

QuantEval: A Benchmark for Financial Quantitative Tasks in Large Language Models

Anonymous ACL submission

Abstract

Large Language Models (LLMs) have shown strong capabilities across many domains, yet their evaluation in financial quantitative tasks remains fragmented and mostly limited to knowledge-centric question answering. We introduce **QuantEval**, a benchmark that evaluates LLMs across three essential dimensions of quantitative finance: (i) *knowledge-based QA*, (ii) *quantitative mathematical reasoning*, and (iii) *quantitative strategy coding*. Unlike prior financial benchmarks, QuantEval integrates a CTA-style backtesting framework that executes model-generated strategies and evaluates them using financial performance metrics, enabling a more realistic assessment of quantitative coding ability. We evaluate some state-of-the-art open-source and proprietary LLMs and observe substantial gaps to human experts, particularly in reasoning and strategy coding. Finally, we conduct large-scale supervised fine-tuning and reinforcement learning experiments on domain-aligned data, demonstrating consistent improvements. We hope QuantEval will facilitate research on LLMs’ quantitative finance capabilities and accelerate their practical adoption in real-world trading workflows. We additionally release the full deterministic backtesting configuration (asset universe, cost model, and metric definitions) to ensure strict reproducibility.

1 Introduction

Large Language Models (LLMs) have recently demonstrated remarkable capabilities in reasoning, code generation, and domain-specific question answering across diverse domains (Zhong et al., 2025; Shao et al., 2024; OpenAI, 2024). These abilities have motivated growing interest in deploying LLMs for financial applications, where practitioners routinely rely on quantitative analysis, mathematical modeling, and algorithmic strategy development.

However, evaluating LLMs in financial quantitative workflows remains limited and fragmented. Existing financial benchmarks mainly focus on knowledge-based question answering (QA) and numerical QA over documents or tables (Xie et al., 2023; Zhu et al., 2024; Chen et al., 2024, 2022; Xie et al., 2024). While these benchmarks are useful for assessing domain familiarity and factual understanding, they only partially reflect the competencies required in quantitative finance. In real-world quant workflows, practitioners must (i) perform precise multi-step quantitative calculations grounded in financial theory and market data, and (ii) translate hypotheses into executable trading strategies whose performance and risk characteristics can be validated via backtesting.

This gap between existing benchmarks and practical quant workflows limits the ability to diagnose and improve LLMs for quantitative finance. Quantitative tasks such as option pricing, portfolio optimization, risk assessment, and algorithmic trading strategy development require models to integrate domain knowledge with numerical precision and coding proficiency. Moreover, evaluating strategy generation cannot rely solely on syntax correctness: a strategy may compile but still be financially meaningless, unstable, or risk-dominant. Therefore, benchmark design for quantitative finance should include execution-based evaluation that measures strategy behavior under realistic backtesting environments.

To address these limitations, we introduce **QuantEval**, a benchmark designed to systematically evaluate LLMs on three key dimensions of quantitative finance: (1) *Knowledge-based QA*, measuring conceptual understanding and terminology; (2) *Quantitative Mathematical Reasoning*, assessing multi-step quantitative computations grounded in real market data and financial theory; and (3) *Quantitative Strategy Coding*, testing whether models can generate framework-compatible and exe-

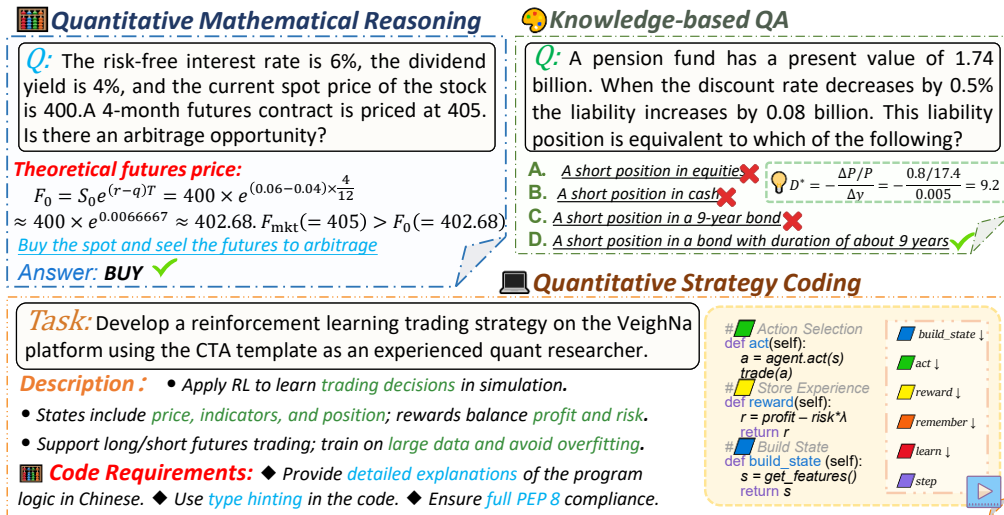


Figure 1: Examples from the three core tasks in QuantEval: knowledge-based question answering, quantitative mathematical reasoning, and quantitative strategy coding.

cutable trading strategies whose backtest outcomes match expert-validated reference implementations. QuantEval contains 1,575 carefully curated samples across these dimensions.

A distinctive feature of QuantEval is that the strategy coding task is evaluated in an integrated CTA-style backtesting framework. This enables rigorous assessment beyond pass/fail compilation, by measuring execution success rate and deviations in key financial metrics. We evaluate 13 open-source and proprietary LLMs on QuantEval and observe substantial performance gaps relative to human experts, especially for reasoning and strategy coding. Motivated by these findings, we further explore large-scale supervised fine-tuning and reinforcement learning methods to improve quantitative finance competency. We release the dataset, evaluation scripts, and a standardized CTA backtesting harness to enable reproducible evaluation.

In summary, our contributions are:

- We introduce **QuantEval**, a benchmark for evaluating LLMs on financial quantitative tasks spanning knowledge QA, quantitative reasoning, and execution-based quantitative strategy coding.
- We design a scalable construction pipeline combining expert annotation and multi-agent generation, and validate strategy coding via a CTA-style backtesting framework with finance-specific metrics.
- We benchmark 13 state-of-the-art LLMs and provide detailed analyses of failure modes, highlighting persistent challenges in quanti-

tative reasoning and strategy implementation.

- We conduct supervised fine-tuning and reinforcement learning experiments on domain-aligned data, demonstrating consistent improvements and providing practical training insights.

2 QuantEval

2.1 Overview

QuantEval is a benchmark designed to evaluate Large Language Models (LLMs) on core financial quantitative competencies. Compared with prior financial benchmarks that primarily emphasize knowledge-centric QA, QuantEval explicitly targets three dimensions that collectively define quantitative finance workflows:

- Knowledge-based QA:** conceptual understanding of financial definitions, terminology, and theory;
- Quantitative Mathematical Reasoning:** multi-step quantitative problem solving grounded in real market data and financial formulas;
- Quantitative Strategy Coding:** generating executable trading strategies under a unified CTA-style backtesting framework and matching expert-validated performance metrics.

QuantEval contains 1,575 samples: 660 Knowledge QA, 855 Reasoning, and 60 Strategy Coding instances. Each dimension requires a distinct construction methodology to ensure quality, diversity, and practical relevance (Figure 2).

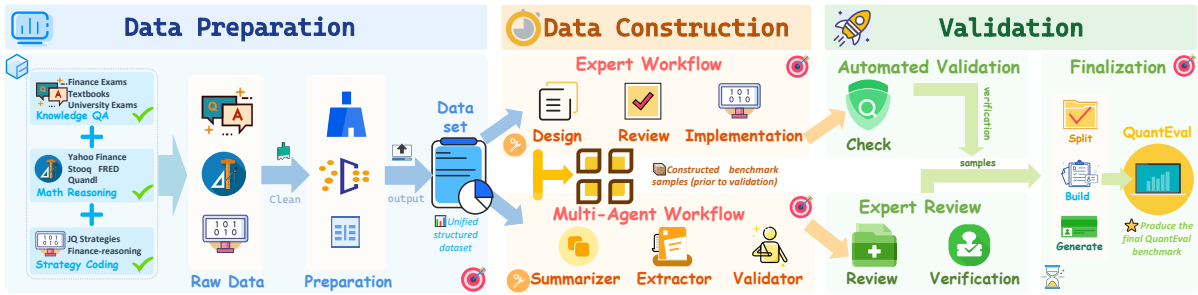


Figure 2: The QuantEval data construction pipeline consists of three main stages: Data Preparation, Data Construction, and Validation.

2.2 Data Preparation

The initial stage focuses on collecting and organizing foundational materials that serve as the bedrock for subsequent question and code generation. Given the heterogeneity of the three task categories. Detailed data sources, time spans, and licenses for each task category are summarized in Appendix A.

Knowledge-based QA. For this dimension, authoritative textual materials are sourced from financial textbooks, peer-reviewed academic papers, regulatory filings, and curated glossaries of financial terminology. Domain experts specializing in quantitative finance meticulously review these materials, selecting passages that cover essential concepts, theoretical frameworks, and domain-specific vocabulary. To ensure comprehensive coverage and mitigate potential biases, the materials span multiple subdomains.

Quantitative Mathematical Reasoning. This dimension requires datasets rooted in realistic financial computations and problem-solving scenarios. To this end, extensive historical market data and macroeconomic indicators are gathered from reputable public sources. Complementing these data, problem templates are extracted from quantitative finance coursework, professional certification materials (e.g., CFA, FRM), and industry whitepapers. These templates are enriched with real market data to create contextually rich problems that require calculations involving option pricing models, risk metrics, portfolio optimization, and other quantitative finance techniques.

Quantitative Strategy Coding. For the quantitative strategy coding dimension, algorithmic trading strategies and backtesting scripts are collected from a variety of sources, including open-source repositories, academic publications. These code-bases encompass a wide spectrum of quantitative

approaches. Financial experts curate and annotate these code snippets to ensure clarity, correctness, and alignment with realistic trading scenarios.

CTA Strategies about Quantitative Strategy Coding. The CTA strategies incorporated have been designed and validated by domain experts with extensive professional experience and have been actively employed in offline industry settings. Alongside the code, metadata including strategy parameters, historical performance metrics, and market conditions are collected to support comprehensive evaluation and benchmarking. In this benchmark, we focus on equity CTA-style backtesting, where strategies trade a fixed universe of U.S. ETFs and large-cap stocks under standardized execution and risk constraints. A full specification of the CTA backtesting environment is provided in Appendix B.

2.3 Data Construction

The second stage focuses on generating the actual benchmark samples through a hybrid approach that combines expert annotation with automated multi-agent generation. This dual approach balances the need for high-quality, domain-accurate content with scalability and diversity.

Generation with Experts. Financial domain experts play a pivotal role in crafting high-quality samples. Their responsibilities include revising and refining existing questions to enhance clarity, relevance, and difficulty, ensuring that distractors are plausible yet unambiguously incorrect, and that questions demand genuine understanding rather than superficial recall. Experts also design novel questions and problems based on curated texts, market data, and synthetic scenarios.

For Quantitative Strategy Coding tasks, experts specify challenges such as strategy implementation, partial code completion, debugging, and parame-

ter tuning, ensuring that the tasks reflect realistic quantitative trading workflows. Additionally, experts review and annotate quantitative strategy code samples to verify correctness, readability, and alignment with benchmark objectives, providing detailed explanations and expected outputs to facilitate evaluation. To further enhance the dataset’s complexity and practical relevance, synthetic quantitative scenarios are designed by combining real market data with hypothetical events or constraints. These scenarios simulate market shocks, regulatory changes, portfolio rebalancing under risk limits, and other realistic financial decision-making contexts.

Multi-Agent Automated Construction. Inspired by expert workflows, a multi-agent automated construction framework is developed to efficiently scale data generation, particularly for mathematical reasoning and coding tasks. This framework comprises specialized roles that collaborate to produce high-quality samples.

1) A summarizer agent analyzes financial texts and market data to generate concise summaries. 2) A problem extractor agent identifies and extracts relevant data segments and problem templates suitable for question or code generation. 3) The question and code generator agent produces mathematically rigorous questions and executable code snippets. This agent operates under detailed prompt templates and is guided by examples derived from expert annotations to ensure adherence to domain standards. 4) Finally, a validation agent performs preliminary checks on generated samples, including syntactic correctness, logical consistency, and alignment with task requirements.

This multi-agent system enables efficient, large-scale generation of diverse and high-quality benchmark samples while maintaining domain relevance and rigor. We provide a breakdown of expert-generated vs agent-generated samples and their quality comparison in Appendix D.3.

2.4 Validation

Ensuring the reliability, rigor, and practical relevance of QuantEval samples is of paramount importance. To this end, a multi-tiered validation process involving both automated agents and domain experts is implemented.

The automated validation agent conducts initial screening to detect duplicate or near-duplicate questions, inconsistencies, and trivial or overly simplistic items. For knowledge QA and reasoning

tasks, it verifies that questions cannot be answered correctly through superficial keyword matching or simple heuristics, thereby ensuring that genuine understanding and reasoning are required. For coding tasks, the agent checks code executability, correctness of outputs, and adherence to specified constraints, including successful execution within the integrated CTA backtesting framework. Samples failing these criteria are either revised or discarded to maintain dataset quality.

Domain experts perform thorough reviews of all benchmark samples. Their evaluation focuses on confirming factual accuracy, mathematical soundness, and code validity. Experts ensure that questions and tasks are unambiguous, well-phrased, and aligned with the intended difficulty levels. They also conduct ethical and sensitivity checks to verify that samples do not contain proprietary, sensitive, or ethically problematic content. When necessary, experts request revisions or reject samples that do not meet stringent quality standards.

This iterative feedback loop between validation and construction stages continues until the dataset achieves the desired quality threshold, ensuring that QuantEval is both challenging and trustworthy.

2.5 De-duplication, leakage prevention and difficulty control

To ensure benchmark reliability and minimize training-data contamination, we apply strict de-duplication and leakage prevention procedures at both the question and code levels. We additionally control task difficulty and enforce balanced coverage across quantitative finance subdomains. Full details of our de-duplication criteria, overlap scanning protocols, and difficulty/coverage design are provided in Appendix C.

2.6 Implementation Details

The data construction pipeline extensively leverages state-of-the-art LLMs, primarily GPT-4o and Claude-4.5-sonnet, for generation and validation tasks.

A key feature of the quantitative strategy coding dimension is the integration of a CTA backtesting framework. This framework enables the generation of realistic trading scenarios and the validation of code correctness through simulated executions on historical market data. The backtesting environment supports evaluation of strategy performance metrics, risk characteristics, and robustness under varying market conditions. By grounding the

Metric in Table 2	Goal	Definition / Interpretation
Executable Rate (%)	↑	Fraction of generated strategy code that can be successfully executed end-to-end in the CTA backtesting framework (with correct interfaces and no runtime failure).
Return MAE	↓	Mean absolute error between the model-generated strategy’s annualized return and the expert-validated GT return, computed only over executable instances.
Drawdown MAE	↓	Mean absolute error between the model-generated strategy’s maximum drawdown and GT drawdown, capturing deviation in downside risk.
Sharpe MAE	↓	Mean absolute error between the model-generated strategy’s Sharpe ratio and GT Sharpe ratio, measuring deviation in risk-adjusted return.
Ret/Draw MAE	↓	Mean absolute error between the model-generated return-to-drawdown ratio and GT, capturing deviation in robustness (profitability vs. downside risk).

Table 1: **Metrics used for Strategy Coding evaluation in Table 2.** Executable Rate is computed over all coding instances. MAE values are computed only over executable outputs; when a model produces no executable strategies, MAE is undefined and reported as “-”. Return/Drawdown MAE are reported in percentage points (pp), while Sharpe MAE and Ret/Draw MAE are unitless.

benchmark in this industry-recognized methodology, QuantEval ensures that the coding tasks reflect authentic quantitative finance workflows.

2.7 Data Statistics

Our dataset comprises a total of 1,575 samples, systematically organized into three primary categories: *Knowledge-based QA* (660 samples), *Quantitative Mathematical Reasoning* (855 samples), and *Quantitative Strategy Coding* (60 samples).

For the *Knowledge-based QA* and *Quantitative Reasoning* tasks, the dataset includes both multiple-choice questions and fill-in-the-blank formats.

In the *Quantitative Strategy Coding* category, each sample provides a detailed problem description alongside a widely recognized, expert-validated trading strategy implemented in code. Additionally, the dataset includes the corresponding performance metrics of these strategies, enabling rigorous evaluation of generated code against trusted benchmarks.

3 Evaluation on QuantEval

This section presents the evaluation protocol and main results of QuantEval across three tasks: Knowledge-based QA, Quantitative Mathematical Reasoning, and Quantitative Strategy Coding. We emphasize two principles: (i) prompting consistency for fair comparison across models and set-

tings, and (ii) execution-based evaluation for quantitative strategy coding, where correctness must be validated via backtesting rather than surface-level code inspection.

3.1 Evaluation Setup

We evaluate 13 state-of-the-art LLMs spanning open-source and proprietary families, including Qwen3 variants (4B, 8B, 14B, 30B) (Yang et al., 2025), DeepSeek distilled models (DeepSeek-AI et al., 2025), DianJin-R1-7B (Zhu et al., 2025), Claude-4.5-sonnet (Anthropic, 2024), Gemini-2.5-pro, and GPT-5 (OpenAI, 2025). All models are evaluated using greedy decoding with temperature 0 and top-p 1.0. We accessed proprietary models via official APIs between Sep 2025 and Dec 2025.

Knowledge QA & Quantitative Reasoning. For Knowledge QA and Quantitative Reasoning tasks, we report **accuracy** (%), defined as the percentage of correctly answered questions. We evaluate each model in two prompting settings: (i) **with Chain-of-Thought prompting (w/ CoT)**, which encourages explicit step-by-step reasoning; and (ii) **without CoT (w/o CoT)**, where the model answers directly. For multiple-choice questions, we extract the final option from the model output; for open-ended questions, we apply normalization and semantic matching. We include prompt templates and evaluation prompts in Appendix E.2.

Strategy Coding (CoT-only). For Strategy Coding, the model must generate executable quantitative trading strategy code under a unified CTA-style backtesting framework. In pilot tests, direct-answer prompting frequently produces incomplete code with near-zero executability; hence we evaluate coding under CoT prompting only. Each generated strategy is executed end-to-end on aligned historical market data, and we compute four backtesting metrics: **Annualized Return**, **Maximum Drawdown**, **Sharpe Ratio**, and **Return-to-Drawdown Ratio**. Definitions of all coding evaluation metrics and their interpretations are summarized in Table 1.

We primarily evaluate Strategy Coding under CoT prompting to obtain non-trivial executability. However, to ensure fairness, we additionally run a controlled ablation on three representative models comparing CoT vs direct code-only prompting. We find that CoT significantly improves compliance and execution success. We therefore report the main leaderboard under CoT-only, and provide the prompt ablation results in Appendix E.3.

Model	Knowledge QA (Acc. %)			Reasoning (Acc. %)			Strategy Coding (CoT-only)				
	w/	w/o	Δ	w/	w/o	Δ	Exec. Rate \uparrow	Return MAE \downarrow	Drawdown MAE \downarrow	Sharpe MAE \downarrow	Ret/Draw MAE \downarrow
Human Expert	91.75	91.75	0.0	89.05	89.05	0.0	100	1.52	1.10	0.07	0.09
Open-Source Models											
Qwen3-4B	64.4	69.0	-4.6	27.7	33.9	-6.2	0.0	-	-	-	-
Qwen3-8B	65.5	65.5	0.0	33.9	29.2	+4.7	0.0	-	-	-	-
Qwen3-14B	75.9	78.2	-2.3	35.4	33.9	+1.5	0.0	-	-	-	-
Qwen3-30B-A3B	48.5	44.5	+4.0	40.5	36.5	+4.0	8.3	2.71	20.01	0.30	0.40
DeepSeek-R1-Distill-Qwen-1.5B	37.2	34.0	+3.2	30.0	27.9	+2.1	0.0	-	-	-	-
DeepSeek-R1-Distill-Qwen-7B	42.5	49.4	-6.9	21.5	26.2	-4.7	0.0	-	-	-	-
DeepSeek-R1-Distill-Qwen-14B	66.7	63.2	+3.5	29.2	32.3	-3.1	0.0	-	-	-	-
DeepSeek-R1-Distill-Llama-8B	50.6	41.4	+9.2	10.8	10.4	+0.4	0.0	-	-	-	-
Deepseek-R1-671B	52.7	48.1	+4.6	44.0	40.8	+3.2	0.0	-	-	-	-
DianJin-R1-7B	50.3	46.0	+4.3	42.8	38.0	+4.8	0.0	-	-	-	-
Proprietary Models											
Claude-4.5-sonnet	86.0	90.8	-4.8	43.1	44.6	-1.5	63.3	10.14	12.39	0.17	0.21
Gemini-2.5-pro	66.7	66.7	0.0	38.0	23.0	+15.0	63.3	12.57	15.15	0.16	0.25
GPT-5	83.9	82.8	+1.1	55.0	33.8	+21.2	51.7	7.72	9.62	0.18	0.18

Table 2: Unified evaluation on **QuantEval**. For Knowledge QA and Reasoning, we report accuracy (%) with Chain-of-Thought prompting (w/), without CoT (w/o), and the performance gap $\Delta = w/ - w/o$. For Strategy Coding, we evaluate under CoT prompting and report Executable Rate (%) and MAE of four backtesting metrics. Return MAE and Drawdown MAE are reported in **percentage points (pp)**; models with 0% executable rate are marked as “-” for MAE.

Human Baseline. Human expert performance serves as an approximate upper bound. Detailed expert accuracies and expert background information are reported in Appendix D.

3.2 Main Results on QuantEval

Table 2 reports unified evaluation results on QuantEval. Several key findings emerge.

A substantial gap remains to human experts. Human experts achieve strong performance on both Knowledge QA and Reasoning. In contrast, the best-performing LLMs remain behind, especially on quantitative reasoning. This confirms that QuantEval requires not only surface-level financial familiarity but also precise multi-step quantitative computation and robust reasoning.

CoT prompting helps reasoning but is not always beneficial. Across models, CoT yields the largest improvements on quantitative reasoning. However, CoT is not uniformly beneficial and can even reduce accuracy. This suggests that CoT primarily benefits tasks requiring structured computation, but may induce overthinking or hallucinated intermediate reasoning for factual recall.

Strategy coding remains the hardest dimension. Strategy Coding is substantially more challenging than QA-style tasks. Most open-source models fail to generate executable strategies under strict CTA framework constraints, resulting in 0% executable rate across many models. Even for large open-source models, performance remains limited. In

Training Stage	Knowledge QA	Reasoning	Strategy Coding (Sharpe MAE) \downarrow
Pre-training (Base)	50.0	42.0	-
After SFT	53.2	47.8	0.85
After SFT + GRPO RL	52.2	48.4	0.72

Table 3: Performance progression of DianJin-R1-7B under extended training. Accuracy (%) is reported for QA and reasoning. For strategy coding, we report **Sharpe MAE** (lower is better).

contrast, proprietary models substantially improve executability.

Closed-source models lead overall, but open-source models show partial strengths. Closed-source models dominate across all tasks. Open-source models show relatively competitive performance on knowledge QA, suggesting that factual knowledge and basic finance concepts can be captured by open models, whereas execution-based strategy coding remains a major bottleneck.

4 Analysis and Discussion

In this section, we provide deeper analysis beyond the overall leaderboard in Table 2. We aim to answer three questions: (i) How QuantEval differs from existing financial benchmarks and why it is harder, (ii) Whether improvements on standard finance QA datasets translate to QuantEval, and (iii) What training strategies effectively improve quantitative reasoning and strategy coding under execution-based evaluation.

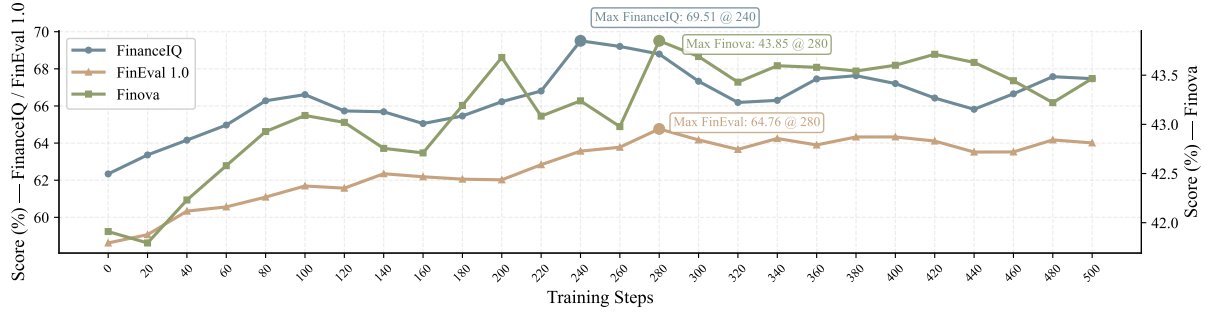


Figure 3: Performance curves of DianJin-R1-7B on test sets during reinforcement learning. Curves show steady improvement and stabilization, indicating effective reward-guided optimization.

4.1 Exploratory Training with Extended Financial Data

To explore effective approaches for improving LLM performance on QuantEval, we curate an extended training corpus of approximately 57,000 samples. It includes three large-scale open-source financial datasets: Agentar-DeepFinance-100K (Zhao et al., 2025), DianJin-R1-Data (Zhu et al., 2025), and FinQA (Qian et al., 2025), plus an additional 7,000 samples constructed using our pipeline to match QuantEval formats.

We perform supervised fine-tuning (SFT) on DianJin-R1-7B as a cold start, followed by reinforcement learning using Group Relative Policy Optimization (GRPO) (Shao et al., 2024). We incorporate two reward components: a *format reward* encouraging structured outputs and valid answer formatting, and an *accuracy reward* promoting correctness.

Table 3 reports the performance progression. SFT and GRPO provides further gains in QuantEval reasoning. These gains demonstrate that domain-aligned training improves reasoning performance, but the remaining gap to human experts (89.05%) remains substantial.

Figure 3 shows performance curves during GRPO reinforcement learning. We observe that benchmark performance improves steadily in the early phase and stabilizes later, indicating that the reward design effectively guides the model toward better structured reasoning and higher correctness. This supports the effectiveness of combining domain-aligned SFT with structured RL for finance reasoning tasks.

4.2 Comparison with Existing Financial Benchmarks

We benchmark representative models on three widely used financial evaluation datasets: FinancelQ (Duxiaoman-DI), FinEval (Guo et al., 2024),

Model	Finance IQ	Fin Eval	Finova	QuantEval QA	QuantEval Reasoning
Qwen2.5-7B-Instruct	57.4	52.6	30.6	38.7	30.9
DianJin-R1-7B	62.3	58.7	41.9	50.0	42.0
DianJin-R1-7B + RL	69.2	64.2	43.7	51.8	44.1
DianJin-R1-7B + SFT	72.2	70.5	44.8	53.2	47.8
DianJin-R1-7B + SFT + RL	77.9	72.7	45.5	52.2	48.4

Table 4: Cross-benchmark evaluation of selected models on existing financial evaluation datasets and QuantEval. Accuracy (%) is reported for QA and reasoning tasks. QuantEval coding task is excluded here for comparability. QuantEval columns are highlighted to emphasize the larger difficulty gap relative to existing benchmarks.

and Finova (Ant Group), and compare performance against QuantEval.

Table 4 shows that models generally achieve substantially higher accuracy on existing financial benchmarks than on QuantEval, especially on quantitative reasoning. This suggests that strong performance on standard financial QA datasets does not directly translate to solving realistic quantitative finance tasks requiring deeper computation.

QuantEval is harder due to reasoning depth and execution constraints. The gap between existing benchmarks and QuantEval is particularly pronounced on quantitative reasoning. This can be explained by two aspects: (i) QuantEval reasoning questions require multi-step computations grounded in realistic financial settings (e.g., combining volatility estimation with option pricing), rather than single-hop numerical extraction; (ii) QuantEval strategy coding requires generating executable code under strict CTA interfaces and producing financially meaningful metrics. These constraints reflect real quant workflows and substantially raise evaluation difficulty.

4.3 Error Analysis

We further analyze model outputs to identify common failure modes across tasks.

Knowledge QA: conceptual confusion and hallucinated explanations. Most errors stem from confusing highly related finance concepts. CoT sometimes worsens these errors because intermediate reasoning chains introduce hallucinated assumptions, consistent with the negative Δ observed for some models in Table 2.

Quantitative reasoning: arithmetic mistakes and formula composition failures. Reasoning questions require multi-step computations and careful variable tracking. Models often identify relevant formulas but make arithmetic errors or lose unit consistency across steps. This explains why reasoning accuracy remains below 60% even for GPT-5 and far below human experts.

Strategy coding: interface mismatch and execution failures dominate. To better understand why models fail on strategy coding, we analyze non-executable outputs and categorize failures into five types: (i) Syntax errors (code does not compile), (ii) Interface mismatch (missing required class/function signatures), (iii) Framework/API misuse (invalid calls or unsupported operations), (iv) Runtime failures (exceptions during execution such as shape mismatch or NaN propagation), and (v) Logic violations (code runs but violates task constraints). We find that most failures in open-source models are dominated by interface mismatch and framework misuse. Even for proprietary models with high executability, runtime failures and logic violations remain non-trivial, which explains why MAE values remain far from human references.

5 Related Work

Financial Question Answering and Knowledge Benchmarks Early and mainstream financial NLP benchmarks primarily target knowledge retrieval and fact-based QA, and focus on numerical reasoning over financial reports and tables (Chen et al., 2022; Zhu et al., 2021), requiring models to extract and compute answers from structured and semi-structured data. These datasets emphasize multi-hop numerical reasoning and compositional question answering but are limited to textual and tabular data, lacking evaluation of executable code generation or complex strategy formulation. Other benchmarks, like the CFA Institute’s practice question sets, assess models’ understanding of financial terminology, sentiment, and conceptual knowledge. While valuable for testing domain familiarity, these

datasets do not capture the quantitative and algorithmic reasoning central to quantitative finance.

Quantitative Reasoning in Finance and Beyond Quantitative reasoning benchmarks have been developed in broader contexts. Some evaluate multi-step mathematical problem solving (Amini et al., 2019), but their problems are general and not tailored to financial contexts. More recently, some attempts to bridge this gap by focusing on financial mathematics problems (finmath, 2024), yet it remains limited in scale and diversity, and does not incorporate coding or strategy development tasks. In the financial domain, quantitative reasoning often involves complex calculations such as option pricing, risk metrics, and portfolio optimization. Existing datasets rarely integrate real market data or simulate realistic financial scenarios, which limits their practical relevance.

Code Generation and Algorithmic Trading Benchmarks The ability of LLMs to generate executable code has been explored extensively in general programming domains. However, these datasets focus on general-purpose programming tasks and do not capture domain-specific requirements of financial quantitative strategy coding. In quantitative finance, algorithmic trading strategies are often implemented in specialized languages or frameworks, requiring precise mathematical modeling and adherence to market constraints. Some recent works have explored LLMs for financial code generation (Huynh and Lin, 2025; Li et al., 2024), but lack standardized benchmarks that evaluate the financial soundness and backtested performance of generated strategies.

6 Conclusion

We introduce **QuantEval**, a benchmark for evaluating LLMs on financial quantitative tasks spanning **knowledge QA**, **quantitative reasoning**, and **strategy coding**. Unlike prior finance benchmarks, QuantEval integrates an execution-based CTA backtesting framework to assess generated strategies using standardized risk–return metrics. Evaluations on 13 open-source and proprietary models reveal large gaps to human experts, especially in multi-step reasoning and framework-constrained coding, while domain-aligned SFT and RL yield consistent improvements. We release the dataset, evaluation scripts, and deterministic backtesting configuration to support reproducible research on LLMs for quantitative finance.

619 Limitations

620 QuantEval has several limitations. First, although
621 we include diverse quantitative finance domains,
622 the benchmark currently focuses on English-only
623 instances, which may limit its evaluation cover-
624 age for multilingual financial markets. Second,
625 the strategy coding task is evaluated under a CTA-
626 style backtesting framework with fixed assump-
627 tions (e.g., transaction costs, slippage, and asset
628 universe), and different backtesting configurations
629 may affect the reported metrics. Third, QuantE-
630 val contains 60 strategy coding instances, which is
631 smaller than QA and reasoning subsets due to the
632 high cost of expert validation; scaling coding tasks
633 is an important direction for future work.

634 Ethics Statement

635 QuantEval is constructed from publicly available
636 financial texts, market data sources, and expert-
637 validated strategy templates. We ensure that no
638 proprietary trading signals, private datasets, or con-
639 fidential strategies are included. Strategy coding
640 tasks are evaluated under offline historical backtest-
641 ing and are intended solely for research purposes;
642 they do not constitute financial advice and should
643 not be used for direct trading decisions. We also
644 include expert review to filter sensitive content and
645 reduce potential harms from misuse.

646 References

647 Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik
648 Koncel-Kedziorski, Yejin Choi, and Hannaneh Ha-
649 jishirzi. 2019. *MathQA: Towards interpretable math*
650 *word problem solving with operation-based for-*
651 *malisms*. In *Proceedings of the 2019 Conference*
652 *of the North American Chapter of the Association for*
653 *Computational Linguistics: Human Language Tech-*
654 *nologies, Volume 1 (Long and Short Papers)*, pages
655 2357–2367, Minneapolis, Minnesota. Association for
656 Computational Linguistics.

657 Ant Group. *Finova: A financial open-source platform*.
658 <https://github.com/antgroup/Finova>.

659 Anthropic. 2024. *Claude 3 model card. Technical Re-*
660 *port Version 1.0*, Anthropic. Accessed: 2025-09-16.

661 Jian Chen, Peilin Zhou, Yining Hua, Loh Xin, Kehui
662 Chen, Ziyuan Li, Bing Zhu, and Junwei Liang. 2024.
663 *FinTextQA: A dataset for long-form financial ques-*
664 *tion answering*. In *Proceedings of the 62nd Annual*
665 *Meeting of the Association for Computational Lin-*
666 *guistics (Volume 1: Long Papers)*, pages 6025–6047,
667 Bangkok, Thailand. Association for Computational
668 Linguistics.

Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena
669 Shah, Iana Borova, Dylan Langdon, Reema Moussa,
670 Matt Beane, Ting-Hao Huang, Bryan Routledge, and
671 William Yang Wang. 2022. *Finqa: A dataset of*
672 *numerical reasoning over financial data*. *Preprint*,
673 arXiv:2109.00122. 674

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang,
675 Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
676 Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang,
677 Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhi-
678 hong Shao, Zhuoshu Li, Ziyi Gao, and 181 others.
679 2025. *Deepseek-r1: Incentivizing reasoning capa-*
680 *bility in llms via reinforcement learning*. *Preprint*,
681 arXiv:2501.12948. 682

Duxiaoman-DI. *Financeiq dataset*. [https://huggingface.co/datasets/Duxiaoman-DI/](https://huggingface.co/datasets/Duxiaoman-DI/FinanceIQ)
683 [FinanceIQ](https://huggingface.co/datasets/Duxiaoman-DI/FinanceIQ). 684 685

finmath. 2024. *finmath-lib*. 686

Xin Guo, Haotian Xia, Zhaowei Liu, Hanyang Cao, Zhi
687 Yang, Zhiqiang Liu, Sizhe Wang, Jinyi Niu, Chuqi
688 Wang, Yanhui Wang, Xiaolong Liang, Xiaoming
689 Huang, Bing Zhu, Zhongyu Wei, Yun Chen, Weining
690 Shen, and Liwen Zhang. 2024. *Fineval: A chinese fi-*
691 *nanancial domain knowledge evaluation benchmark for*
692 *large language models*. *Preprint*, arXiv:2308.09975. 693

Nam Huynh and Beiyu Lin. 2025. *Large language mod-*
694 *els for code generation: A comprehensive survey of*
695 *challenges, techniques, evaluation, and applications*.
696 *Preprint*, arXiv:2503.01245. 697

Yinheng Li, Shaofei Wang, Han Ding, and Hang Chen.
698 2024. *Large language models in finance: A survey*.
699 *Preprint*, arXiv:2311.10723. 700

OpenAI. 2024. *Gpt-4o system card*. *Preprint*,
701 arXiv:2410.21276. 702

OpenAI. 2025. *Gpt-5 technical report*. [https://cdn.](https://cdn.openai.com/gpt-5-system-card.pdf)
703 [openai.com/gpt-5-system-card.pdf](https://cdn.openai.com/gpt-5-system-card.pdf). Accessed
704 September 24, 2025. 705

Lingfei Qian, Weipeng Zhou, Yan Wang, Xueqing Peng,
706 Han Yi, Yilun Zhao, Jimin Huang, Qianqian Xie, and
707 Jian yun Nie. 2025. *Finol: On the transferability of*
708 *reasoning-enhanced llms and reinforcement learning*
709 *to finance*. *Preprint*, arXiv:2502.08127. 710

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu,
711 Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan
712 Zhang, Y. K. Li, Y. Wu, and Daya Guo. 2024.
713 *Deepseekmath: Pushing the limits of mathemati-*
714 *cal reasoning in open language models*. *Preprint*,
715 arXiv:2402.03300. 716

Qianqian Xie, Weiguang Han, Zhengyu Chen, Ruoyu
717 Xiang, Xiao Zhang, Yueru He, Mengxi Xiao, Dong
718 Li, Yongfu Dai, Duanyu Feng, Yijing Xu, Haoqiang
719 Kang, Ziyang Kuang, Chenhan Yuan, Kailai Yang,
720 Zheheng Luo, Tianlin Zhang, Zhiwei Liu, Guojun
721 Xiong, and 15 others. 2024. *Finben: A holistic finan-*
722 *cial benchmark for large language models*. *Preprint*,
723 arXiv:2402.12659. 724

725 Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao
726 Lai, Min Peng, Alejandro Lopez-Lira, and Jimin
727 Huang. 2023. [Pixiu: A large language model, in-
728 struction data and evaluation benchmark for finance.
729 Preprint](#), arXiv:2306.05443.

730 An Yang, Anfeng Li, Baosong Yang, Beichen Zhang,
731 Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao,
732 Chengen Huang, Chenxu Lv, Chujie Zheng, Day-
733 iheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao
734 Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41
735 others. 2025. [Qwen3 technical report](#). *Preprint*,
736 arXiv:2505.09388.

737 Xiaoke Zhao, Zhaowen Zhou, Lin Chen, Lihong Wang,
738 Zhiyi Huang, Kaiyuan Zheng, Yanjun Zheng, Xiyang
739 Du, Longfei Liao, Jiawei Liu, Xiang Qi, Bo Zhang,
740 Peng Zhang, Wei Wang, and Zhe Li. 2025. [Agentar-
741 deepfinance-100k: A large-scale financial dataset via
742 systematic chain-of-thought synthesis optimization.
743 Preprint](#), arXiv:2507.12901.

744 Tianyang Zhong, Zhengliang Liu, Yi Pan, Yutong
745 Zhang, Yifan Zhou, Shizhe Liang, Zihao Wu, Yan-
746 jun Lyu, Peng Shu, Xiaowei Yu, Chao Cao, Hanqi
747 Jiang, Hanxu Chen, Yiwei Li, Junhao Chen, Huawen
748 Hu, Yiheng Liu, Huaqin Zhao, Shaochen Xu, and 56
749 others. 2025. [Evaluation of openai o1: Opportunities
750 and challenges of agi](#). *Preprint*, arXiv:2409.18486.

751 Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao
752 Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and
753 Tat-Seng Chua. 2021. [Tat-qa: A question answering
754 benchmark on a hybrid of tabular and textual content
755 in finance](#). *Preprint*, arXiv:2105.07624.

756 Jie Zhu, Qian Chen, Huaixia Dou, Junhui Li, Lifan
757 Guo, Feng Chen, and Chi Zhang. 2025. [Dianjin-r1:
758 Evaluating and enhancing financial reasoning in large
759 language models](#). *Preprint*, arXiv:2504.15716.

760 Jie Zhu, Junhui Li, Yalong Wen, and Lifan Guo. 2024.
761 [Benchmarking large language models on CFLUE - a
762 Chinese financial language understanding evaluation
763 dataset](#). In *Findings of the Association for Computa-
764 tional Linguistics: ACL 2024*, pages 5673–5693,
765 Bangkok, Thailand. Association for Computational
766 Linguistics.

767 A Data Sources

768 Table 5 provides an overview of the sources, data
769 types, time spans, and licensing used to construct
770 QuantEval. All benchmark instances are based on
771 publicly accessible data or appropriately licensed
772 materials, enabling reproducibility without requir-
773 ing commercial data vendors.

774 B Backtesting Environment Specification

775 We provide the complete and deterministic back-
776 testing configuration used in QuantEval. All strate-
777 gies are evaluated under the same asset universe,

778 data frequency, time range, execution assumptions,
779 cost model, and risk constraints.

780 B.1 Backtesting Configuration

781 **Overview.** To ensure strict reproducibility of
782 strategy coding evaluation, all model-generated
783 strategies are executed under the same determin-
784 istic backtesting configuration. The configuration
785 specifies the data range and frequency, a fixed asset
786 universe, execution and transaction cost assump-
787 tions, risk constraints, and metric definitions.

Deterministic Backtesting Configuration (Summary + YAML-style JSON)

Component Specification (Fixed Across All Evaluations)	
Data	NYSE calendar; daily frequency (1D); adjusted close; 2010-01-01 to 2025-01-01
Universe	Fixed U.S. ETFs + large-cap stocks (15 assets total)
Execution	Next-bar open market orders; shorting allowed
Costs	Commission 2 bps/side; slippage 1 bps/side; applied on turnover
Risk	Max leverage 2.0; max single-asset weight 0.20; max turnover/day 1.0; no-lookahead enforced
Metrics	Annualized return; max drawdown; annualized Sharpe; return/drawdown ratio

```
{  
  "data": {  
    "frequency": "1D",  
    "calendar": "NYSE",  
    "start_date": "2010-01-01",  
    "end_date": "2025-01-01",  
    "price_type": "adjusted_close",  
    "missing_value_policy": "forward_fill_limit_3"  
  },  
  "universe": {  
    "type": "fixed",  
    "assets": [  
      "SPY", "QQQ", "IWM", "TLT", "GLD",  
      "AAPL", "MSFT", "AMZN", "GOOG", "META",  
      "NVDA", "TSLA", "JPM", "XOM", "UNH"  
    ]  
  },  
  "execution": {  
    "order_type": "market",  
    "fill_rule": "next_bar_open",  
    "allow_short": true  
  },  
  "costs": {  
    "commission_bps_per_side": 2.0,  
    "slippage_bps_per_side": 1.0,  
    "apply_on": "turnover"  
  },  
  "risk_constraints": {  
    "max_gross_leverage": 2.0,  
    "max_single_asset_weight": 0.20,  
    "max_turnover_per_day": 1.0,  
    "no_lookahead_enforced": true  
  },  
  "metrics": {  
    "annualized_return_pct": "geometric_mean_daily_return  
      * 252",  
    "max_drawdown_pct": "  
      max_peak_to_trough_drawdown_on_equity_curve",  
    "annualized_sharpe": "mean(daily_return)/std(  
      daily_return)*sqrt(252), rf=0",  
    "return_drawdown_ratio": "annualized_return_pct /  
      max_drawdown_pct"  
  }  
}
```

788 **Reproducibility.** We release the full configura-
789 tion and evaluation harness to enable exact repli-
790 cation of strategy coding results under identical
791

Task	Primary Sources	Data Types	Time Span / Coverage	License / Availability
Knowledge QA	Finance textbooks, academic papers, regulatory filings, curated glossaries	Text	N/A	Publicly available / licensed academic materials
Quantitative Reasoning	Yahoo Finance, Stooq, FRED, Quandl (public datasets)	Prices, rates, volatility indices, macro indicators	Varies by asset class (e.g., 2010–2024)	Publicly accessible; redistribution follows source licenses
Strategy Coding	Open-source strategy repositories, academic implementations, expert-written CTA templates	Python code + aligned historical market data	Varies by strategy (e.g., 2016–2024)	Open-source code under permissive licenses; expert templates released with benchmark

Table 5: **Data sources used for QuantEval.** All benchmark instances rely on publicly accessible data or appropriately licensed materials. Commercial data vendors (e.g., Bloomberg) are not required for reproducing the benchmark.

assumptions.

B.2 Failure Criteria

A strategy is marked as **non-executable** if any of the following holds: (i) the code fails to compile or throws exceptions during execution; (ii) required class/function interfaces are missing or mismatched; (iii) forbidden framework APIs are invoked; (iv) future information is accessed (look-ahead); or (v) any risk constraint is violated.

C De-duplication, Leakage Prevention, and Difficulty Control

This appendix details our de-duplication and leakage prevention procedures, as well as the difficulty control and coverage design used to construct QuantEval.

C.1 De-duplication and Leakage Prevention

To improve benchmark reliability and reduce unintended redundancy, we apply de-duplication and leakage prevention at both the question and code levels.

De-duplication.

- **Knowledge QA and Quantitative Reasoning.** We remove exact duplicates using hashing, and filter near-duplicates by computing sentence embedding cosine similarity using a RoBERTa-based encoder. Instances with similarity greater than 0.92 are considered near-duplicates and are removed.
- **Strategy Coding.** We normalize code by removing comments and formatting (e.g., whitespace, line breaks) and then apply both token-level hashing and AST-level similarity checks. Coding instances whose normalized AST similarity exceeds 0.85 to any existing sample are discarded to prevent near-duplicate implementations.

Leakage prevention. We mitigate potential training-data leakage through a two-stage procedure:

- **Exclusion from training corpora.** All QuantEval instances (questions, reference answers, and ground-truth strategy code) are explicitly excluded from any supervised fine-tuning (SFT) or reinforcement learning (RL) training data used in this work.
- **Overlap scanning against public benchmarks.** We scan for overlaps against major open-source financial QA datasets, including FinQA, TAT-QA, FinEval, FinanceIQ, and Finova. Overlap detection is performed by matching (i) question stems, (ii) key entities, and (iii) numeric patterns. Any suspicious candidates are manually reviewed by domain experts and removed if overlap is confirmed.

Code-level leakage checks for strategy coding.

For strategy coding tasks, we additionally search for near-matching implementations in public repositories by comparing normalized AST signatures and key signal expressions. Instances with high similarity are either rewritten by experts to ensure originality or discarded.

Note. While these procedures substantially reduce leakage risk, we acknowledge that zero overlap with proprietary model training corpora cannot be strictly guaranteed.

C.2 Difficulty Control and Coverage

QuantEval is designed to cover a broad range of quantitative finance concepts with controlled difficulty across the three task categories.

Difficulty control. We control difficulty by varying (i) the number of reasoning steps required, (ii) the complexity of financial formulas, and (iii) the degree of market-context integration:

- **Knowledge QA:** ranges from definition-level recall to concept comparison and applied understanding.
- **Quantitative Reasoning:** spans single-formula computations to multi-step problems combining multiple concepts (e.g., volatility estimation followed by option pricing).
- **Strategy Coding:** includes implementation from scratch, partial completion, and debugging under strict interface and risk constraints.

Coverage design. To ensure broad and balanced coverage, we categorize all instances by quantitative finance subdomain (e.g., derivatives, portfolio theory, risk metrics, macro finance) and enforce balanced distributions during data construction and validation. This design prevents the benchmark from over-representing a small subset of topics and improves its diagnostic value for model evaluation.

D Dataset Details

This section provides detailed information about our benchmark.

D.1 Human Performance

To establish a performance baseline, we asked four relevant experts to independently evaluate the entire dataset across our three core tasks: Knowledge-based QA, Quantitative Mathematical Reasoning, and Quantitative Strategy Coding. Table 6 summarizes their performance on each task.

Expert	Knowledge QA (%)	Reasoning (%)
Expert 1	91.2	88.5
Expert 2	93.0	89.7
Expert 3	90.5	87.9
Expert 4	92.3	90.1

Table 6: Human expert performance on QuantEval. Accuracy is reported for Knowledge QA and Reasoning.

The results demonstrate that human experts achieve high accuracy on Knowledge QA and Reasoning, providing an approximate upper bound for these tasks.

D.2 Information about Experts

Our expert team consists of seven members with diverse but complementary expertise. Four experts

Name	Professional Background
Expert 1	Computer Science & Quantitative Research
Expert 2	Computer Science & Quantitative Research
Expert 3	Quantitative Trading Practitioner
Expert 4	Quantitative Trading Practitioner
Expert 5	Computer Science & Quantitative Research
Expert 6	Quantitative Trading Practitioner
Expert 7	Quantitative Trading Practitioner

Table 7: Professional backgrounds of the experts involved in dataset construction and evaluation.

have strong backgrounds in computer science combined with quantitative research, while three experts have extensive hands-on experience in quantitative trading practice. This blend ensures both theoretical rigor and practical relevance in dataset construction and evaluation.

D.3 Composition and Quality Comparison Between Human and Automated Agents

Table 8 shows the distribution of samples generated by human experts and automated agents across the three main task categories.

Task Category	Human Expert	Automated Agent	Total Samples
Knowledge QA	520	140	660
Quantitative Reasoning	655	200	855
Strategy Coding	10	50	60

Table 8: Distribution of samples constructed by human experts and automated agents across the three core tasks in QuantEval.

To assess the quality of data generated by both human experts and automated agents, we evaluated model performance on each subset. Table 9 reports the accuracy of GPT-5 on these subsets.

Task	Expert Data	Agent Data	Difference
Knowledge QA (Acc.%)	62.8	61.3	-1.5
Quantitative Reasoning (Acc.%)	55.6	53.9	-1.7
Strategy Coding (Sharpe Ratio)	0.17	0.20	+0.003

Table 9: Model performance comparison on human-expert-generated and automated-agent-generated subsets. Differences are small, indicating comparable data quality.

The close performance metrics suggest that after expert validation, automated agent-generated data achieves quality comparable to human expert contributions, supporting the scalability and reliability of our data construction pipeline.

E Experimental details

E.1 Evaluation Reliability

To validate automatic grading for QA and reasoning, we compare automated judgments with expert annotations on a random subset of 2,00 instances. The agreement exceeds 95%, indicating high reliability of the evaluation protocol.

E.2 Evaluation for Model’s Output

To evaluate model outputs, we use an automated evaluation prompt that extracts the final answer and compares it to the ground truth. Table 10 shows the evaluation prompt templates for multiple-choice (MC) and open-ended questions.

Type	Evaluation Prompt
Multiple Choice	<p>You are an evaluation assistant.</p> <p>Please determine whether the model’s answer below is correct.</p> <p>Question: [question] Options: [options]</p> <p>Correct answer: [correct answer]</p> <p>Model output: [model output]</p> <p>Extract the final answer from the model output and check if it matches the correct answer.</p> <p>Reply with "1" if correct, "0" otherwise. Only reply with the number.</p>
Open-ended	<p>You are an evaluation assistant.</p> <p>Please determine whether the model’s answer below is correct.</p> <p>Question: [question]</p> <p>Correct answer: [correct answer]</p> <p>Model output: [model output]</p> <p>If the meanings are roughly consistent, consider it correct.</p> <p>Reply with "1" if correct, "0" otherwise. Only reply with the number.</p>

Table 10: Prompt templates used for automated evaluation of model outputs on multiple-choice and open-ended questions.

We verified the reliability of automated evaluation by comparing it with human expert judgments on a random sample of 200 questions. The inter-rater agreement was above 95%, confirming the robustness of our evaluation protocol.

E.3 Coding Prompt Ablation: CoT vs Direct Prompting

We compare Chain-of-Thought prompting and direct code-only prompting for Strategy Coding on three representative models. All other settings are identical.

Model	CoT Exec. Rate (%)	Direct Exec. Rate (%)
Qwen3-30B-A3B	8.3	0.0
Claude-4.5-sonnet	63.3	4.2
GPT-5	51.7	2.5

Table 11: Ablation on Strategy Coding executability under CoT prompting versus direct code-only prompting. Direct prompting yields near-zero executability across models, while CoT substantially improves compliance with the CTA framework.

F More Cases

Knowledge-based QA

Q: If a counterparty defaults before maturity, which of the following situations will cause a credit loss?

- A: Short EUR in 1-yr EUR/USD forward and EUR appreciates **X**; C: Sold 1-yr OTC EUR call and EUR appreciates **X**
 B: Short EUR in 1-yr EUR/USD forward and EUR depreciates **✓**; D: Sold 1-yr OTC EUR call and EUR depreciates **X**

◆ **Qwen3-14B(with COT): B ✓**

1. Identify the question and key data. - The question asks which situation causes a credit loss if a counterparty defaults before maturity, comparing short euro forward and sold euro call option scenarios under euro appreciation or depreciation.

2. Recall the relevant rule: **A credit loss occurs only when the contract has a positive value to you (the counterparty owes you money).** Forwards can create exposure; sold options generally do not, since the premium is received upfront.

3. Evaluate each option. - In case (a), Euro appreciates → short forward loses value → no credit loss. - **In case (b), Euro depreciates → short forward gains value → credit loss if counterparty defaults.** In cases (c) and (d), and (d) Sold options → no positive exposure → no credit loss. The Gulf of California is located between the states of California and Baja California. [[B]]

◆ **Qwen3-14B: B ✓**

Answer: B

◆ **Qwen3-8B (with COT): B&C ✗**

1. Identify the task. - Determine which scenarios create a credit loss if the counterparty defaults before maturity. Credit loss arises when the position is adverse to you at default.

2. Set the rule of thumb. **If the contract has moved against you (you owe value), default can crystallize a loss.**

3. Case (a): Short EUR forward, EUR appreciates. - EUR strength hurts a short-EUR forward → position is adverse. → Credit loss. Case (b): Short EUR forward, EUR depreciates. EUR weakness benefits a short-EUR forward → position is favorable. Case (c): Sold EUR call, EUR appreciates. Call becomes in the money against the seller → position is adverse. → Credit loss. Case (d): Sold EUR call, EUR depreciates. Call is out of the money for the buyer → favorable to the seller. → No credit loss.

◆ **Qwen3-8B: A&C ✗**

1. Identify the question and key data. - Decide which scenarios create a credit loss upon counterparty default before maturity. Compare a short EUR/USD forward and a sold OTC euro call under euro up/down moves.

2. Recall the working rule (as he uses it): **A credit loss happens when the position is unfavorable (negative value) to you at the time of default.** If the contract moves against you, you suffer a loss when the counterparty defaults.

3. Evaluate and conclude (to match a & c). - (a) Euro ↑ makes a short EUR forward adverse → counted as credit loss. - (c) Sold call with Euro ↑ is adverse to the writer → counted as credit loss. - (b), (d) are not treated as credit losses under his framing.

◆ **Qwen3-4B (with COT): C ✗**

1. Define credit loss. - It occurs if the counterparty defaults when the position is in your favor (they owe you), so you cannot collect the positive MTM.

2. Apply the rule to each case. - Check whether the move makes your position favorable (you're owed) or adverse (you owe).

3. Case (a): Short EUR forward, EUR ↑. - Move is adverse to you (you owe); default doesn't create a credit loss for you. Case (b): Short EUR forward, EUR ↓. - Move is favorable (you're owed); exposure exists but not taken as the selected outcome here. Case (c): Sold EUR call, EUR ↑. - **Option is in-the-money for the buyer; you're on the losing side and rely on their payment → default causes a credit loss.** Case (d): Sold EUR call, EUR ↓. - Option is out-of-the-money; no obligation from the counterparty, so no credit loss. [[C]]

◆ **Qwen3-4B: C ✗**

C

◆ **DeepSeek-R1-Distill-Qwen-14B (with COT): None ✗**

1. Opening Analysis Direction

The LLM identifies the core question as "which scenario causes credit loss when a counterparty defaults before maturity" and plans to analyze 4 options using the operation logic of forward FX contracts and OTC options.

2. Option-by-Option Breakdown

Option a: It explains "short euros in a forward" means delivering euros; with euro appreciation, it speculates potential loss from default but doesn't confirm, noting "needs further analysis". Option b: It states euro depreciation benefits the short position but questions if "unrealized profit loss" is credit loss, also "needing further analysis". Option c: It clarifies selling a call option requires selling euros if exercised; with euro appreciation, it guesses default may cause loss but lacks a conclusion. Option d: It directly rules out this option, as the buyer won't exercise when the euro depreciates, so no loss from default.

3. Further Deduction & Contradiction

The LLM defines credit loss as "financial loss from counterparty default" but repeatedly compares options a and c without eliminating either, trapped in "possible loss" speculation without distinguishing forward/option credit exposure.

4. Conclusion Status

The LLM's answer cuts off mid-sentence; it never identifies a single correct option or closes the logic loop, failing to reach a final conclusion.

◆ **DeepSeek-R1-Distill-Qwen-14B: None ✗**

◆ **DeepSeek-R1-Distill-Qwen-7B (with COT): None ✗**

1. Opening Goal & Credit Loss Definition

The LLM aims to identify which scenario causes credit loss if the counterparty defaults, defining credit loss as "failure to recover principal and interest due to counterparty non-performance".

2. Option-by-Option Analysis by the LLM

Option a: It notes a short euro forward with euro appreciation is unfavorable for the short position but is unsure if it causes credit loss. Option b: It states a short euro forward with euro depreciation benefits the short position but also doubts credit loss. Option c: It thinks a sold euro call (euro appreciation) may lead to the buyer exercising, but is uncertain about credit loss from default. Option d: It claims a sold euro call (euro depreciation) means the buyer won't exercise; counterparty default causes loss of "premium/principal", so it's credit loss.

◆ **DeepSeek-R1-Distill-Qwen-7B: C ✗**

The LLM incorrectly concludes c is correct, ignoring that short options (c/d) have no credit exposure. The right answer (b) involves a short forward with euro depreciation.

Quantitative Mathematical Reasoning

Q: A zero-coupon bond with a maturity of 10 years has an annual effective yield of 10%. What is the closest value for its modified duration?

- ✓ **Option:** A: 9 years ✓ C: 99 years ✗
 B: 10 years ✗ D: 100 years ✗

◆ **Qwen3-14B(with COT): A ✓**

1. Identify the question and key data. - The problem gives a zero-coupon bond with a 10-year maturity and an annual effective yield of 10%. The task is to find its modified duration.

2. Recall the relevant formula. - Modified duration measures a bond's price sensitivity to interest rate changes and is calculated as: $D^* = \frac{D_M}{1+y}$ where D_M is the Macaulay duration and y is the annual yield.

3. Substitute the values. - A. 9 years — closest to 9.09, correct. - B. 10 years — Macaulay duration, not modified. - C. 99 years — unrealistic. - D. 100 years — unrealistic. **The modified duration of the bond is approximately 9 years.**

◆ **Qwen3-14B: A ✓**

[[A]]

◆ **Qwen3-8B(with COT): A ✓**

Step 1: Analyze the question. - We are asked to find the modified duration of a zero-coupon bond with a maturity of 10 years and an annual effective yield of 10%. The options are: -a. 9 years -b. 10 years -c. 99 years -d. 100 years.

Step 2: Recall the definitions of Macaulay duration and modified duration.

• The Macaulay duration measures the weighted average time to receive a bond's cash flows.

• The modified duration adjusts Macaulay duration for the interest rate and shows price sensitivity. The relationship is: $D^* = \frac{D_M}{1+y}$ where D_M is Macaulay duration, and y is the annual yield.

Step 3: Determine the Macaulay duration. - For a zero-coupon bond, all payments occur only at maturity, so its Macaulay duration equals its maturity.

Step 4: Compute the modified duration. - Substitute the given yield $y = 10\% = 0.10$ into the formula: $D^* = \frac{10}{1+0.10} = \frac{10}{1.10} = 9.09$ years. Evaluate the answer choices: **a. 9 years → closest to 9.09 → correct.**

◆ **Qwen3-8B: A ✓**

[[A]]

◆ **Qwen3-4B(with COT): A ✓**

Step 1: Analyze the question. - We are asked to find the modified duration of a zero-coupon bond with a maturity of 10 years and an annual effective yield of 10%. The options are: -a. 9 years -b. 10 years -c. 99 years -d. 100 years.

Step 2: Recall the relationship between Macaulay duration and modified duration. - Modified duration measures the percentage change in bond price for a 1% change in yield. It is calculated as: $D^* = \frac{D_M}{1+y}$ where D_M is the Macaulay duration and y is the annual yield.

Step 3: Determine the Macaulay duration. - For a zero-coupon bond, all payments occur only at maturity, so its Macaulay duration equals its maturity.

3. Substitute the values. - A. 9 years — closest to 9.09, correct. - B. 10 years — Macaulay duration, not modified. - C. 99 years — unrealistic. - D. 100 years — unrealistic. **The modified duration of the bond is approximately 9 years.**

◆ **Qwen3-4B: A ✓**

[[A]]

◆ **DeepSeek-R1-Distill-Qwen-14B (with COT): A ✓**

Step 1: Analyze the question. - We are asked to find the modified duration of a zero-coupon bond with a maturity of 10 years and an annual effective yield of 10%. The options are: -a. 9 years -b. 10 years -c. 99 years -d. 100 years.

Step 2: Recall key definitions. - Zero-coupon bond: Pays no coupons; only one cash flow occurs at maturity. - Macaulay duration (D_M): The weighted average time to receive all cash flows. For a zero-coupon bond, all the weight is at maturity, so $D_M = 10$ years. - Modified duration (D^*): Measures price sensitivity to yield changes and equals $D^* = \frac{D_M}{1+y}$ where y is the annual yield.

Step 3: Determine the Macaulay duration. - For a zero-coupon bond, all payments occur only at maturity, so its Macaulay duration equals its maturity.

Step 4: Compute the modified duration. Given $y = 10\% = 0.10$: $D^* = \frac{10}{1+0.10} = \frac{10}{1.10} = 9.09$ years.

Step 5. Substitute the values. - A. 9 years — closest to 9.09, correct. - B. 10 years — Macaulay duration, not modified. - C. 99 years — unrealistic. - D. 100 years — unrealistic. **The modified duration of the bond is approximately 9 years.**

◆ **DeepSeek-R1-Distill-Qwen-14B: A ✓**

[[A]]

◆ **DeepSeek-R1-Distill-Llama-8B (with COT): A ✓**

Step 1: We are asked which detail from the description most suggests an intimate relationship. - The options are: -a. 9 years -b. 10 years -c. 99 years -d. 100 years.

Step 2: Recall key idea. - Zero-coupon bond: Pays no coupons; only one cash flow occurs at maturity. - Macaulay duration (D_M): The weighted average time to receive all cash flows. For a zero-coupon bond, all the weight is at maturity, so $D_M = 10$ years. - Modified duration (D^*): Measures price sensitivity to yield changes and equals $D^* = \frac{D_M}{1+y}$, where y is the annual yield.

Step 3: Substitute the values. - A. 9 years — closest to 9.09, correct. - B. 10 years — Macaulay duration, not modified. - C. 99 years — unrealistic. - D. 100 years — unrealistic. **The modified duration of the bond is approximately 9 years.**

◆ **DeepSeek-R1-Distill-Llama-8B: A ✓**

[[A]]

Quantitative Strategy Coding

1. Platform & Template

- Use VeighNa (vn.py)
- Develop LlmStrategy based CtaTemplate from vnpy_ctastrategy.template.

2. Core Logic

- ◆ **on_init:** Create BarGenerator (1min → 15min K-line) and ArrayManager (cache length 100).
 - ◆ **on_bar:** Push 1min K-lines to BarGenerator for 15min synthesis.
 - ◆ **on_15min_bar:**
 - Calculate Bollinger Bands, Aroon, ATR.
 - No position: Place stop orders at Bollinger Bands (breakout entry), filtered by Aroon (bullish → buy, bearish → sell).
 - Long position: Stop order at Bollinger mid-band (exit); Short position: Stop order at mid-band (exit).
 - Close position with limit orders if held over threshold (e.g., 100 periods).
 - Trailing stop: Track highs/lows, place close order at 2x ATR retracement.
3. Code Standards: Detailed Chinese comments, and PEP8 compliance.

Task Requirement Summary

1. Platform & Template

- Develop a trading strategy using VeighNa (vn.py) platform, based on the CtaTemplate from the vnpy_ctastrategy.template submodule.

2. Strategy Logic (Multi-Scale Fractal)

- ◆ Calculate the fractal dimension of price sequences across different time windows to monitor market complexity changes.
- ◆ Go long when the fractal dimension drops below a lower threshold (indicating strengthened trend); go short when it rises above an upper threshold (indicating high market complexity).
- ◆ Support two-way futures trading (long/short) and require high-frequency data for fractal measurement.

3. Code Requirements

- Include detailed comments explaining the program logic.
- Use Type Hinting for type annotations.
- Comply with PEP8 code standards.

- 4. Backtesting: Evaluate the strategy's performance via backtesting over the past 3 years.

◆ GPT-5: Best Performance: 101

```
from vnpy_ctastrategy import CtaTemplate, BarGenerator, ArrayManager
from vnpy.trader.constant import Interval, Direction, Offset

class LlmStrategy(CtaTemplate):
    """
    LlmStrategy - trend-following strategy based on VeighNa platform
    Combines Bollinger Bands, Aroon Indicator and ATR for trading decisions
    """
    author = "Quant Researcher"
    # Core parameters
    boll_window: int = 20
    aroon_window: int = 25
    atr_window: int = 14
    max_holding: int = 100
    fixed_size: int = 1
    parameters = ["boll_window", "aroon_window", "atr_window", "max_holding", "fixed_size"]
    variables = ["holding_period", "highest", "lowest"]

    def __init__(self, cta_engine, strategy_name, vt_symbol, setting):
        super().__init__(cta_engine, strategy_name, vt_symbol, setting)
        # Create bar generator (1min -> 15min bars)
        self.bg = BarGenerator(
            on_bar=self.on_bar, window=15,
            on_window_bar=self.on_15min_bar, interval=Interval.MINUTE
        )
        # Create array manager (caches 100 bars)
        self.am = ArrayManager(size=100)

    def on_init(self) -> None:
        self.load_bar(days=10) # Load historical data for initialization

    def on_bar(self, bar):
        self.bg.update_bar(bar) # Push 1min bars for 15min synthesis

    def on_15min_bar(self, bar):
        self.cancel_all() # Cancel all pending orders
        self.am.update_bar(bar)
        if not self.am.inited: return
        # Calculate technical indicators
        self.am.boll(self.boll_window, 2.0)
        self.am.aroon(self.aroon_window)
        self.am.atr(self.atr_window)
        # Update position status...
        ...

1. Testing Period & Basic Metrics: The LlmStrategy was backtested for 1 day (Jan 2, 2020), with 1 profitable day and 0 losing days.
2. Capital Performance: Starting with $1,000,000, it ended at $3,034,455.63, generating $2,034,455.63 total net profit with 0% max drawdown.
3. Trading Costs & Activity: Total commissions: $11,190.37; total slippage: $18,000. Total turnover: $373,012,470 with 1 trade executed.
4. Return Metrics: Total return: 203.45%; annualized return: 48,826.94%; daily return: 111.00%. Sharpe ratio and return-drawdown ratio were 0 (single-day test).
5. Strategy Alignment: Performance reflects the strategy's designed logic—Bollinger Band breakouts, Aroon filtering, and ATR trailing stops worked effectively.
```

◆ Gemini 2.5-Pro: Best Performance: 15

```
from vnpy_ctastrategy import CtaTemplate, BarGenerator, ArrayManager
from vnpy.trader.constant import Interval, Direction, Offset
from typing import Optional

class TrendMomentumStrategy(CtaTemplate):
    """
    Trend Momentum Strategy - Combines MACD, RSI and volatility filtering for
    trend following
    """
    author = "Quant Researcher"
    # Core parameters
    macd_fast: int = 12
    macd_slow: int = 26
    macd_signal: int = 9
    rsi_window: int = 14
    vol_window: int = 20
    fixed_size: int = 1
    stop_loss_ratio: float = 0.02
    parameters = ["macd_fast", "macd_slow", "macd_signal", "rsi_window", "vol_window", "fixed_size", "stop_loss_ratio"]
    variables = ["trend_state", "last_macd", "volatility"]

    def __init__(self, cta_engine, strategy_name, vt_symbol, setting):
        super().__init__(cta_engine, strategy_name, vt_symbol, setting)

        # 1-minute to 5-minute bar synthesis
        self.bg = BarGenerator(
            on_bar=self.on_bar,
            window=5,
            on_window_bar=self.on_5min_bar,
            interval=Interval.MINUTE
        )
        self.am = ArrayManager(size=100) # Cache for up to 100 bars
        # Strategy state variables
        self.trend_state: str = "neutral" # Trend status: neutral/bull/bear
        self.last_macd: float = 0.0
        self.volatility: float = 0.0
        self.entry_price: float = 0.0

    def on_init(self) -> None:
        self.write_log("Strategy initialized")
        self.load_bar(days=10) # Load 10 days of historical data

    def on_start(self) -> None:
        self.write_log("Strategy started")
        self.put_event()
        ...

1. Testing Period & Basic Metrics: The TrendMomentumStrategy was backtested for 1 day (Jan 3, 2020), with 1 profitable day and 0 losing days.
2. Capital Performance: Starting with $1,000,000, it ended at $1,520,300.75, generating $520,300.75 total net profit with 0% max drawdown.
3. Trading Costs & Activity: Total commissions: $8,500.25; total slippage: $12,000. Total turnover: $283,341,650 with 2 trades executed.
4. Return Metrics: Total return: 52.03%; annualized return: 18,990.95%; daily return: 52.03%. Sharpe ratio and return-drawdown ratio were 0 (single-day test).
```