

Summoning the Oracle to Slay It: Mitigating Look-Ahead Bias in Financial Backtesting with Large Language Models

Weixian Waylon Li
University of Edinburgh
waylon.li@ed.ac.uk

Mengyu Wang
University of Edinburgh
mengyu.wang@ed.ac.uk

Tiejun Ma
University of Edinburgh
tiejun.ma@ed.ac.uk

Abstract

Backtesting large language models (LLMs) on historical financial data is unreliable when their pre-training data include the evaluated events. An LLM trained in 2024 may already encode how stocks moved during 2018–2020. We name this failure *parametric look-ahead bias* and propose FINCAD, an inference-time adaptation of Context-Aware Decoding that attenuates contributions from memorised historical outcomes without retraining. FINCAD pairs an adversarial bias-discovery pipeline that learns a model-specific memory-activating prior prompt with an entity- and date-adaptive rule that scales the CAD strength using a per-(entity, date) confidence signal. Across five 7–14B LLMs and five mega-cap equities, the largest model-level mean in-sample return correction is -67.1% . For the three larger models, 2025 out-of-sample returns remain within \$8K and mean Sharpe within ± 0.10 of baseline; mean general-benchmark accuracy remains positive or within -1.7 points for four of five models. On an eleven-model leaderboard, FINCAD raises the subset-averaged in-sample/out-of-sample Spearman correlation from $+0.779$ to $+0.846$, yielding rankings that are more closely aligned with post-cutoff performance.

1 Introduction

Large language models (LLMs) are increasingly applied as autonomous financial agents that read news, filings, and prices and decide whether to buy, sell, or hold (Saha et al., 2025). In quantitative finance, backtests gate capital allocation, so a result inflated by undetected bias translates directly into capital misallocation and live-trading losses; standardised corrections exist for backtest overfitting and multiple testing (Harvey and Liu, 2014; Bailey et al., 2014; Bailey and López de Prado, 2014; López de Prado, 2018), survivorship truncation (Heckman, 1979; Brown et al., 1992), and point-in-time data leakage (Li et al., 2026a;

Kong et al., 2026). However, limited work has studied the bias introduced by LLMs. A state-of-the-art LLM with a 2025 cutoff has already seen how NVIDIA, Microsoft, and Netflix moved during 2010–2020, so any backtest overlapping the training window inherits an additional bias that lives *inside* the model’s weights and is invisible to data-pipeline audits. We call this failure mode *parametric look-ahead bias*, and its consequence is a trap: uncorrected long-horizon backtests are systematically misleading, whereas the only bias-free alternative, the short post-cutoff period, is too limited to support statistically reliable conclusions (Li et al., 2026a).

A natural starting point is Context-Aware Decoding (CAD) (Shi et al., 2024), which sharpens the contribution of the input by subtracting an off-context prior at the logit level. CAD was originally proposed to suppress *unfaithful* priors (hallucinations); ours is the opposite regime, where the prior the model carries about “NVDA in 2018” is not a hallucination but an accurate memory of the future that backtesting must explicitly reject. Two challenges follow: x_{prior} must reliably activate that memory rather than answer the question, and α must vary across entities and dates since memorisation is non-uniform on both axes, so a fixed α would simultaneously under-correct memorised cases and over-penalise context-grounded reasoning on out-of-sample dates.

We propose FINCAD¹, an inference-time decoding scheme that mitigates parametric look-ahead bias without retraining by explicitly eliciting a memory-enriched prior and subtracting it from the context-conditioned logits with an adaptive strength. Conceptually, the discovery stage “summons the oracle” by making the historical prior explicit; the title’s “slaying” refers to attenuating that prior’s contribution during decoding. An *adver-*

¹Code: <https://github.com/waylonli/FinCAD>

serial bias-discovery pipeline optimises a model-specific instruction T_{prior}^* that activates the model’s historical-outcome recall on a held-out probe set, yielding an explicit subtractable prior. An *entity- and date-adaptive* rule for α compares the model’s per-(entity, date) directional entropy against an entity-level baseline, so the CAD strength rises with date-variance and per-date excess confidence and attenuates when this signal is absent.

Our contributions are:

- We propose FINCAD, the first decoding-time adaptation of CAD designed for mitigating parametric look-ahead bias, with an adversarial discovery pipeline for x_{prior} and an entity- and date-adaptive rule for α (§3.1, §3.2).
- Across five LLMs (7B–14B), five mega-cap equities, and five reasoning benchmarks, FINCAD produces model-dependent in-sample (2010–2020) performance corrections while mean benchmark accuracy remains positive or within -1.7 points for four of five models (§4.1, §4.2).
- FINCAD improves the descriptive alignment between in-sample and post-cutoff model rankings: across eleven LLMs on the S&P 500, the subset-averaged Spearman correlation between in-sample (2010–2020) and 2025 Sharpe rankings rises from $+0.779$ to $+0.846$ (§4.3).

2 Related Work

A growing line of work casts LLMs as autonomous financial investors (Guo et al., 2024; Saha et al., 2025), including a recent wave of dedicated trading agents (Yu et al., 2023; Zhang et al., 2024; Yang et al., 2024; Ding et al., 2024; Yu et al., 2024; Wang and Ma, 2024; Wang et al., 2024a, 2025; Xiao et al., 2025; Li and Ma, 2025; Fatouros et al., 2025; Li et al., 2025, 2026c) and reinforcement-learning-enhanced variants (Ding et al., 2023; Koa et al., 2024; Xiao et al., 2025; Deng et al., 2026), most of which evaluate on short historical windows without bias control, leaving open whether reported gains reflect genuine reasoning or contamination of the evaluation regime.

Input-side protections (survivorship, data-snooping, data-leakage look-ahead) are tracked by FINSABER (Li et al., 2026a), but none addresses the bias that resides in the model’s weights themselves. This *parametric look-ahead* is documented

by Benhenda (2026), who shows that LLMs recall in-cutoff S&P 500 prices verbatim, and by Gao et al. (2026), who provide a complementary statistical test; concurrently, Merchant and Levy (2025) propose a decoding-time mitigation that fine-tunes a pair of auxiliary models on labelled forget/retain partitions, whereas FINCAD uses a single base model with no auxiliary fine-tuning.

Our work connects to the broader literature on LLM memorisation (Carlini et al., 2021, 2023; Li et al., 2026b) and *machine unlearning* for LLMs (Liu et al., 2024), including parameter-editing methods (Meng et al., 2022, 2023) and fine-tuning-based forget protocols (Jang et al., 2023; Eldan and Russinovich, 2023; Maini et al., 2024). All require curated forget sets, retraining cycles, or structured (subject, relation, object) edits, and risk collateral damage to neighbouring knowledge; they are operationally impractical at the eleven-model leaderboard scale of financial backtests, where each model would need a bespoke unlearning pass per ticker per evaluation date. Decoding-side interventions instead operate at inference without modifying weights (Li et al., 2023; Chuang et al., 2024), and FINCAD belongs to this family. We build specifically on Context-Aware Decoding (Shi et al., 2024), whose contrast between with-context and without-context branches matches our setting: a trading agent’s memory of how a stock actually moved should not leak into its reading of the financial context provided at decision time.

3 Methodology

Let $\mathcal{M}$ denote a language model with vocabulary $\mathcal{V}$. For a given input string x , $\mathcal{M}(x) \in \mathbb{R}^{|\mathcal{V}|}$ denotes the model’s next-token logit vector, and we write $\tilde{\mathbf{y}} = \text{softmax}(\mathcal{M}(x))$ for the corresponding next-token output distribution over $\mathcal{V}$. Context-Aware Decoding (Shi et al., 2024) replaces this ordinary output distribution with a contrastively adjusted distribution. In our FINCAD, the adjustment strength varies per (entity, date) pair (s, t) , where s is the stock symbol and t the decision date:

$$\tilde{\mathbf{y}}_{s,t} = \text{softmax}\left((1 + \alpha(s, t)) \mathcal{M}(x_{\text{ctx}}) - \alpha(s, t) \mathcal{M}(x_{\text{prior}})\right) \quad (1)$$

where $\tilde{\mathbf{y}}_{s,t}$ denotes the FINCAD-adjusted next-token output distribution for entity s at date t . The original CAD targets unfaithful hallucinations and treats the prior as noise to be discounted. We repurpose it for the opposite regime, where a memory-




Figure 1: FINCAD pipeline: (1) MIPROv2 discovers a memory-activation instruction T_{prior}^* on $\mathcal{D}_{\text{calib}}$; (2) a completion probe and entity calibration set a per-(entity, date) strength $\alpha(s, t)$; (3) CAD subtracts the memory-enriched prior logits with strength $\alpha(s, t)$ from the context-conditioned logits.

enriched prior about historical outcomes is actively elicited and subtracted. This leads to two design choices. x_{prior} must maximally activate that memory rather than answer the question; and $\alpha(s, t)$ should vary across both entities and time. We address these in §3.1 and §3.2 respectively; Figure 1 summarises the complete procedure, with full pseudocode in Algorithm 1 (Appendix C).

3.1 Adversarial Bias Discovery

The prior prompt x_{prior} determines what knowledge is subtracted during decoding; if it fails to activate the model’s memorised future knowledge, the subtraction has no effect and look-ahead bias persists, so we need a prior that reliably triggers parametric recall of historical outcomes.

For a decision over entity s at date t , the two branches of Equation 1 share structural fields but differ in instruction and content. The context branch carries the full agent input:

$$x_{\text{ctx}} = T_{\text{ctx}} \oplus (s, t) \oplus F_{\text{task}} \oplus D(s, t) \quad (2)$$

where $\oplus$ denotes string concatenation, T_{ctx} is the system prompt framing the model as a quantitative analyst, F_{task} is the task-specific instruction (e.g., “Analyse the following financial data and decide whether to buy, sell, or hold”) together with the output-format specification (“Respond with valid JSON”), and $D(s, t)$ is the financial context the agent reasons over. To realise the shared substructure with x_{prior} at decoding time, F_{task} is extracted and copied verbatim into the prior input below. The prior branch strips $D(s, t)$ and replaces T_{ctx} with an

adversarially optimised memory-activation instruction T_{prior}^* designed to activate parametric recall:

$$x_{\text{prior}} = T_{\text{prior}}^* \oplus (s, t) \oplus F_{\text{task}} \quad (3)$$

T_{prior}^* is optimised once per model. Sharing the (s, t) and F_{task} fields across the two branches keeps both inputs over the same structural tokens, so the logit subtraction in Equation 1 targets the *content* of the response rather than its format or task framing. During autoregressive decoding, each sampled token is appended to both branches; every subsequent step therefore contrasts the same candidate tokens conditioned on a shared generated prefix.

Optimisation objective. We frame the search for T_{prior}^* as a discrete prompt optimisation problem over a calibration dataset $\mathcal{D}_{\text{calib}}$ of historical (entity, date, outcome) triples. Each example presents the model with an entity s_i and date t_i and asks for the direction label $d_i \in \{\text{up}, \text{down}\}$ from memory alone, with no financial context. The objective is:

$$\hat{d}_i = \arg \max_d \mathcal{M}(T_{\text{prior}} \oplus (s_i, t_i) \oplus F_{\text{task}}^{\text{calib}})_d \quad (4)$$

$$T_{\text{prior}}^* = \arg \max_{T_{\text{prior}}} \sum_i \delta(\hat{d}_i, d_i) \quad (5)$$

where d is a direction label drawn from $\{\text{up}, \text{down}\}$; for each label d , $\mathcal{V}_d \subset \mathcal{V}$ collects its tokenizer variants (casing and leading-space differences), and we write $\mathcal{M}(X)_d := \max_{v \in \mathcal{V}_d} \mathcal{M}(X)_v$ as shorthand for the maximum logit over those variants. The sum runs over all triples $(s_i, t_i, d_i) \in \mathcal{D}_{\text{calib}}$, T_{prior} is the candidate memory-activation instruction being optimised, $F_{\text{task}}^{\text{calib}}$ is a fixed calibration

task-shell that specifies the permitted output labels but contains *no* prediction instruction, and $\delta(\hat{d}_i, d_i)$ equals 1 when the predicted direction matches the ground-truth label and 0 otherwise. Crucially, only T_{prior} is optimised; $F_{\text{task}}^{\text{calib}}$ remains fixed, and the optimised T_{prior}^* is later reused at inference with a different (backtest) F_{task} via Eq. 3.

Calibration dataset. We construct $\mathcal{D}_{\text{calib}}$ from historical S&P 500 price data covering the period 2005–2015, chosen to fall well within the pre-training data of modern LLMs. For each ticker s , we sample dates at quarter-end frequency and compute the forward return over the subsequent 63 trading days (approximately one calendar quarter): $r_{s,t} = (p_{s,t+63} - p_{s,t}) / (p_{s,t})$, where $p_{s,t}$ is the split- and dividend-adjusted closing price. We discard examples where $|r_{s,t}| < 0.05$, since small returns yield ambiguous direction labels. The remaining examples are labelled $d_{s,t} = \text{up}$ if $r_{s,t} > 0$ and down otherwise. We balance the two classes by downsampling the majority and cap the dataset at 200 examples, split 80/20 into training and validation.

Optimiser. We solve the optimisation in Equation 5 using MIPROv2 (Opsahl-Ong et al., 2024), a multi-stage instruction optimisation algorithm from the DSPy framework (Khatab et al., 2024). MIPROv2 generates a pool of candidate instructions by prompting a proposer model, evaluates each candidate’s accuracy on the training set, and selects the best-performing candidate on the held-out validation set. We initialise the search with the seed instruction: “Recall what you know from your training data about this financial entity and predict the outcome.”

To keep T_{prior}^* task-agnostic, we hide task-specific information from MIPROv2’s proposer: training data (ticker names, dates, labels) and bootstrap traces are withheld, so the proposer sees only a generic program skeleton and the current seed. This prevents the proposer from leaking task-specific output formats (e.g., “predict whether the stock went up or down”) into T_{prior}^* . The evaluation metric, however, uses the full calibration input (including the output labels in $F_{\text{task}}^{\text{calib}}$), so candidates are still scored on their ability to surface memorised knowledge. The optimised T_{prior}^* is serialised as a JSON artefact and loaded at inference time; the per-model T_{prior}^* values appear in Appendix B.3.

3.2 Entity-Adaptive Penalty

A fixed penalty α applies the same subtraction strength to every entity and time period. This is problematic for two reasons. First, parametric look-ahead bias is non-uniform across entities: the model has strong parametric recall for well-known stocks (e.g., AAPL, NVDA) but near-random predictions for obscure or recently listed companies. Applying a high α to the latter subtracts noise from the context-conditioned logits, degrading output quality. Second, memorisation is non-uniform across time: the model has detailed knowledge of historical periods covered by its training data but no memorised prices for dates beyond its knowledge cutoff. A date-agnostic probe would assign the same high α to a well-known entity regardless of whether the backtest date falls inside or outside the training window, penalising the prior even when it carries no look-ahead information. We therefore set α dynamically per entity *and* date, based on how confident the model is in its own memory.

Completion-based probing. Given entity s and date t , we construct a completion probe by prepending the optimised instruction T_{prior}^* to a natural-language prefix that elicits a directional continuation: $x_{\text{probe}} = T_{\text{prior}}^* \oplus$ “After t , s stock went”.

We run a single forward pass and read the per-label logits $\mathcal{M}(x_{\text{probe}})_d$ (the max-over-variants shorthand introduced in §3.1):

$$\ell_{\uparrow} = \mathcal{M}(x_{\text{probe}})_{\text{up}}, \quad \ell_{\downarrow} = \mathcal{M}(x_{\text{probe}})_{\text{down}} \quad (6)$$

These directional logits determine only the scalar gate $\alpha(s, t)$; they are not mapped to the agent’s buy/sell/hold actions. Action generation remains the same-token, full-vocabulary contrast in Equation 1.

Date-variance scaling. We convert the extracted logits into a normalised binary entropy $\hat{H} \in [0, 1]$:

$$p_{\uparrow} = \frac{\exp(\ell_{\uparrow})}{\exp(\ell_{\uparrow}) + \exp(\ell_{\downarrow})}, \quad p_{\downarrow} = 1 - p_{\uparrow} \quad (7)$$

$$\hat{H} = -\frac{1}{\log 2} (p_{\uparrow} \log p_{\uparrow} + p_{\downarrow} \log p_{\downarrow}) \quad (8)$$

where $\hat{H} = 0$ indicates full confidence and $\hat{H} = 1$ indicates maximum uncertainty.

A naïve approach would threshold on raw entropy: penalise when $\hat{H}$ is low and abstain when it is high. However, raw entropy conflates two distinct sources of confidence. The model may be

confident because it has *memorised the specific outcome* for entity s at date t (temporal memorisation, the target of bias removal), or because it holds a *persistent belief* about s regardless of t (a brand prior, e.g., “NVIDIA is a strong company”). Subtracting a brand prior removes legitimate reasoning without addressing look-ahead bias.

We disentangle these two signals through a **date-variance analysis**. Before each backtest, we probe entity s across a set of N calibration dates $\{t_1, \dots, t_N\}$ drawn from the 2006–2015 period and compute:

$$\bar{H}_s = \frac{1}{N} \sum_{i=1}^N \hat{H}(s, t_i) \quad (9)$$

$$\sigma_s = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{H}(s, t_i) - \bar{H}_s)^2} \quad (10)$$

The **date-variance** σ_s measures how much the model’s directional confidence fluctuates across dates for a *fixed* entity. When σ_s is high, the model’s confidence varies with the date, a pattern consistent with date-specific knowledge. When σ_s is low, the model is similarly confident across dates, a pattern consistent with a stable entity prior.

The entity-adaptive penalty combines two multiplicative factors:

$$s_\sigma = \min(1, \sigma_s / \sigma_{\text{ref}}) \quad (11)$$

$$s_H = \max(0, (\bar{H}_s - \hat{H}(s, t)) / \Delta_{\text{range}}) \quad (12)$$

$$\alpha(s, t) = \max(\alpha_{\min}, \min(\alpha_{\max} s_\sigma s_H, \alpha_{\text{cap}})) \quad (13)$$

where $\alpha_{\max}$ scales the nominal peak penalty, α_{cap} is a hard upper bound on the resulting strength, $\alpha_{\min}$ is a non-negative floor (set to 0 in our experiments), and σ_{ref} and Δ_{range} are model-specific hyperparameters from profiling (described in the “Model-specific normalisation” paragraph below). The first factor (date-variance scale, Eq. 11) attenuates α for entities whose confidence is stable across dates, precisely those driven by brand priors rather than memorisation. The second factor (per-date confidence, Eq. 12) activates the penalty only at dates where the model is *more confident than its own average* for that entity.

This design has three key properties. First, it is **entity-adaptive**: each entity’s α is calibrated to its own memorisation profile, avoiding both over-subtraction on brand-prior-driven entities and

under-subtraction on strongly memorised entities. Second, it is **temporally selective**: even for a memorised entity, $\alpha > 0$ only at dates where the model is unusually confident, avoiding blanket subtraction across the entire backtest period. Third, it is **temporally attenuating**: when post-cutoff dates produce entropy above $\bar{H}_s$, $\bar{H}_s - \hat{H} < 0$ and the penalty falls to $\alpha = 0$.

Entity calibration. At the start of each backtest, we probe the target entity across $N = 12$ quarterly dates from 2006–2015 and compute $\bar{H}_s$ and σ_s (Eqs. 9–10). The calibration dates are fixed and entity-agnostic, and lie within the period used to optimise T_{prior}^* , so the probe measures the same T_{prior}^* -activated recall that decoding will subtract.

Model-specific normalisation. σ_{ref} and Δ_{range} come from a one-time profiling pass that records the entropy distribution over representative (entity, date) pairs. We set σ_{ref} to the standard deviation of the post-cutoff entropy distribution and Δ_{range} to the in-sample entropy standard deviation. Neither value uses target labels or backtest outcomes, but using post-cutoff model confidence makes this normalisation label-free and transductive. An inductive variant can instead estimate σ_{ref} on an earlier post-cutoff calibration window and freeze it before evaluation on a later window; we do not evaluate that variant here.

For reporting, we summarise the per-(entity, date) penalty over a backtest window by its mean: $\bar{\alpha}_{\text{IS}}$ averages $\alpha(s, t)$ over all in-sample dates of a backtest, and $\bar{\alpha}_{\text{OOS}}$ is the analogous mean over out-of-sample dates.

Computational structure. FINCAD splits the work into offline and online stages, which clarifies the cost profile. The offline work is done once per model (T_{prior}^* optimisation, profiling of σ_{ref} and Δ_{range}) and once per entity before each backtest ($\bar{H}_s$ and σ_s from twelve dated probes). The online work consists of one completion probe per (entity, date) decision to set $\alpha(s, t)$, followed at each decoding step by the two-branch logit blend in Eq. 1. This adds one prior-branch forward pass per generated token plus one probe per (entity, date) decision.

4 Experiments and Results

Our evaluation is structured around a central concern: does FINCAD mitigate look-ahead bias with-

out broadly degrading the model’s task performance?

A non-selective penalty could reduce backtest returns while also degrading tasks that require general reasoning. The three experiments examine complementary properties. **Experiment 1** (§4.1, *reasoning preservation*) measures changes on standard reasoning benchmarks under a fixed model-specific strength α_{bench} . **Experiment 2** (§4.2, *in-sample performance correction*) evaluates the entity- and date-adaptive penalty $\alpha(s, t)$ on in-sample and post-cutoff dates. **Experiment 3** (§4.3, *ranking alignment*) measures whether the FINCAD-adjusted in-sample model ranking is more closely aligned with the post-cutoff ranking.

Experiments 2 and 3 share a single-stock backtesting framework²: the LLM acts as a quantitative analyst, receives a structured price-derived summary at each trading date, and emits a buy/sell/hold signal that is executed at the next opening price under a long-only portfolio. We extend the framework with notional transaction costs and a liquidity cap for realism; full details are deferred to Appendix A. At inference, FINCAD adds one prior-branch forward pass per generated token and one completion probe per (entity, date) decision; the remaining calibration is performed offline (§3.2).

4.1 Experiment 1: Reasoning Preservation

We evaluate FINCAD across five instruction-tuned LLMs spanning 7B–14B parameters: **Phi-4-14B**, **Qwen2.5-14B**, **Llama-3.1-8B**, **Starling-7B**, and **DeepSeek-7B-Chat**. Each model is run on five standard benchmarks under both baseline decoding and FINCAD: FinEval (Ke et al., 2025) (financial multiple-choice across 20 subtasks), MMLU-Pro (Wang et al., 2024b), GSM8K (Cobbe et al., 2021), Competition-Math (Hendrycks et al., 2021), and HumanEval (Chen et al., 2021). Because these benchmarks contain no dated financial entity on which to evaluate $\alpha(s, t)$, we apply the fixed model-specific strengths α_{bench} listed in Table 1. We distinguish this benchmark setting from $\bar{\alpha}_{\text{IS}}$, which is the empirical mean of the adaptive strengths over dated backtest decisions.

Table 1 reports the full five-model sweep. Mean accuracy remains positive or within -1.7 points of baseline for four of five models; DeepSeek-7B-Chat is the exception (-4.3 points), driven primar-

	 Phi	 Qwen	 Llama	 Starl.	 DSeek
α_{bench}	.18	.81	.26	.62	.90
 GSM8K	+0.5	+0.4	−4.9	−6.9	−21.3
 MMLU-Pro	−0.2	+2.1	+6.6	+0.2	+2.4
 Comp. Math	+0.8	−2.0	−7.8	−1.0	−4.6
 FinEval	+7.1	−0.6	+0.2	−1.0	+2.8
 HumanEval	+3.0	+0.6	+0.6	0.0	−0.6
<i>mean Δ</i>	+2.2	+0.1	−1.0	−1.7	−4.3

Table 1: General-benchmark accuracy changes (FINCAD – baseline, percentage points) at the fixed model-specific strength α_{bench} . Columns denote Phi-4-14B, Qwen2.5-14B, Llama-3.1-8B, Starling-7B, and DeepSeek-7B-Chat. Greedy decoding, seed 42.

ily by a -21.3 -point change on GSM8K. Table 6 of Appendix D pairs these per-benchmark changes with the mean baseline and FINCAD accuracy for all five models.

4.2 Experiment 2: In-Sample Performance Correction

We first run the single-stock backtesting framework on five heavily memorised mega-cap equities (NVDA, MSFT, AAPL, NFLX, AMZN) for the same five models introduced in Experiment 1. These tickers are widely represented in financial news, filings, and analyst coverage, increasing the likelihood that their histories appear in pre-training corpora. Their prominence also makes them common testbeds in recent LLM-trading work (Yu et al., 2023; Ding et al., 2024; Xiao et al., 2025). Each ticker is evaluated over both the in-sample period (2010–2020, where evaluated dates precede every model’s training cutoff) and the post-cutoff period (calendar year 2025); Table 2 reports Phi-4-14B and Qwen2.5-14B, with the other three models in Appendix D. This paired design measures whether FINCAD produces larger corrections on dates that may overlap pre-training data than on dates beyond each model’s reported cutoff.

In-sample performance correction. For Qwen2.5-14B, FINCAD reduces returns on *all five* tickers with a mean relative drop of -45.4% of baseline return and -25.6% of baseline Sharpe, reaching -78.2% (NVDA) and -60.8% (MSFT) on the most heavily covered names. The mean calibrated $\bar{\alpha}_{\text{IS}} = 1.06$ confirms that the entity-adaptive penalty is firing strongly on the dates Qwen2.5-14B is most over-confident. For Phi-4-14B, four of five tickers drop (-19.1% mean

²<https://github.com/virattt/ai-hedge-fund>

return, -6.3% Sharpe), with mean $\bar{\alpha}_{IS} = 0.23$ indicating lower memorisation.

Post-cutoff behaviour. On the 2025 OOS window, Phi-4-14B receives $\bar{\alpha}_{OOS} = 0$ on every trading day, leaving the FINCAD trajectory byte-identical to baseline ($\Delta = 0$ across all five tickers). Qwen2.5-14B receives $\bar{\alpha}_{OOS} = 0.16$ on average, with non-zero firings only on MSFT (+\$8K, +0.39 Sharpe) and AAPL (+\$2K, +0.06 Sharpe); three of five tickers remain byte-identical. The post-cutoff effect on these two models is therefore zero or mildly positive, whereas the in-sample corrections are larger. The temporal dynamics of $\alpha(s, t)$ for the full five LLMs are reported in Appendix D. Experiment 1 (§4.1) provides a complementary measure of whether FINCAD broadly degrades task performance.

Decision-level check. On the first rebalance date of three saved Phi-4 trajectories, FINCAD applies $\alpha = 3.0$ (NVDA), 3.0 (MSFT), and 2.4 (AAPL), yet preserves the baseline buy action and executed share count in all three cases; confidence changes by $+5$, -5 , and 0 points, respectively.

Generalising to mid- and small-cap stocks. We extend the analysis beyond mega-caps with multiple mid-caps and small-caps. The tier-level $\bar{\alpha}_{IS}$ decreases from mega- to mid- to small-cap entities for Phi-4 ($0.23 \rightarrow 0.15 \rightarrow 0.14$) and Qwen2.5-14B ($1.06 \rightarrow 0.90 \rightarrow 0.80$; Appendix D.2). The in-sample drops on Qwen2.5-14B likewise attenuate from -45.4% (mega) to -26.7% (mid) and -0.9% (small); Phi-4 remains in its smaller-correction regime across the three tiers. The full results for all five backtested models and all mid- and small-caps are in Appendix D.2.

Entity-specific activation. Table 3 reports $\bar{\alpha}_{IS}$ for each model-ticker pair in the mega-cap evaluation. Every model exhibits within-model variation (max-min from 0.40 to 1.74), and the ticker attaining the maximum differs across models. Thus, the adaptive penalty does not reduce to a single model-level coefficient or a uniform ordering of entities.

4.3 Experiment 3: Cross-Model Ranking Alignment

We evaluate FINCAD across 11 instruction-tuned LLMs from 10 families spanning 7B–14B parameters (Phi-4, Llama-3.1, Qwen2.5-14B, Qwen2.5-7B, GLM-4, Mistral, Yi-1.5, OLMo-2, Starling,

Ticker	IS (2010–2020)				OOS (2025)			
	H\$K	B\$K	C\$K	$\Delta\%$	H\$K	B\$K	C\$K	$\Delta\%$
 <i>Phi-4-14B (medium memorisation)</i>								
NVDA	1379	1300	813	-37.5	137	109	109	0.0
MSFT	654	548	554	$+1.2$	114	112	112	0.0
AAPL	1106	1109	893	-19.5	110	122	122	0.0
NFLX	4075	4646	3610	-22.3	105	102	102	0.0
AMZN	1355	934	773	-17.2	104	103	103	0.0
<i>mega mean</i>				-19.1				0.0
POOL (<i>mid</i>)	1276	901	928	$+3.0$	68	67	74	$+9.8$
SAM (<i>small</i>)	801	692	562	-18.8	65	55	64	$+15.7$
 <i>Qwen2.5-14B (high memorisation)</i>								
NVDA	1379	1089	237	-78.2	137	128	128	0.0
MSFT	654	652	255	-60.8	114	100	108	$+8.0$
AAPL	1106	982	806	-17.9	110	115	117	$+1.1$
NFLX	4075	2572	1491	-42.0	105	101	101	0.0
AMZN	1355	1237	889	-28.2	104	112	112	0.0
<i>mega mean</i>				-45.4				$+1.8$
POOL (<i>mid</i>)	1276	1027	753	-26.7	68	81	73	-9.7
SAM (<i>small</i>)	801	415	412	-0.9	65	69	76	$+10.2$

Table 2: Per-ticker in-sample (2010–2020) and post-cutoff (2025) single-stock backtest results across three cap tiers, daily rebalancing, $T=1.0$, seed 42. **H\$K**: passive buy-and-hold ending value (initial \$100K). **B\$K/C\$K**: Baseline / FINCAD ending value. $\Delta\%$: $(C-B)/|B|$. The per-ticker Sharpe-ratio breakdown is in Appendix D.

Falcon3, DeepSeek; per-model family/size/cutoff in Table 11 of Appendix D.3). All models have reported training cutoffs before January 2025, so the 2025 evaluation period follows every reported cutoff.

The two simpler baselines are anonymisation (every ticker $\rightarrow$ [ticker N], every company name $\rightarrow$ [company N]; e.g. SPY $\rightarrow$ [ticker 1]) and prompt-injection (the system prompt’s look-ahead disclaimer is replaced with a strengthened directive; verbatim text in Appendix B.2). Together they compare three intervention granularities: anonymisation is a coarse input-side intervention, prompt-injection a coarse instruction-level intervention, and FINCAD a decoder-side intervention that scales logit subtraction using the model’s per-(entity, date) confidence signal.

Since any particular selection of models can affect the correlation, we adopt an exhaustive subset sensitivity analysis: we enumerate all $\binom{11}{7} = 330$ subsets of seven models and, for each subset, compute Spearman’s ρ and Kendall’s τ between each in-sample condition and the post-cutoff ranking; we write ρ_B for the baseline per-subset correlation and ρ_C for the corresponding FINCAD value. We report the subset-averaged means $\bar{\rho}$ and $\bar{\tau}$ over the

	Phi	Qwen	Llama	Starl.	DSeek
NVDA	0.22	0.73	0.33	0.87	1.12
MSFT	0.22	2.16	0.20	0.52	1.78
AAPL	0.10	1.33	0.22	0.82	0.50
NFLX	0.50	0.42	0.68	1.28	0.85
AMZN	0.12	0.63	0.32	0.62	1.03
<i>mean</i>	0.23	1.06	0.35	0.82	1.06
<i>max-min</i>	0.40	1.74	0.48	0.75	1.28

Table 3: Per-ticker mean adaptive strength $\bar{\alpha}_{IS}$ over the in-sample mega-cap backtests (2010–2020). Model columns follow Table 1; bold marks the maximum within each model.




Figure 2: Each marker is one of 11 LLMs at its (IS, OOS) Sharpe on SPY; the red line is the OLS fit and the dashed line is $y=x$ for reference.

330 subsets and supplement them with leave-one-model-out sensitivity analysis. Because the subsets overlap, they are treated as descriptive perturbations of the model pool rather than independent observations for inferential testing.

Ranking alignment. Table 4 reports the per-model SPY Sharpe ratios under each in-sample condition together with the out-of-sample baseline, and Table 5 reports the exhaustive correlation analysis. FINCAD is the only mitigation with higher subset-averaged alignment: by Sharpe ratio, $\bar{\rho}$ rises from +0.779 (baseline) to +0.846 ($\Delta\bar{\rho} = +0.067$). The two simpler baselines move the descriptive alignment in the opposite direction: anonymisation lowers $\bar{\rho}$ to +0.547 and prompt-injection to +0.482. Sortino-ranked alignment behaves the same way. In leave-one-model-out analysis, FINCAD improves the Spearman correlation in 10 of 11 omissions; the exception is the omission of Qwen2.5-14B. We report these quantities as observed effect sizes and sensitivity analyses; formal inference requires additional independent models and post-cutoff periods. Across subset sizes $k \in \{5, \dots, 11\}$, the observed Spearman improvement remains between +0.061

Model	IS-B	IS-An	IS-PI	IS-C	OOS
Phi-4-14B	0.620	0.578	0.673	0.585	0.892
Llama-3.1-8B	0.501	0.329	0.229	0.248	0.860
Qwen2.5-14B	0.636	0.653	0.685	0.314	0.732
Qwen2.5-7B	0.230	0.422	0.595	0.123	-0.247
GLM-4-9B	-0.154	-0.080	0.321	-0.272	0.156
Mistral-7B	0.685	0.364	0.640	0.682	0.926
Yi-1.5-9B	0.620	0.652	0.560	0.578	1.061
OLMo-2-7B	-0.266	-0.231	0.103	-0.404	-0.308
Starling-7B	0.335	0.353	0.337	0.058	0.812
Falcon3-7B	0.382	0.614	0.551	0.337	0.859
DeepSeek-7B	0.173	-0.352	-0.514	-0.242	-0.249

Table 4: Per-model SPY Sharpe ratios under each in-sample condition (Baseline, Anonymisation, Prompt-Injection, FINCAD) and the out-of-sample baseline. $T=1.0$, seed 42. These per-model values underlie the subset-averaged $\bar{\rho} / \bar{\tau}$ statistics in Table 5.

Method	Spearman $\bar{\rho}$		Kendall $\bar{\tau}$	
	Sharpe	Sortino	Sharpe	Sortino
Baseline	+0.779	+0.642	+0.673	+0.491
Anonymisation	+0.547 $\downarrow$	+0.488 $\downarrow$	+0.418 $\downarrow$	+0.382 $\downarrow$
Prompt-injection	+0.482 $\downarrow$	+0.320 $\downarrow$	+0.345 $\downarrow$	+0.236 $\downarrow$
FINCAD	+0.846$\uparrow$	+0.694$\uparrow$	+0.709$\uparrow$	+0.527$\uparrow$

Table 5: Subset-averaged correlation ($\binom{11}{7}=330$ overlapping subsets) between each in-sample mitigation’s ranking and the post-cutoff baseline ranking. $\uparrow/\downarrow$ shows the direction of the observed change relative to baseline.

and +0.073 (Appendix F).

Analysis. Figure 2 visualises this realignment as an alignment scatter. A higher Spearman ρ corresponds to markers more tightly arranged in a monotonic pattern, meaning the IS-axis ordering of models agrees with the OOS-axis ordering. One illustrative correction is on Qwen2.5-14B (orange square): under baseline it ranks IS top-3 (Sharpe 0.636) but delivers only the 7th-best OOS Sharpe (0.732); FINCAD corrects its IS Sharpe to 0.314, demoting its IS rank to a position consistent with its OOS rank. Models with smaller memorisation footprints whose in-sample edge survives OOS (Mistral-7B, Yi-1.5, Phi-4-14B) barely shift between the two panels, and the OOS top-3 is preserved in both.

4.4 Takeaway

The three experiments characterise complementary aspects of FINCAD: changes on general benchmarks (§4.1), the relative magnitude of in-sample and post-cutoff corrections (§4.2), and the alignment of in-sample and post-cutoff model rankings

(§4.3). We view FINCAD as a contamination control conceptually parallel to the deflated Sharpe (Bailey and López de Prado, 2014) and survivorship corrections (Heckman, 1979; Brown et al., 1992), providing a tool for re-evaluating agent pipelines when model training data overlap the backtest period.

5 Conclusion

To mitigate *parametric look-ahead bias*, we proposed FINCAD, an inference-time adaptation of Context-Aware Decoding that attenuates a memory-enriched prior without retraining, by combining an adversarial bias-discovery pipeline that learns a model-specific prior prompt with an entity- and date-adaptive rule for the CAD strength α .

Across multiple 7–14B LLMs, FINCAD produces model-dependent in-sample corrections, with limited post-cutoff changes for the three larger models and limited mean general-benchmark changes for four of five models. In the leaderboard analysis, the subset-averaged Spearman correlation between in-sample and post-cutoff Sharpe rankings rises from +0.779 to +0.846, improving their descriptive alignment.

More broadly, FINCAD is a complementary intervention that integrates with existing trading-agent evaluation pipelines and works alongside input-side protections when training-cutoff overlap is unavoidable. Its operational template may extend to other counterfactual evaluations involving named entities and dated outcomes, but such extensions require domain-specific calibration and validation.

Limitations

First, FINCAD only *partially* mitigates parametric look-ahead bias: the calibrated logit subtraction reduces the contribution of memorised future knowledge but does not guarantee its complete removal. Residual memorisation may persist whenever the prior prompt fails to fully activate the model’s parametric recall, or whenever memorisation is encoded in pathways that the chosen probe does not detect. The entropy-based signal is a heuristic proxy for date-specific recall: no counterfactual ground truth permits identifiable precision or recall for $\alpha(s, t)$, or an exact decomposition of a portfolio change into memorised and context-grounded components. We therefore interpret the results as system-level mitigation evidence rather than pointwise causal

attribution or an exact estimate of contamination. The reported normalisation is also label-free but transductive because it uses post-cutoff model confidence; the inductive, temporally frozen variant described in §3.2 remains to be evaluated.

Second, our evaluation is restricted to instruction-tuned LLMs in the 7B–14B parameter range. Whether the same calibration mechanism transfers cleanly to substantially larger models (e.g., 70B+ or frontier closed-source models) remains open. The eleven-model evaluation required several thousand GPU-hours; evaluation across multiple frontier-scale models remains future work. FINCAD also requires access to next-token logits. APIs exposing action-token log probabilities could support paired external contrasts, while text-only APIs would require repeated sampling with greater cost and variance; neither black-box variant is evaluated here.

Third, our backtests are restricted to price-derived summaries in long-only, single-stock US-equity portfolios (mega-, mid-, and small-cap names and SPY), with one temperature and one seed per condition. Transfer to agents using news and filings, multi-asset portfolios, other markets, or longer action trajectories remains an empirical question; the magnitude of the bias and the appropriate α range may differ from those observed here.

FINCAD provides an operational intervention for studying parametric look-ahead bias in LLM-based backtesting. Existing work documents the bias qualitatively (Gao et al., 2026; Benhenda, 2026), while direct identification of its contribution to a specific portfolio trajectory remains open.

Acknowledgments

Generative AI tools, including OpenAI Codex, were used solely to polish author-written prose and to assist with code implementation and debugging. All suggestions and generated code were reviewed and verified by the authors, who take full responsibility for the manuscript and released artifacts.

References

David H. Bailey, Jonathan M. Borwein, Marcos López de Prado, and Qiji Jim Zhu. 2014. Pseudomathematics and financial charlatanism: The effects of backtest overfitting on out-of-sample performance. *Notices of the American Mathematical Society*, 61(5):458–471.

- David H. Bailey and Marcos M. López de Prado. 2014. [The deflated sharpe ratio: Correcting for selection bias, backtest overfitting, and non-normality](#). *The Journal of Portfolio Management*, 40(5):94–107.
- Mostapha Benhenda. 2026. [Look-ahead-bench: a standardized benchmark of look-ahead bias in point-in-time llms for finance](#). *Preprint*, arXiv:2601.13770.
- Stephen J. Brown, William Goetzmann, Roger G. Ibbotson, and Stephen A. Ross. 1992. [Survivorship bias in performance studies](#). *The Review of Financial Studies*, 5(4):553–580.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang. 2023. [Quantifying memorization across neural language models](#). In *The Eleventh International Conference on Learning Representations*.
- Nicholas Carlini, Florian Tramèr, Eric Wallace, Matthew Jagielski, Ariel Herbert-Voss, Katherine Lee, Adam Roberts, Tom Brown, Dawn Song, Úlfar Erlingsson, Alina Oprea, and Colin Raffel. 2021. [Extracting training data from large language models](#). In *30th USENIX Security Symposium (USENIX Security 21)*, pages 2633–2650. USENIX Association.
- Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, and 39 others. 2021. [Evaluating large language models trained on code](#). *Preprint*, arXiv:2107.03374.
- Yung-Sung Chuang, Yujia Xie, Hongyin Luo, Yoon Kim, James R. Glass, and Pengcheng He. 2024. [Dola: Decoding by contrasting layers improves factuality in large language models](#). In *The Twelfth International Conference on Learning Representations*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Zheyang Deng, Weixiang Yan, Changlong Yu, and Jiashu Wang. 2026. [Alphaquant: An end-to-end tool-augmented agentic reinforcement learning framework for stock trading](#). *Preprint*, arXiv:2510.14264.
- Qianggang Ding, Haochen Shi, and Bang Liu. 2024. Tradexpert: Revolutionizing trading with mixture of expert llms.
- Yujie Ding, Shuai Jia, Tianyi Ma, Bingcheng Mao, Xizhe Zhou, Liuliu Li, and Dongming Han. 2023. Integrating Stock Features and Global Information via Large Language Models for Enhanced Stock Return Prediction. Papers 2310.05627, arXiv.org.
- Ronen Eldan and Mark Russinovich. 2023. [Who’s Harry Potter? Approximate unlearning in LLMs](#). *Preprint*, arXiv:2310.02238.
- George Fatouros, Kostas Metaxas, John Soldatos, and Manos Karathanassis. 2025. Marketsenseai 2.0: Enhancing stock analysis through llm agents.
- Zhenyu Gao, Wenxi Jiang, and Yutong Yan. 2026. [Detecting lookahead bias in llm forecasts](#). *Preprint*, arXiv:2512.23847.
- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V. Chawla, Olaf Wiest, and Xi-angliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI 2024, Jeju, South Korea, August 3-9, 2024*, pages 8048–8057.
- Campbell R. Harvey and Yan Liu. 2014. Evaluating trading strategies. *The Journal of Portfolio Management*, 40(5):108–118.
- James J. Heckman. 1979. [Sample selection bias as a specification error](#). *Econometrica*, 47(1):153–161. Full publication date: Jan., 1979.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. [Measuring mathematical problem solving with the MATH dataset](#). In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)*.
- Joel Jang, Dongkeun Yoon, Sohee Yang, Sungmin Cha, Moontae Lee, Lajanugen Logeswaran, and Minjoon Seo. 2023. [Knowledge unlearning for mitigating privacy risks in language models](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14389–14408, Toronto, Canada. Association for Computational Linguistics.
- Zixuan Ke, Yifei Ming, Xuan-Phi Nguyen, Caiming Xiong, and Shafiq Joty. 2025. [Demystifying domain-adaptive post-training for financial LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 31033–31059, Suzhou, China. Association for Computational Linguistics.
- Omar Khattab, Arnav Singhvi, Paridhi Maheshwari, Zhiyuan Zhang, Keshav Santhanam, Sri Vardhamanan, Saiful Haq, Ashutosh Sharma, Thomas T. Joshi, Hanna Moazam, Heather Miller, Matei Zaharia, and Christopher Potts. 2024. Dspy: Compiling declarative language model calls into self-improving pipelines. In *The Twelfth International Conference on Learning Representations*.
- Kelvin J. L. Koa, Yunshan Ma, Ritchie Ng, and Tat-Seng Chua. 2024. Learning to generate explainable stock predictions using self-reflective large language

- models. In *Proceedings of the ACM on Web Conference 2024, WWW 2024, Singapore, May 13-17, 2024*, pages 4304–4315.
- Yaxuan Kong, Hoyoung Lee, Yoontae Hwang, Alejandro Lopez-Lira, Bradford Levy, Dhagash Mehta, Qingsong Wen, Chanyeol Choi, Yongjae Lee, and Stefan Zohren. 2026. [Evaluating llms in finance requires explicit bias consideration](#). *Preprint*, arXiv:2602.14233.
- Weixian Waylon Li, Hyeonjun Kim, Mihai Cucuringu, and Tiejun Ma. 2026a. [Can llm-based financial investing strategies outperform the market in long run?](#) In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1, KDD '26*, page 2711–2722, New York, NY, USA. Association for Computing Machinery.
- Weixian Waylon Li and Tiejun Ma. 2025. [Learn to rank risky investors: A case study of predicting retail traders' behaviour and profitability](#). *ACM Trans. Inf. Syst.*, 44(1).
- Weixian Waylon Li, Yuchen Niu, Yongxin Yang, Keshuang Li, Tiejun Ma, and Shay B Cohen. 2026b. [Spectral attention steering for prompt highlighting](#). In *The Fourteenth International Conference on Learning Representations*.
- Weixian Waylon Li, Jiaxin Zhang, Xianan Jim Yang, Tiejun Ma, and Yiwen Guo. 2026c. [Time is not a label: Continuous phase rotation for temporal knowledge graphs and agentic memory](#). In *The 2026 Conference on Empirical Methods in Natural Language Processing*.
- Xiang Lisa Li, Ari Holtzman, Daniel Fried, Percy Liang, Jason Eisner, Tatsunori Hashimoto, Luke Zettlemoyer, and Mike Lewis. 2023. [Contrastive decoding: Open-ended text generation as optimization](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12286–12312, Toronto, Canada. Association for Computational Linguistics.
- Ying Li, Mengyu Wang, Miguel de Carvalho, Sotirios Sabanis, and Tiejun Ma. 2025. [FinGEAR: Financial mapping-guided enhanced answer retrieval](#). In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 7239–7255, Suzhou, China. Association for Computational Linguistics.
- Sijia Liu, Yuanshun Yao, Jinghan Jia, Stephen Casper, Nathalie Baracaldo, Peter Hase, Yuguang Yao, Chris Yuhao Liu, Xiaojun Xu, Hang Li, Kush R. Varshney, Mohit Bansal, Sanmi Koyejo, and Yang Liu. 2024. [Rethinking machine unlearning for large language models](#). *Preprint*, arXiv:2402.08787.
- Marcos López de Prado. 2018. *Advances in Financial Machine Learning*. Wiley.
- Pratyush Maini, Zhili Feng, Avi Schwarzschild, Zachary C. Lipton, and J. Zico Kolter. 2024. [TOFU: A task of fictitious unlearning for LLMs](#). *Preprint*, arXiv:2401.06121.
- Kevin Meng, David Bau, Alex J Andonian, and Yonatan Belinkov. 2022. [Locating and editing factual associations in GPT](#). In *Advances in Neural Information Processing Systems*.
- Kevin Meng, Arnab Sen Sharma, Alex J Andonian, Yonatan Belinkov, and David Bau. 2023. [Mass-editing memory in a transformer](#). In *The Eleventh International Conference on Learning Representations*.
- Humzah Merchant and Bradford Levy. 2025. [A fast and effective solution to the problem of look-ahead bias in LLMs](#). In *NeurIPS 2025 Workshop: Generative AI in Finance*.
- Robert Novy-Marx and Mihail Velikov. 2016. [A taxonomy of anomalies and their trading costs](#). *The Review of Financial Studies*, 29(1):104–147.
- Krista Opsahl-Ong, Michael J Ryan, Josh Purtell, David Broman, Christopher Potts, Matei Zaharia, and Omar Khattab. 2024. [Optimizing instructions and demonstrations for multi-stage language model programs](#). In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 9340–9366, Miami, Florida, USA. Association for Computational Linguistics.
- Preetha Saha, Jingrao Lyu, Arnav Saxena, Tianjiao Zhao, and Dhagash Mehta. 2025. [Large language model agents for investment management: Foundations, benchmarks, and research frontiers](#). In *Proceedings of the 6th ACM International Conference on AI in Finance, ICAIF '25*, page 736–744, New York, NY, USA. Association for Computing Machinery.
- Weijia Shi, Xiaochuang Han, Mike Lewis, Yulia Tsvetkov, Luke Zettlemoyer, and Wen-tau Yih. 2024. [Trusting your evidence: Hallucinate less with context-aware decoding](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 783–791, Mexico City, Mexico. Association for Computational Linguistics.
- Mengyu Wang, Shay B. Cohen, and Tiejun Ma. 2024a. [Modeling news interactions and influence for financial market prediction](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3302–3314, Miami, Florida, USA. Association for Computational Linguistics.
- Mengyu Wang and Tiejun Ma. 2024. [Mana-net: Mitigating aggregated sentiment homogenization with news weighting for enhanced market prediction](#). In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, CIKM '24*, page 2379–2389, New York, NY, USA. Association for Computing Machinery.
- Mengyu Wang, Sotirios Sabanis, Miguel de Carvalho, Shay B. Cohen, and Tiejun Ma. 2025. [One more question is enough, expert question decomposition](#)

(EQD) model for domain quantitative reasoning. In *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 20355–20369, Suzhou, China. Association for Computational Linguistics.

Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. [MMLU-pro: A more robust and challenging multi-task language understanding benchmark](#). In *The Thirty-eight Conference on Neural Information Processing Systems Datasets and Benchmarks Track*.

Yijia Xiao, Edward Sun, Di Luo, and Wei Wang. 2025. [Tradingagents: Multi-agents llm financial trading framework](#). *Preprint*, arXiv:2412.20138.

Hongyang Yang, Boyu Zhang, Neng Wang, Cheng Guo, Xiaoli Zhang, Likun Lin, Junlin Wang, Tianyu Zhou, Mao Guan, Runjia Zhang, and Christina Dan Wang. 2024. Finrobot: An open-source ai agent platform for financial applications using large language models.

Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W. Suchow, and Khaldoun Khashanah. 2023. [Finmem: A performance-enhanced llm trading agent with layered memory and character design](#). *Preprint*, arXiv:2311.13743.

Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yuechen Jiang, Yupeng Cao, Zhi Chen, Jordan W. Suchow, Zhenyu Cui, Rong Liu, Zhaozhao Xu, Denghui Zhang, Koduvayur Subbalakshmi, GUO-JUN XIONG, Yueru He, Jimin Huang, Dong Li, and Qianqian Xie. 2024. [Fincon: A synthesized LLM multi-agent system with conceptual verbal reinforcement for enhanced financial decision making](#). In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*.

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. 2024. 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, KDD 2024, Barcelona, Spain, August 25-29, 2024*, pages 4314–4325.

A Backtesting Framework

This appendix documents the single-stock backtesting framework used in Experiments 2 and 3. The same execution and accounting rules are applied to baseline and FINCAD conditions so that their observed differences are not induced by the simulator.

Agent and signal. At each trading date the LLM acts as a quantitative analyst for a single ticker. It receives a structured financial summary built from price data only, with no news, filings, or macro feeds: the latest closing price, trailing returns over multiple windows (1m, 3m, 6m, 1y), simple and exponential moving averages, annualised realised volatility, the 52-week high/low and the current price’s position within that range, the count of positive days over the last 21 trading sessions, and (when available) the 21-day average dollar volume. Given this summary, the agent must return a JSON object with three fields: `action` $\in$ {buy, sell, hold}, `quantity` (integer share count), and `confidence` (integer in $[0, 100]$). Outputs are parsed with a lightweight tolerant JSON extractor; on parse failure the procedure retries with a reduced α .

Execution timing (no intraday look-ahead). The agent at rebalance date t sees only data with timestamp strictly less than t (i.e., up to the previous day’s close). Trades are then executed at the *next* trading day’s adjusted opening price, never the same day’s close that produced the signal. Daily portfolio values are marked-to-market using adjusted close. This ordering ensures the agent never observes the price at which it trades, which is the standard guard against intraday look-ahead.

Initial capital and portfolio. Each backtest starts with \$100,000 of cash and zero shares. The portfolio is long-only and consists of cash plus a position in a single ticker; short selling, leverage, and options are disabled. Cash earns no interest in the simulation (the risk-free rate enters only via Sharpe/Sortino computation, not via cash accrual).

Commission model. Transaction costs are modelled as 10 basis points (bps) of trade notional, charged one-way on both buys and sells. This is a central estimate for large- and mid-cap US equities covering broker commission, half-spread, and market impact (Novy-Marx and Velikov, 2016). The model is deliberately notional-based rather than per-share: split-adjusted prices can be orders of magnitude smaller than nominal prices, so a per-share fee would inflate effective commissions on heavily split tickers. A notional fee is split-invariant. Commission is deducted mechanically from the portfolio cash and is *not* mentioned in the LLM prompt, to avoid introducing asymmetric context between the context and the CAD prior. Insuffi-

cient cash is handled by reducing the requested share count until the order plus commission fits the available cash.

Liquidity cap. To prevent the agent from placing economically unrealistic orders, every order is capped at 1% of the trailing 20-day average daily volume (ADV) on the execution date. Orders requesting more than this cap are clipped to the cap; the remainder is dropped (no carry-over). The cap is disabled if volume data is unavailable for the ticker.

Performance metrics. We report total return, CAGR, annualised volatility, Sharpe ratio, Sortino ratio, maximum drawdown, and ending portfolio value, computed from the daily mark-to-market series with 252 trading days per year. The risk-free rate is fixed at 3% annualised (a long-run US T-bill average over our evaluation window), used in both Sharpe and Sortino. Each backtest is benchmarked against a buy-and-hold reference: an all-in purchase at the first execution open, held to the final close, with the same commission charged on the initial purchase.

Decoding configuration. All backtesting experiments use temperature $T=1.0$ with a fixed random seed for reproducibility. We initially tried greedy decoding ($T=0$) and found that it masks CAD’s effect: the logit subtraction shifts the full distribution but rarely flips the argmax token, so a large fraction of trading decisions remain identical between baseline and FINCAD. At $T=1.0$ the distribution shift propagates into sampled tokens, so CAD can meaningfully alter the agent’s outputs. Experiment 1 (general-reasoning benchmarks) instead uses $T=0$ to match standard evaluation protocols on those benchmarks.

The verbatim agent prompts, the two Experiment-3 mitigation variants, and the per-model discovered T_{prior}^* instructions are listed together in Appendix B.

B Prompts

This appendix collects the verbatim prompts used across all experiments: the baseline backtesting agent prompt (§B.1), the two Experiment-3 mitigation variants (§B.2), and the model-specific memory-activation instructions T_{prior}^* discovered by MIPROv2 (§B.3). Curly-brace tokens (e.g. {ticker}, {date}) are runtime substitution slots

filled per (entity, date) at backtest time; everything else is verbatim.

B.1 Backtesting agent prompt (baseline)

The baseline single-stock agent (Experiments 2–3) receives a system message and a context body, and must respond with a single JSON object. The system message establishes the role, the look-ahead disclaimer, and the output schema; the context body supplies the price-derived financial summary, the current portfolio state, and the action-space limits.

System message

You are a portfolio manager for a single-stock strategy.
You will be given financial data for {ticker} as of {date}.
Base your decision ONLY on the data provided – do not use any knowledge about events after {date}.
You must pick one action and a quantity within the allowed limits.
Return your decision as a JSON object with exactly these keys:

- "action": one of "buy", "sell", or "hold"
- "quantity": integer number of shares to trade (0 for hold)
- "confidence": integer between 0 and 100
- "reasoning": string with a concise rationale (max 100 chars)

Respond with valid JSON only.

Context body

```
=== Financial Data for {ticker} as of {date}
===
{financial_summary}
=== Portfolio State ===
Cash: ${cash}
Current Shares: {shares}
Portfolio Value: ${portfolio_value}
=== Allowed Actions ===
{allowed_actions}
Based solely on the data above, what is your trading decision?
```

The financial_summary slot is populated by the price-derived summary documented in Appendix A (trailing returns, moving averages, realised volatility, 52-week range, positive-day count, ADV).

B.2 Mitigation prompts (Experiment 3)

Experiment 3 evaluates two input-side mitigations alongside FINCAD. Both edit the prompts above; all other agent state (rebalance schedule, output schema, decoding configuration) is identical to the baseline.

Anonymisation. Every occurrence of a ticker symbol is replaced with the placeholder [ticker N] and every occurrence of a company name with [company N], where N is a per-entity counter assigned in order of first appearance (e.g., $\text{TSLA} \rightarrow [\text{ticker } 1]$, $\text{Tesla} \rightarrow [\text{company } 1]$). Replacement is applied to the entire prompt content (system message and financial-summary body) before the prompt is sent to the model. Example: the system message for a TSLA decision becomes “*You will be given financial data for [ticker 1] as of {date}*” rather than “... data for TSLA ...”.

Prompt-injection. The baseline system message contains a brief look-ahead disclaimer (“*Base your decision ONLY on the data provided – do not use any knowledge about events after {date}.*”). For the prompt-injection condition we replace the entire role-and-disclaimer block of the system message with the strengthened version below; the JSON output schema (the lower half of the system message) is identical to the baseline.

Prompt-injection system message (strengthened)

You are a portfolio manager for a single-stock strategy operating *as of {date}*. You will be given financial data for {ticker} as of {date}. Your information set is strictly limited to events and prices up to and including {date}; you have NO knowledge of anything that happened after that date. You must NOT recall, retrieve, or rely on any memory of {ticker}'s future behaviour. Do not reason about post-{date} price movements, news, earnings, products, or events. If your training data contains such information, treat it as inaccessible. Decisions based on post-{date} memory are a violation of these instructions. You must pick one action and a quantity within the allowed limits. Return your decision as a JSON object with exactly these keys:

- "action": one of "buy", "sell", or "hold"
- "quantity": integer number of shares to trade (0 for hold)
- "confidence": integer between 0 and 100
- "reasoning": string with a concise rationale (max 100 chars)

Respond with valid JSON only.

B.3 Optimised memory-activation instruction

For each model in our pool we run the adversarial bias-discovery pipeline of §3.1 once to obtain a model-specific memory-activation instruction T_{prior}^* . The same T_{prior}^* is reused as the prefix of the prior prompt in every Experiment-2

and Experiment-3 backtest for that model. Each T_{prior}^* below is the verbatim instruction selected by MIPROv2 on the held-out validation split of $\mathcal{D}_{\text{calib}}$. For three of our five backtested models the optimiser retained the seed instruction unchanged, a finding consistent with the seed already being a strong general-purpose memory probe.

Phi-4-14B

Imagine a high-stakes scenario in which the rapid financial stability of a major corporation is at risk due to emerging economic volatility. Your objective is to swiftly analyze the available information and insights you possess from your training data about this financial entity. Based on this analysis, predict a critical outcome regarding the corporation's financial resilience.

Qwen2.5-14B-Instruct

Recall what you know from your training data about this financial entity and predict the outcome.

Llama-3.1-8B-Instruct

Recall what you know from your training data about this financial entity and predict the outcome.

Starling-LM-7B

In a high stakes scenario where the Language Model must predict the outcome of a significant financial decision, the following instruction can be used to prompt the model: “Predict the potential impact on the stock market and investment opportunities for a major corporation if it successfully acquires a rival company in a multi-billion-dollar deal, considering historical trends, market conditions, and potential regulatory challenges.”

DeepSeek-7B-Chat

Recall what you know from your training data about this financial entity and predict the outcome.

The remaining six cross-model-validation models (Qwen2.5-7B, GLM-4-9B, Mistral-7B, Yi-1.5-9B, OLMo-2-7B, Falcon3-7B) use T_{prior}^* instructions discovered with the same pipeline; their full text is omitted for brevity but is available in the released discovery artefacts .

C FINCAD Inference Algorithm

Algorithm 1 gives the complete inference-time procedure for a single (entity, date) decision under FINCAD, combining the prior-prompt construction (§3.1), the entity-adaptive penalty (§3.2), and the context-aware decoding loop with adaptive retry on parse failure.

D Per-Model Detail: Reasoning Preservation and Single-Stock Backtests

This appendix provides the full per-model evidence underlying the main-paper claims in §4. Section D.1 gives the full five-model per-benchmark changes and mean absolute accuracies (Table 6), complementing the compact main-paper summary in Table 1. Section D.2 extends Experiment 2 with the temporal dynamics of the adaptive strength (Figure 3 and Table 7) and full per-ticker IS/OOS breakdowns for the three models not shown in Table 2 (Table 8: Llama-3.1-8B, Starling-7B, DeepSeek-7B-Chat), as well as the full mid- and small-cap robustness panel across all five backtested models (Tables 9 and 10). Section D.3 lists the 11 backtested LLMs and their training cutoffs (Table 11).

D.1 Experiment 1: Per-Benchmark Detail

Table 6 reports the per-benchmark change for all five models together with each model’s mean baseline and FINCAD accuracy. The fixed α_{bench} used for each model is included alongside the model name.

D.2 Experiment 2: Cross-Model and Per-Ticker Detail

Temporal dynamics of the calibrated penalty.

Figure 3 shows the daily calibrated α across the entire IS period (2010–2020) for every (model, ticker) pair, with a 63-day rolling mean overlaid. Two qualitatively different temporal patterns emerge. For Phi-4-14B, Llama-3.1-8B, Starling-7B, and DeepSeek-7B-Chat, α is concentrated in the *early* part of the IS window and decays towards zero in later years. Qwen2.5-14B exhibits the opposite pattern: α is low in 2010–2012 and rises from 2014 onwards. Table 7 reports the year-bucketed numbers underlying this contrast.

Cross-model summary. The mean calibrated $\bar{\alpha}_{\text{IS}}$ rises monotonically with the model’s measured memorisation level, from 0.23 for Phi-4-14B (the lowest-memorisation backbone in our pool) to 1.06 for Qwen2.5-14B and DeepSeek-7B-Chat at the high end. For four of five models (Phi-4, Qwen2.5-14B, Starling-7B, DeepSeek-7B) the mean in-sample return drops under FINCAD, with the magnitude of the drop scaling with $\bar{\alpha}_{\text{IS}}$. Llama-3.1-8B is the lone exception: with mean $\bar{\alpha}_{\text{IS}} = 0.35$ FINCAD instead acts as a positive regulariser

Algorithm 1 FINCAD Inference

Require: Entity s , date t , model $\mathcal{M}$, system prompt T_{ctx} , data summary $D(s, t)$, optimised instruction T_{prior}^* , task format F_{task} , entity stats $\bar{H}_s, \sigma_s$ (from entity calibration), model params $\sigma_{\text{ref}}, \Delta_{\text{range}}, \alpha_{\text{min}}, \alpha_{\text{max}}, \alpha_{\text{cap}}$

Ensure: FINCAD-adjusted generated response

```

1:                                     ▷ Construct prompts
2:  $x_{\text{ctx}} \leftarrow T_{\text{ctx}} \oplus (s, t) \oplus F_{\text{task}} \oplus D(s, t)$  ▷ Eq. 2
3:  $x_{\text{prior}} \leftarrow T_{\text{prior}}^* \oplus (s, t) \oplus F_{\text{task}}$  ▷ Eq. 3
4:                                     ▷ Entity-adaptive penalty (§3.2)
5:  $x_{\text{probe}} \leftarrow T_{\text{prior}}^* \oplus \text{“After } t, s \text{ stock went”}$  ▷ Completion probe
6:  $\hat{H} \leftarrow \text{BINARYENTROPY}(\mathcal{M}, x_{\text{probe}})$  ▷ Eq. 8
7:  $\alpha \leftarrow \alpha_{\text{max}} \cdot \min(1, \sigma_s / \sigma_{\text{ref}}) \cdot \max(0, (\bar{H}_s - \hat{H}) / \Delta_{\text{range}})$  ▷ Eq. 13
8:  $\alpha \leftarrow \max(\alpha_{\text{min}}, \min(\alpha, \alpha_{\text{cap}}))$ 
9: ▷ Context-aware decoding with adaptive retry
10: repeat
11:   for each decoding step do
12:      $V_{\text{ctx}} \leftarrow \mathcal{M}(x_{\text{ctx}}); V_{\text{prior}} \leftarrow \mathcal{M}(x_{\text{prior}})$ 
13:      $\tilde{V} \leftarrow (1 + \alpha) V_{\text{ctx}} - \alpha V_{\text{prior}}$  ▷ logit blend
14:      $v \leftarrow \text{SAMPLE}(\text{softmax}(\tilde{V}))$  ▷ sample from  $\tilde{y}$ , Eq. 1
15:     Append  $v$  to both  $x_{\text{ctx}}$  and  $x_{\text{prior}}$ 
16:   end for
17:   if output is parseable then
18:     return generated sequence
19:   end if
20:    $\alpha \leftarrow 0.8 \alpha$  ▷ Reduce and retry (up to 5 times)
21: until parsed or retries exhausted

```

(mean return +26.5%, mean Sharpe +0.07), driven by a single AAPL outlier (+115.3%). This result illustrates that an in-sample correction need not decrease returns: changing the contribution of the prior can also yield a more profitable sampled trajectory. On out-of-sample (2025) all three larger models (Phi-4-14B, Qwen2.5-14B, Llama-3.1-8B) preserve the baseline within $\pm \$8\text{K}$ and the mean Sharpe within ± 0.10 ; the two smaller chat models (Starling-7B, DeepSeek-7B) show larger OOS variance, but the directional pattern of the mean is consistent.

Per-ticker detail. Table 8 extends Table 2 of the main paper to the three models not shown

Model (α_{bench})	 GSM8K	 MMLU-P	 C-Math	 FinEval	 HumEval	B	C	Δ
 Phi-4-14B (.18)	+0.5	−0.2	+0.8	+7.1	+3.0	59.8	62.0	+2.2
 Qwen2.5-14B (.81)	+0.4	+2.1	−2.0	−0.6	+0.6	68.1	68.2	+0.1
 Llama-3.1-8B (.26)	−4.9	+6.6	−7.8	+0.2	+0.6	57.9	56.8	−1.0
 Starling-7B (.62)	−6.9	+0.2	−1.0	−1.0	0.0	50.4	48.6	−1.7
 DeepSeek-7B-Chat (.90)	−21.3	+2.4	−4.6	+2.8	−0.6	46.4	42.2	−4.3

Table 6: Full five-model reasoning-preservation results. The five task columns report the accuracy change (FINCAD − baseline, percentage points); B and C report each model’s mean baseline and FINCAD accuracy across the five benchmarks. Greedy decoding, seed 42.

Model	2010–12	2013–15	2016–18	2019
 Phi-4-14B	0.68	0.05	0.03	0.03
 Qwen2.5-14B	0.40	1.12	1.49	1.52
 Llama-3.1-8B	0.63	0.45	0.10	0.00
 Starling-7B	1.31	0.79	0.62	0.04
 DeepSeek-7B	1.37	1.29	0.64	0.66

Table 7: Year-bucketed mean α averaged across the five tickers.

there. Column conventions are identical: B\$K and C\$K denote the ending portfolio value (\$K) under Baseline and FINCAD; $\Delta\%$ is the relative return change; ΔSh is the absolute Sharpe-ratio change. On Starling-7B FINCAD produces the strongest IS drop in our pool: returns collapse on all five tickers, with NVDA losing 90% of its baseline ending value. OOS preservation is also weaker than for the larger models, with $\bar{\alpha}_{\text{OOS}}=0.12$ producing residual drops on 4/5 tickers (memorisation that the OOS calibrator does not zero out). On DeepSeek-7B-Chat the AMZN OOS gain (+22.8%, $\Delta\text{Sh} +1.19$) is an artefact of the calibrator firing on a small number of dates where the baseline was holding a substantially negative-Sharpe position rather than a genuine OOS improvement.

Mid- and small-cap robustness. Table 2 of the main paper extends each panel with one mid-cap (POOL) and one small-cap (SAM) probe ticker. We additionally backtested every (model, ticker) combination across a six-ticker mid/small panel: KMX (CarMax), POOL (Pool Corporation), EXPE (Expedia) on the mid-cap side, and CROX (Crocs), SAM (Boston Beer), BJRI (BJ’s Restaurants) on the small-cap side. Tickers were chosen to have the full 2010–2025 daily price coverage that the IS+OOS protocol requires. Table 9 reports the per-(model, ticker) calibrated penalty $\bar{\alpha}_{\text{IS}}$. Two patterns generalise from the mega-cap pool. First, $\bar{\alpha}_{\text{IS}}$ scales with each model’s overall memorisation strength (Phi-4

Tkr	IS (2010–2020)				OOS (2025)			
	B\$K	C\$K	$\Delta\%$	ΔSh	B\$K	C\$K	$\Delta\%$	ΔSh
 Llama-3.1-8B ($\bar{\alpha}_{\text{IS}}=0.35$)								
 NVDA	680	837	+23.1	+0.4	142	142	0.0	.00
 MSFT	382	383	+0.3	+0.0	110	110	0.0	.00
 AAPL	339	730	+115.3	+0.33	100	100	0.0	.00
 NFLX	1164	1329	+14.1	+0.3	103	103	0.0	.00
 AMZN	533	424	−20.5	−0.7	114	114	0.0	.00
mean			+26.5	+0.7			0.0	.00
 Starling-7B ($\bar{\alpha}_{\text{IS}}=0.82$)								
 NVDA	872	87	−90.1	−.75	112	78	−30.3	−.98
 MSFT	416	151	−63.6	−.48	114	94	−18.0	−1.15
 AAPL	498	296	−40.5	−.18	109	105	−4.3	−.20
 NFLX	1394	304	−78.2	−.35	107	122	+13.6	+0.50
 AMZN	703	259	−63.2	−.34	101	72	−28.8	−1.34
mean			−67.1	−.42			−13.6	−.63
 DeepSeek-7B-Chat ($\bar{\alpha}_{\text{IS}}=1.06$)								
 NVDA	294	387	+31.7	+0.9	109	84	−23.4	−.97
 MSFT	295	108	−63.5	−.71	100	100	0.0	−.13
 AAPL	318	250	−21.3	−.19	94	97	+2.7	−.32
 NFLX	759	375	−50.6	−.30	105	94	−10.6	−.63
 AMZN	424	285	−32.8	−.16	82	101	+22.8	+1.19
mean			−27.3	−.25			−1.7	−.17

Table 8: Per-ticker IS / OOS single-stock backtest results for the three models not shown in Table 2. Daily rebalancing, $T=1.0$, seed 42.

mid/small mean 0.14; Qwen2.5-14B and Starling-7B both ≈ 0.9 ; DeepSeek-7B-Chat 0.11). Second, the within-model spread on mid/small entities is non-trivial (Phi-4 max−min 0.27; Qwen2.5-14B 0.80), so the calibrator continues to discriminate among individual mid- and small-cap names rather than collapsing to a uniform setting. The corresponding per-ticker IS+OOS backtests for all five models are in Table 10; for completeness this table includes the POOL and SAM rows already shown in the main paper for Phi-4-14B and Qwen2.5-14B alongside the remaining four mid/small tickers and the three other models. Per-ticker IS deltas are noisier on small-caps because of lower trade volumes, smaller baseline portfolios, and the cascade of even




Figure 3: Daily calibrated α across the IS period for every (model, ticker) pair. Faint scatter = per-day α ; solid line = 63-day rolling mean; per-panel mean is annotated. Rows = LLMs (top to bottom: Phi-4-14B, Qwen2.5-14B, Llama-3.1-8B, Starling-7B, DeepSeek-7B-Chat); columns = tickers (NVDA, MSFT, AAPL, NFLX, AMZN). α axis is shared across all panels (range 0–4).

Model	Mid-cap			Small-cap			mean
	KMX	POOL	EXPE	CROX	SAM	BJRI	
Phi-4-14B	0.27	0.15	0.08	0.00	0.14	0.17	0.14
Qwen2.5-14B	0.45	0.90	1.18	0.38	0.80	1.00	0.79
Llama-3.1-8B	0.68	0.69	0.55	0.79	0.75	0.43	0.65
Starling-7B	0.73	0.71	1.26	0.59	0.87	1.33	0.92
DeepSeek-7B	0.06	0.10	0.25	0.19	0.03	0.02	0.11

Table 9: Per-ticker mean calibrated $\bar{\alpha}_{IS}$ across mid- and small-cap entities, all five backtested models. Compare with the mega-cap grid in Table 3.

small α firings through portfolio state; mean rows in the bottom of each panel capture the per-cap-class signal with less noise.

D.3 Experiment 3: Cross-Model Roster

Table 11 lists the 11 instruction-tuned LLMs used in the cross-model alignment analysis (§4.3), with family, parameter count, and reported pre-training cutoff. All reported cutoffs fall before January 2025, so the 2025 evaluation window follows every model’s reported cutoff.

E Per-Model Equity Curves

Figure 2 in the main paper visualises the IS→OOS alignment shift induced by FINCAD via a per-model scatter of (IS, OOS) Sharpe pairs. Figure 4 complements this with the underlying in-sample equity curves for all ten models in the cross-model analysis. Within each panel, the faded line is baseline decoding and the bold line is FINCAD. The gap between the two lines is a visual proxy for the look-ahead bias FINCAD removes: it is near-zero for Phi-4, Llama-3.1, and Yi-1.5, but substantial for Qwen2.5-14B, Qwen2.5-7B, Starling-7B, and Falcon3-7B.

F Subset-Size Sensitivity

Table 12 reports the exhaustive subset analysis for Sharpe ratio across all subset sizes $k \in \{5, \dots, 11\}$ from the $n=11$ models. The observed difference $\Delta\rho = \bar{\rho}_C - \bar{\rho}_B$ remains between +0.06 and +0.07 across these subset sizes. Because subsets at a given k overlap, this analysis measures sensitivity to model-pool composition rather than sampling uncertainty over independent observations.




Figure 4: In-sample SPY equity curves (2010–2020, daily) for all ten models. Faded: baseline decoding; bold: FinCAD. Curves normalised to 1.0 at the start date.

G Artifact Licences, Provenance, and Attribution

This section records the provenance, declared licence status, and experimental use of the principal datasets, codebases, and model weights. Dataset and model licences are reported from their official repository or model cards at the time of access.

Price data (FINSABER). The 2010–2025 daily price panel and the survivorship-, leakage-, and look-ahead-aware data conventions used throughout Experiments 2 and 3 are sourced from the FINSABER-V2 data release³ (Li et al., 2026a). Its dataset card declares the licence as **other** and notes that users remain responsible for the terms of the underlying data providers. We use the data for academic evaluation and direct users to the upstream release; no price data are redistributed with FINCAD.

³<https://huggingface.co/datasets/finsaber-team/FINSABER-V2-Data>

General benchmark datasets. Experiment 1 uses openai/gsm8k, TIGER-Lab/MMLU-Pro, HuggingFaceH4/MATH-500, openai/openai_humaneval, and FinEval through their Hugging Face releases. The original Salesforce/FinEval dataset card⁴ declares the **CC BY-NC 4.0 License** and notes that users remain responsible for the terms of its constituent source datasets. For reproducibility, we release the 20 multiple-choice subsets converted into the unified schema used by our evaluator. This derived release⁵ does not supersede the upstream licence or source-specific terms. The dataset cards for GSM8K, MMLU-Pro, and HumanEval declare the **MIT License**. The MATH-500 card does not declare a licence and identifies OpenAI’s MIT-licensed PRM800K repository as its immediate source. Except for the separate

⁴<https://huggingface.co/datasets/Salesforce/FinEval>

⁵<https://huggingface.co/datasets/waylonli/fineval-processed>

Tkr	IS (2010–2020)				OOS (2025)			
	B\$K	C\$K	$\Delta\%$	ΔSh	B\$K	C\$K	$\Delta\%$	ΔSh
 <i>Phi-4-14B (mid+small mean $\bar{\alpha}_{IS}=0.14$)</i>								
KMX	201	163	-18.8	-.07	44	47	+5.5	+.23
POOL	901	928	+3.0	+0.0	67	74	+9.8	+.47
EXPE	165	206	+24.8	+0.6	149	149	0.0	.00
CROX	182	302	+65.4	+1.2	59	73	+23.4	+.60
SAM	692	562	-18.8	-.07	55	64	+15.7	+.84
BJRI	132	225	+70.1	+1.6	114	109	-4.4	-.11
mean			+21.0	+0.3			+8.3	+.34
 <i>Qwen2.5-14B (mid+small mean $\bar{\alpha}_{IS}=0.79$)</i>								
KMX	226	282	+24.7	+0.8	55	64	+17.2	+.42
POOL	1027	753	-26.7	-.09	81	73	-9.7	-.33
EXPE	135	103	-24.0	-.11	140	136	-3.4	-.05
CROX	606	912	+50.6	+1.0	58	98	+69.7	+1.11
SAM	415	412	-0.9	+0.1	69	76	+10.2	+.49
BJRI	117	132	+12.8	+0.2	101	143	+42.0	+.96
mean			+6.1	+0.0			+21.0	+.43
 <i>Llama-3.1-8B (mid+small mean $\bar{\alpha}_{IS}=0.65$)</i>								
KMX	421	106	-74.9	-.51	51	51	0.0	.00
POOL	640	1017	+58.8	+1.8	63	63	0.0	.00
EXPE	198	184	-7.1	-.02	164	164	0.0	.00
CROX	311	321	+3.0	+0.1	63	79	+25.9	+.41
SAM	502	444	-11.6	-.06	67	70	+5.9	+.29
BJRI	146	137	-6.2	-.02	135	106	-21.4	-.58
mean			-6.3	-.07			+1.7	+.02
 <i>Starling-7B (mid+small mean $\bar{\alpha}_{IS}=0.92$)</i>								
KMX	90	53	-41.5	-.42	61	72	+18.5	-.34
POOL	737	192	-74.0	-.58	77	100	+30.7	+1.30
EXPE	113	18	-84.2	-.74	117	82	-30.4	-1.38
CROX	189	209	+10.8	.00	89	95	+6.4	+0.04
SAM	335	129	-61.6	-.37	73	100	+36.4	+1.59
BJRI	152	75	-50.7	-.36	88	106	+21.1	+.62
mean			-50.2	-.41			+13.8	+.31
 <i>DeepSeek-7B-Chat (mid+small mean $\bar{\alpha}_{IS}=0.11$)</i>								
KMX	97	120	+24.1	+1.2	75	82	+8.9	+.84
POOL	289	314	+8.5	+0.4	72	76	+6.2	+.32
EXPE	30	66	+124.3	+4.3	106	121	+14.6	+.53
CROX	111	96	-13.4	-.03	72	72	0.0	.00
SAM	117	271	+132.1	+4.1	71	75	+5.3	+.20
BJRI	101	26	-74.6	-.57	103	170	+64.5	+1.54
mean			+33.5	+0.7			+16.6	+.57

Table 10: Full mid- and small-cap per-ticker results for all five backtested models. The POOL and SAM rows for Phi-4-14B and Qwen2.5-14B are reproduced here from Table 2 of the main paper for completeness. Daily rebalancing, $T=1.0$, seed 42. Per-ticker Δ is noisier on small-caps due to lower trade volumes and the cascade of small α firings through portfolio state; mean rows summarise each panel.

preprocessed FinEval release described above, these datasets are accessed for evaluation only and are not redistributed with FINCAD.

Context-Aware Decoding. Our logit-blending update (Eq. 1) follows the Context-Aware Decoding formulation of Shi et al. (2024), whose reference implementation is publicly available⁶. We re-implement the decoding loop on top of Hug-

⁶<https://github.com/xhan77/context-aware-decoding>

Model	Family	Size	Cutoff
 Phi-4-14B	Phi	14B	Jun '24
 Llama-3.1-8B-Inst.	Llama	8B	Mar '24
 Qwen2.5-14B-Inst.	Qwen	14B	Jun '24
 Qwen2.5-7B-Inst.	Qwen	7B	Jun '24
 GLM-4-9B-Chat	GLM	9B	Mar '24
 Mistral-7B-v0.3	Mistral	7B	May '24
 Yi-1.5-9B-Chat	Yi	9B	May '24
 OLMo-2-7B-Inst.	OLMo	7B	Nov '24
 Starling-LM-7B	Starling	7B	Mar '24
 Falcon3-7B-Inst.	Falcon	7B	Sep '24
 DeepSeek-7B-Chat	DeepSeek	7B	Nov '23

Table 11: Models for cross-model validation (11 models from 10 families, cutoffs before Jan 2025).

k	$\binom{11}{k}$	$\bar{\rho}_B$	$\bar{\rho}_C$	$\Delta\rho$
5	462	+0.758	+0.818	+0.061
6	462	+0.770	+0.834	+0.064
7	330	+0.779	+0.846	+0.067
8	165	+0.786	+0.855	+0.069
9	55	+0.792	+0.862	+0.070
10	11	+0.796	+0.868	+0.072
11	1	+0.800	+0.873	+0.073

Table 12: Descriptive Sharpe-ratio ranking alignment (Spearman ρ) across subset sizes. B = baseline; C = FINCAD.

gingFace Transformers in order to support (i) the entity-adaptive penalty (§3.2), (ii) the adversarial prior-discovery pipeline (§3.1), and (iii) the adaptive-retry mechanism on JSON parse failure (Algorithm 1). The reference repository did not contain a licence file when accessed. We use the published formulation as prior work and do not copy source files from the repository.

FINCAD implementation. Our implementation is licensed under the **MIT License** and is released⁷. The repository includes the inference and discovery code, experiment configurations, and discovery JSON artifacts; datasets and model weights remain available from their respective providers.

Single-stock backtesting harness (ai-hedge-fund). The structured-summary prompt template, JSON action schema, and execution-at-next-open conventions used in Experiments 2 and 3 are adapted from the open-source ai-hedge-fund project⁸, released under the **MIT License**. We retain the agent-as-quantitative-analyst design but implement our own trading loop, a one-way

⁷<https://github.com/waylonli/FinCAD>

⁸<https://github.com/virattt/ai-hedge-fund>

transaction-cost model of 10 bps of trade notional (Appendix A), and the long-only single-stock constraint that defines the simulation envelope on top of HuggingFace Transformers and pandas.

Prompt optimisation (DSPy / MIPROv2). The adversarial prior-discovery pipeline (§3.1) uses MIPROv2 (Khattab et al., 2024) as released in DSPy under the **MIT License**. We use DSPy as a library without modification and write a project-specific signature, metric, and proposer configuration as described in §3.1.

Inference stack. Model inference uses HuggingFace Transformers (Apache License 2.0) and PyTorch (BSD-3-Clause License); ancillary data handling uses pandas (BSD-3-Clause).

Pre-trained model licenses. All eleven LLMs evaluated in §4.3 are accessed through HuggingFace under their providers' stated terms; we use the pre-trained weights for non-commercial academic research consistent with each licence. The majority (Qwen2.5-14B-Instruct, Qwen2.5-7B-Instruct, Mistral-7B-v0.3, Yi-1.5-9B-Chat, and OLMo-2-7B-Instruct) are released under the **Apache License 2.0**. Phi-4-14B is released under the **MIT License**. The remaining five models carry custom or community licences: Llama-3.1-8B-Instruct under the **Llama 3.1 Community Licence**, GLM-4-9B-Chat under the **GLM-4 Model Licence**, Falcon3-7B-Instruct under the **TII Falcon-LLM Licence 2.0**, DeepSeek-7B-Chat under the **DeepSeek Model Licence**, and Starling-LM-7B under **Apache 2.0 with a research-only restriction**. Licences are as reported on each model's HuggingFace card at the time of writing.