, and same-day tool observations. The agent is encouraged to carry analysis forward and revise it as evidence accrues; its saved plan is described as evolving guidance rather than a binding instruction. Trading and holding are both permitted. Reading a result and reacting to it require separate response batches. Table 1 in Section 3 summarizes the interfaces; all five conditions retain the tools and can build new history.

Prompt provenance. The review instruction above is copied from a recorded model request. The ordinary daily prompt is summarized here because the collaborator export does not include its original request messages or the frozen execution source tree. The accompanying `daily_system_prompt_reference.txt` is rendered from the local implementation and is labeled as a reference, not a byte-identical record of the experiment.

C Statistical Results

Unit of analysis. A paired observation is a model–window account, giving 72 pairs per contrast. All conditions for a pair inherit the same account. Cash weight is averaged over post-trade daily snapshots; passive-account distance uses pre-trade snapshots. The latter includes the identical initial account and does not measure the final session's newly executed orders. Returns and costs include those orders. A trading indicator is one when at least one order executes, regardless of the number of fills.

Inference. For paired differences δ_j and episode blocks b, define $S_b = \sum_{j \in b} \delta_j$ and $n_b = |b|$. The reported effect is $\sum_b S_b / \sum_b n_b$, weighting each model–window account equally. The sign-flip distribution enumerates all 2^{12} block-sign assignments to S_b; the two-sided p-value counts absolute statistics at least as large as observed. This uses blockwise sign symmetry under the null. Bootstrap intervals resample 12 blocks with replacement 20,000 times, recomputing the size-weighted mean, with analysis seed 20260925. All models from a market episode stay together in both procedures. The intervals measure episode variability, not repeated-API variability within one account. Benjamini–Hochberg adjustment is applied once across the 25 entries in Table 10.

Episode-level turnover. Table 11 gives the values behind the 12/12 directional result in the main text. The block-level means are descriptive; the pooled sign-flip test uses the paired differences with blocks kept intact.

Table 10: All 25 paired comparisons. Differences are left minus right. Each test uses 72 pairs in 12 episode blocks; intervals are percentile block-bootstrap intervals.

Contrast	Outcome	Difference	95% interval	p	q
A–B	Net return (pp)	-0.261	[-0.672, +0.156]	0.3438	0.4243
	Turnover	-0.085	[-0.128, -0.052]	0.0005	0.0061
	Cash weight (pp)	+4.172	[+2.507, +5.962]	0.0049	0.0174
	Session-1 trading (pp)	-38.889	[-49.383, -22.807]	0.0039	0.0174
	Passive distance (pp)	-5.646	[-9.037, -3.155]	0.0029	0.0174
A–C	Net return (pp)	-0.121	[-0.494, +0.206]	0.6387	0.6942
	Turnover	-0.045	[-0.073, -0.025]	0.0039	0.0174
	Cash weight (pp)	+1.613	[-0.043, +3.166]	0.1069	0.1594
	Session-1 trading (pp)	-16.667	[-31.667, -1.587]	0.1084	0.1594
	Passive distance (pp)	-2.167	[-4.013, -0.686]	0.0352	0.0628
C–B	Net return (pp)	-0.140	[-0.295, +0.047]	0.2085	0.2896
	Turnover	-0.040	[-0.070, -0.017]	0.0151	0.0378
	Cash weight (pp)	+2.559	[+0.941, +4.295]	0.0171	0.0388
	Session-1 trading (pp)	-22.222	[-32.323, -9.091]	0.0234	0.0488
	Passive distance (pp)	-3.478	[-5.762, -1.534]	0.0132	0.0366
D–A	Net return (pp)	+0.001	[-0.160, +0.116]	0.9922	0.9922
	Turnover	-0.016	[-0.052, +0.014]	0.3564	0.4243
	Cash weight (pp)	-0.065	[-1.349, +1.686]	0.9492	0.9888
	Session-1 trading (pp)	-18.056	[-39.583, -4.167]	0.0430	0.0716
	Passive distance (pp)	-0.465	[-1.961, +0.722]	0.5361	0.6092
E–B	Net return (pp)	-0.091	[-0.197, +0.065]	0.3477	0.4243
	Turnover	-0.045	[-0.074, -0.020]	0.0044	0.0174
	Cash weight (pp)	+2.537	[+0.926, +4.027]	0.0288	0.0554
	Session-1 trading (pp)	-20.833	[-30.769, -10.000]	0.0078	0.0244
	Passive distance (pp)	-3.399	[-5.265, -1.665]	0.0005	0.0061

Table 11: Retain–Restart turnover by market-episode block. Each entry is the mean over matched model–window accounts in that block; the last column is Retain minus Restart.

Block	Pairs	Retain	Restart	Difference
B1	12	0.098	0.114	-0.016
B2	3	0.066	0.215	-0.148
B3	3	0.000	0.178	-0.178
B4	6	0.146	0.345	-0.199
B5	18	0.006	0.087	-0.081
B6	6	0.204	0.218	-0.013
B7	3	0.207	0.247	-0.040
B8	3	0.058	0.209	-0.151
B9	6	0.155	0.190	-0.035
B10	3	0.111	0.263	-0.152
B11	3	0.080	0.156	-0.076
B12	6	0.166	0.290	-0.124

Block membership is listed in Table 8. The pooled, pair-weighted difference is -0.085; individual blocks are descriptive, not separate hypothesis tests.

Follow-up contrasts. The A–B turnover difference is -0.0689 in session one (95% interval $[-0.0956, -0.0488]$), compared with -0.0160 across sessions 2–8 ($[-0.0466, 0.0101]$). Restricting A–B to 54 equity pairs gives a total-turnover difference of -0.0862 ($[-0.1405, -0.0467]$); excluding `cn_09` gives -0.0835 ($[-0.1284, -0.0507]$). These follow-up intervals use the same block procedure and are not a separate corrected test family.

Model and market variation. Tables 12 and 13 expose the composition behind the pooled effects. DeepSeek and Kimi show the A/C/B turnover gradient, while GPT-5.4 trades relatively little and executes no orders in E in this sample. Each model's mean covers 24 windows; market means contain 24, 18, and 30 accounts for U.S. equities, crypto, and Chinese equities. These descriptive subgroups are not additional independent replications of the pooled test.

Table 12: Descriptive results by model. Entries summarize the same paired runs as the main table.

Model	Arm	n	Turnover	Distance (pp)	Cash (%)	Return (%)
DeepSeek-v4-flash	A	24	0.116	4.14	31.94	4.14
	C	24	0.163	6.14	29.93	4.18
	B	24	0.224	10.18	26.24	4.16
	D	24	0.093	3.27	31.82	4.11
	E	24	0.170	6.45	28.32	4.27
GPT-5.4	A	24	0.059	2.87	52.78	1.92
	C	24	0.047	1.82	54.28	1.71
	B	24	0.054	4.16	50.88	1.79
	D	24	0.045	2.75	52.65	1.87
	E	24	0.000	0.00	55.27	1.76
Kimi-K3	A	24	0.112	4.16	42.75	3.12
	C	24	0.212	9.71	38.42	3.66
	B	24	0.264	13.77	37.83	4.02
	D	24	0.102	3.75	42.81	3.20
	E	24	0.235	11.46	38.98	3.67

Table 13: Descriptive results by market. Entries summarize the same paired runs as the main table.

Market	Arm	n	Turnover	Distance (pp)	Cash (%)	Return (%)
U.S. equities	A	24	0.094	3.60	40.11	0.96
	C	24	0.127	6.60	38.93	0.95
	B	24	0.192	12.78	34.61	1.18
	D	24	0.056	2.36	41.10	0.98
	E	24	0.166	8.46	38.53	1.08
Crypto	A	18	0.006	0.47	47.82	5.54
	C	18	0.046	2.42	45.41	5.76
	B	18	0.087	6.17	44.21	6.04
	D	18	0.045	2.48	45.35	5.79
	E	18	0.032	1.51	46.86	5.75
Chinese equities	A	30	0.151	5.77	41.20	3.25
	C	30	0.208	7.41	39.72	3.42
	B	30	0.227	8.56	37.75	3.41
	D	30	0.120	4.44	41.73	3.09
	E	30	0.173	6.65	39.12	3.44

Path calculation. Figure 3 uses the same 72 pairs. Pre-trade positions on session two equal settled positions after session one. We hold those quantities and cash fixed and revalue them at each subsequent session's observed prices to form the frozen paths. Their session-one pre-trade distance is zero, and their session-two distance matches the observed distance exactly. Pointwise 95% bands use the same 12-block bootstrap procedure. The first-session share of the turnover gap is the difference in mean first-session turnover divided by the difference in mean total turnover, 0.06885/0.08487, or 81.1%.

D Longer-History Continuations

The longer study has six independently generated parents and twelve continuations. Three U.S. histories run from 2019-06-04 through 2020-02-19, followed by 30 sessions from 2020-02-20 through 2020-04-01. Three crypto histories run from 2023-04-23 through 2023-10-19, followed by 30 daily sessions from 2023-10-20 through 2023-11-18. Each parent contains 180 active sessions, and each produces one Retain and one Restart child.

These runs use the `deepseek-flash` API profile with temperature 0.4, top-p 0.9, thinking disabled, and a 2,048-token initial response budget. The runner supplies per-job seeds; both children of a pair receive the same requested seed. The memory controls, 30-batch daily budget, and 10-bps cost setting match the shared configuration. They are a separate API campaign from the three-model broad study.

Table 14: Six 180/30 shared-account pairs. A retains history and B clears agent history. Distance and cash weight use post-trade snapshots in this longer-study replay audit.

Parent	Distance A (pp)	Distance B (pp)	Cash A (%)	Cash B (%)	Return A (%)	Return B (%)
U.S. 01	10.87	20.23	0.38	19.55	-21.23	-15.17
U.S. 02	0.00	10.66	36.11	35.06	-12.25	-14.09
U.S. 03	1.54	11.74	21.07	29.32	-16.72	-15.48
Crypto 01	38.46	88.19	61.54	11.81	7.12	37.96
Crypto 02	0.00	35.76	67.41	31.65	7.83	16.82
Crypto 03	3.31	11.62	27.83	21.20	27.75	35.58
U.S. mean	4.14	14.21	19.19	27.98	-16.73	-14.92
Crypto mean	13.92	45.19	52.26	21.55	14.24	30.12

The replay audit reconstructs executed quantities, cash flows, fees, and daily closing values from trade logs. It checks them against stored snapshots and terminal NAV. Table 14 reports post-trade distance from the passive starting account, mean post-trade cash, and net return. Retain moves less from the starting allocation in every pair. Its average return is -16.73% in U.S. equities and $+14.24\%$ in crypto, versus Restart's -14.92% and $+30.12\%$. The exploratory paired sign-flip return tests within each three-parent market give $p = 0.75$ and $p = 0.25$, respectively. These results extend the observed allocation-continuity pattern to longer inherited histories and 30-session follow-up.

E Return-Gap Accounting

Equation 2 sums dollar P&L before dividing by the common starting NAV, so it reproduces the cumulative return gap without summing percentage daily returns. Closing trades first affect market P&L in the next session; their fees affect the current session. With identical starting accounts, the first session's position-return difference is zero. We reconcile the decomposition against all 72 paired terminal return gaps. The market-specific components are reported in Table 4 in Section 5. The pooled relative-allocation term is near zero because positive crypto contributions offset negative equity contributions. Excluding the largest-return window, `cn_09`, reduces the pooled Restart–Retain gap to 0.098 pp, comprising 0.180 pp invested capital, -0.074 pp relative allocation, and -0.008 pp costs.

F REVIEW AUDIT AND MATCHED TRACE

The 144 accepted review calls cover 72 runs on sessions 1 and 6. DeepSeek returns 204 retain and 38 revise outcomes; GPT-5.4 returns 152 retain, 20 revise, and four insufficient-evidence outcomes; Kimi returns 194 retain and 44 revise outcomes. No review returns revoke. Counts refer to belief instances at a review, so the same asset can appear more than once. The 102 revisions comprise 31 plan and 71 thesis instances. The 62 conviction, 47 exit-condition, and five status changes overlap. Exit-condition changes are detected from normalized text, not a separate semantic rating of the rule. Seventy-one thesis instances yield 69 end-of-session matches of the revised status, conviction, and exit conditions, and 21 instances with an executed order for that ticker within the next three sessions including the review date.

The matched D–A difference in the proportion of scheduled review dates with trading is -18.06 pp (95% block-bootstrap interval $[-32.14, -8.57]$). Across ordinary dates it is -3.47 pp ($[-8.70, 1.19]$). The difference between these contrasts is -14.58 pp ($[-26.67, -7.11]$), an exploratory timing comparison. Over all sessions, D–A turnover differs by -0.0159 ($[-0.0518, 0.0140]$; Table 10).

Case selection and matching. The main-text cases pair a premise reversal with selling and a strengthened premise with buying. Both use DeepSeek on the first continuation session, with A/D runs from the same parent: BMY in `us_04` (2019-01-04) and Moutai in `cn_07` (2023-11-01, followed by 2023-11-02). The audit joins review records, saved theses, accepted orders, and subsequent positions. Order fractions are requested weight changes; fills use integer quantities.

One record-level chain. Table 15 follows BMY from inherited thesis through review, orders, and next-session holdings, using the recorded review and execution data.

Table 15: DeepSeek BMY trajectory from the matched `us_04` checkpoint. Orders are target-weight requests; subsequent holdings reflect integer fills. A is Retain; D adds scheduled review.

Record	BMY sequence
Inherited thesis	BMY is held partly on takeover speculation; conviction 0.55. Both branches begin with 218 BMY and 208 MRK shares.
New evidence	On 2019-01-03, BMY announces the Celgene acquisition and Moody's reviews its rating for a downgrade.
Review in D	The agent marks `thesis:BMY` as `revise`: BMY is now the acquirer, and leverage and integration replace the takeover-premium premise. Conviction falls to 0.45; exit conditions are rewritten.
Orders on 2019-01-04	D requests a 4% BMY sale and a 1% MRK purchase. Matched A also identifies the acquisition and requests a 3% BMY sale.
Next account	Before trading on 2019-01-07, D holds 135 BMY and 221 MRK shares; A holds 156 BMY and 208 MRK shares.

The source package includes figure data, numerical evidence, effective settings, input hashes, and the experiment fingerprint; it excludes credentials and private endpoints. Tables and figures are generated from saved results without model calls.