FinChain: A Symbolic Benchmark for Verifiable Chain-of-Thought Financial Reasoning

Anonymous ACL submission

Abstract

Multi-step symbolic reasoning is essential for robust financial analysis; yet, current benchmarks largely overlook this capability. Existing datasets such as FinQA and ConvFinQA emphasize final numerical answers while neglecting the intermediate reasoning required for transparency and verification. To address this gap, we introduce FinChain, the first benchmark specifically designed for verifiable Chain-of-Thought (CoT) evaluation in finance. FinChain spans 58 topics across 12 financial domains, each represented by parameterized symbolic templates with executable Python traces that enable fully machine-verifiable reasoning and scalable, contamination-free data generation. To assess reasoning capacity, we propose ChainEval, a dynamic alignment metric that jointly evaluates both the final-answer correctness and the step-level reasoning consistency. Evaluating 23 leading LLMs reveals that even frontier proprietary systems exhibit clear limitations in symbolic financial reasoning, while domain-adapted and math-enhanced fine-tuned models substantially narrow this gap. Overall, FinChain exposes persistent weaknesses in multi-step financial reasoning and provides a foundation for developing trustworthy, interpretable, and verifiable financial AI.

FinChain Template (Compound Interest)

#Question:
`investor_name` invested `principal` in `project_name`. The investment grows at an annual interest rate of `rate`% compounded annually over `time` years. Calculate the `compound interest (CI)`.

#Variables:
- `investor_name` = sample(investors)
- `project_name` = sample(projects)
- `principal` = range(1000, 5000)
- `rate` = uniform(2, 10)
- `time` = range(1, 5)

#Chain-of-Thought Solution:
Step 1: Compute the compound amount: `amount` = `principal` $\times \left(1 + \frac{\texttt{rate}}{100}\right)^{\texttt{time}}$
Step 2: Compute the compound interest: `CI` = `amount` − `P`

Figure 1: **Symbolic template for generating compound interest problems in FinChain.**

1 Introduction

Large language models (LLMs) have shown strong performance across diverse tasks (Zhao et al., 2023a; Xie et al., 2023b), including applications in finance, healthcare, and law (Chen et al., 2024b). In finance, effective analysis often requires processing large volumes of text from reports, news, and social media, which frequently reflect or influence financial phenomena such as market sentiment (Nie et al., 2024). Prior work in financial NLP has mainly adapted general-purpose tasks to financial documents, including information extraction (Shah et al., 2023), sentiment analysis (Pei et al., 2022), and text classification (Sy et al., 2023), where models produce relatively simple targets such as entity spans or sentiment labels. In contrast, financial reasoning requires generating chain-of-thought (CoT) traces that justify each step in solving a problem, as illustrated in Figure 1.

Existing benchmarks such as FinQA (Chen et al., 2021) and ConvFinQA (Chen et al., 2022) frame reasoning as numerical question answering, emphasizing final answers without requiring explicit or verifiable intermediate reasoning. Although some examples include partial traces, these are neither comprehensive nor rigorously structured. Lacking explicit step-level supervision limits diagnostic power: such benchmarks cannot reveal where reasoning fails or distinguish genuine multi-step inference from pattern matching. In mathematical

reasoning, the symbolic-template paradigm introduced by GSM-Symbolic (Mirzadeh et al., 2024) re-templates problems from GSM8K (Cobbe et al., 2021) to provide explicit intermediate steps. We adopt a similar approach but construct our dataset entirely from scratch for finance. As shown in Figure 1, each symbolic template encodes a parameterized financial problem (e.g., compound interest) with named variables and numeric inputs, paired with executable Python code computing both intermediate steps and final results. This design enables scalable, contamination-free generation of consistent examples for training and evaluation. The financial domain spans diverse topics with distinct reasoning styles. To capture this diversity, we organize our dataset using a fine-grained taxonomy (Figure 2) covering 12 domains (e.g., Corporate Finance, Sustainable Finance, Crypto) and 58 topics. For each topic, we develop five parameterized templates: two easy, two intermediate, and one advanced, varying in symbolic complexity and required expertise. This forms the most detailed taxonomy of financial reasoning tasks to date. Each templated instance includes (1) a scenario card describing the topic, difficulty, and sampled inputs (e.g., "*Discounted cash flow valuation, advanced*") and (2) an executable chain of reasoning steps in Python grounded in domain-specific formulas. Because every operation is explicit and executable, the benchmark is fully machine-verifiable: hallucinated, skipped, or incorrect steps can be automatically detected. This contrasts with datasets like FinQA and ConvFinQA, which supervise only final answers.

To enable rigorous evaluation, we introduce ChainEval, a dynamic-alignment metric assessing both final-answer correctness and intermediate-step consistency. Our evaluation across proprietary and open-source LLMs reveals a consistent trend: while frontier models such as `GPT-5`, `Claude 4.5`, and `Gemini 2.5 Pro` achieve the highest accuracy, they still struggle with complex symbolic reasoning. Smaller or domain-specific models perform worse, indicating that financial fine-tuning alone is insufficient; robust multi-step reasoning also requires adequate capacity for symbolic and multi-hop inference. Our main contributions are as follows:

- We present the first from-scratch symbolic benchmark for financial reasoning, built using a fine-grained financial taxonomy spanning 12 domains and 58 topics.
- We propose ChainEval, an alignment metric evaluating both final-answer correctness and intermediate reasoning consistency.
- We benchmark 23 open-source and proprietary LLMs, showing that even state-of-the-art systems face substantial difficulty with symbolic financial reasoning, especially on advanced multi-step templates.

2 Related Work

2.1 Financial NLP

Progress in financial NLP has been driven by both modeling and benchmarking. Early work focused on extraction and classification with models such as FinBERT (Araci, 2019), while later efforts expanded to personal finance (Hean et al., 2025), credit scoring (Feng et al., 2023), and risk-awareness benchmarking (Yuan et al., 2024). Datasets like FiNER-ORD, REFinD, FinARG, and ECTSum support tasks in NER, relation extraction, argument mining, and summarization (Shah et al., 2023; Kaur et al., 2023; Mukherjee et al., 2022; Xie et al., 2024). Large financial language models have further advanced the field. BloombergGPT (Wu et al., 2023) achieved broad domain performance, FinGPT (Liu et al., 2023) emphasized open-source adaptability, and FinMA (Xie et al., 2023a) delivered competitive results with a compact architecture. Corresponding benchmarks such as FLANG (Shah et al., 2022), FinBen (Xie et al., 2024), and FinMTEB (Tang and Yang, 2025) broadened evaluation coverage across diverse tasks, while BizBench (Koncel-Kedziorski et al., 2023) and PIXIU (Xie et al., 2023a) revealed limitations in quantitative and multimodal reasoning.

Despite this progress, limitations remain in multi-step reasoning, long-context understanding, and cross-market generalization (Chen et al., 2024b). **These challenges motivate the need for benchmarks that assess capability of models to perform faithful, auditable reasoning grounded in financial knowledge.**

2.2 Financial Reasoning

Real-world financial problems require precise numerical reasoning. FinQA (Chen et al., 2021) and ConvFinQA (Chen et al., 2022) supervise arithmetic program generation but provide only weak step-level supervision, yielding traces that are neither explicit nor verifiable. FinTextQA (Chen et al., 2024a) introduces long-form financial questions

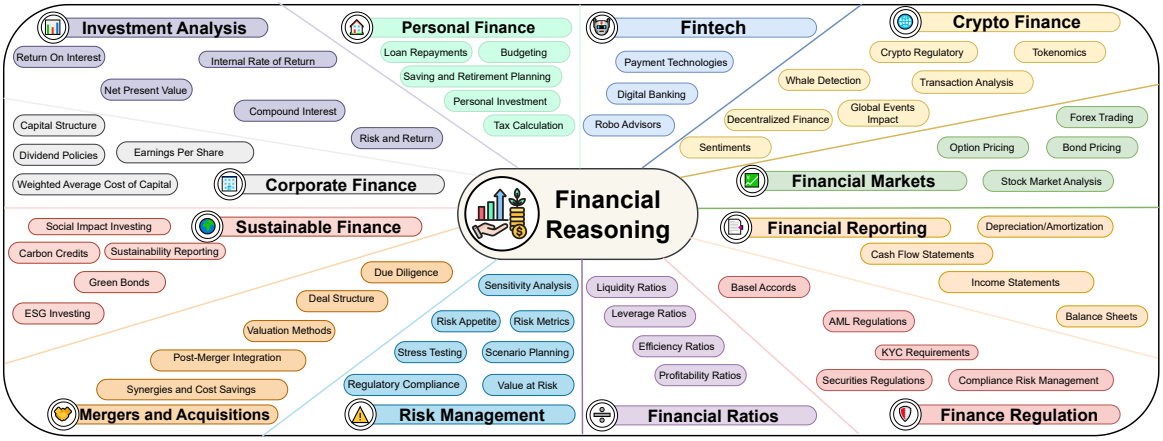

Figure 2: **FinChain taxonomy of financial reasoning topics.** Our benchmark spans 58 topics organized into 12 major domains, ranging from traditional areas like *Corporate Finance* and *Financial Reporting* to emerging fields such as *Crypto Finance* and *Sustainable Finance*. This hierarchical structure enables fine-grained evaluation of symbolic reasoning across diverse financial domains.

from textbooks and regulatory sources, focusing on explanatory retrieval rather than traceable computation. Bridging text and numerical reasoning, TAT-QA (Zhu et al., 2021) and MultiHiertt (Zhao et al., 2022) combine textual and tabular evidence, while DocMath-Eval (Zhao et al., 2023c) and FinanceMath (Zhao et al., 2023b) move toward interpretable, symbolic evaluation. However, these datasets still lack explicit, step-level, domain-aware supervision within authentic financial contexts. Model-centric advances such as FinR1 (Liu et al., 2025) and Fino1 (Qian et al., 2025) improve performance through program-based reinforcement learning and instruction tuning but remain limited in multi-step reasoning.

Overall, existing benchmarks offer either coarse program supervision (FinQA, ConvFinQA), fine-grained but domain-agnostic traces (GSM-Symbolic, DocMath-Eval), or retrieval-based QA without traceability (FinTextQA, TAT-QA, MultiHiertt, FinanceMath). None provide systematic, step-level supervision grounded in financial equations. **FinChain fills this gap by introducing a symbolic, executable benchmark that supervises each intermediate step, supports automatic alignment and human assessment, and spans 54 topics across 12 financial domains.**

3 FinChain

3.1 Data Creation Process

We begin by identifying and defining financial domains based on established literature (Bodie et al., 2025) and expert input within the team. This process yields 12 distinct domains. To generate topics within each domain, we extract relevant passages from the literature and prompt ChatGPT[1] with the domain name to propose candidate financial topics. Financial experts then curate these outputs, resulting in 58 topics in total (mean 4.8 per domain). The resulting taxonomy is illustrated in Figure 2. Following Mirzadeh et al. (2024), we implement executable Python methods that generate questions and solutions with chain-of-thought reasoning using ChatGPT. We upload the full paper and prompt ChatGPT to summarize the data creation process, then instruct it to produce ten template based questions per topic, explicitly varying complexity (easy, intermediate, and advanced) according to the number of reasoning steps. The template creation prompt is provided in Appendix A. Finally, we manually select five representative templates per topic (two easy, two intermediate, and one advanced) to ensure balanced difficulty in FinChain.

3.2 Data Validation

Directly prompting ChatGPT to mimic our style and generate questions sometimes fails. After iteratively inspecting generations and refining prompts, we introduced additional constraints to address recurring issues:

Cross-national inconsistencies. Prompts occasionally produced hyper-localized financial contexts (e.g., currencies, exchange rates, indices, terminol-

[1]We use GPT-4o, the latest version available at the time of writing.

ogy from various countries). We standardized all questions to U.S.-based financial settings.

Precision mismatch. Questions and solutions often displayed rounded values while computations used full precision. We aligned computational outputs with the displayed precision.

Incomplete input specification. Some questions omitted variables required for calculation. We edited these cases to include all necessary inputs.

Unit consistency. Currency notation (e.g., $) was inconsistently applied across questions and solutions. We standardized units throughout.

Non-informative steps. Certain solutions decomposed a single calculation into trivial substeps or only presented results. We revised such cases to reflect substantive reasoning steps.

Multiple targets. Some questions requested multiple values, which complicates the evaluation process. We constrained questions to require a single target value.

3.3 Expert Check

To further enhance data quality, we recruited ten financial experts who volunteered to participate in this project. The team included seven graduate students in economics, finance, and related quantitative disciplines, and three industry professionals with practical experience in quantitative research, financial engineering, and risk management. All reviewers were carefully selected through an internal vetting process to ensure domain expertise and professional credibility. Their diverse profiles across academia, asset management, auditing, and financial technology helped maintain annotation consistency and domain coverage. Further demographic and professional details of the annotators are provided in Appendix B. These experts were tasked with carefully reviewing the ChatGPT-generated templates after the data validation phase. Before starting the main annotation, we conducted a pilot study to ensure consistent evaluation standards among annotators, followed by the distribution of the remaining tasks.

Annotation Platform To facilitate efficient and user-friendly annotation, we developed a Streamlit-based annotation platform. Implementation details are presented in Appendix C.

Pilot Study We randomly selected 20 templates for the pilot study and asked all annotators to review them. The goal was to align their annotation standards. After calibration, all annotators agreed that the selected templates were correct.

Main Annotation After the pilot study, the remaining 270 templates were randomly distributed among the annotators, with each template reviewed by a single expert. Out of the total 290 templates, 29 were identified as incorrect. For those templates, we use the versions which are fixed by our financial experts. Detailed statistics of the identified issues are summarized in Appendix D.

4 ChainEval

We propose ChainEval, an evaluation framework assessing model outputs along two axes: final-answer correctness and reasoning-step alignment. Inspired by prior work on reasoning consistency (Lyu et al., 2023; Golovneva et al., 2023), we extend it by explicitly verifying intermediate results through step–answer matching. Unlike earlier approaches that focus mainly on textual consistency, our framework also checks the final numeric answer, providing a holistic view of model performance in both reasoning faithfulness and end-task fidelity. To our knowledge, this combination has not been explored in financial reasoning.

4.1 Preliminaries

We define the gold solution S^* and the predicted solution $\hat{S}$ as sequences of m and n reasoning steps, respectively:

$$S^* = (s_1^*, \ldots, s_m^*), \quad \hat{S} = (\hat{s}_1, \ldots, \hat{s}_n), \tag{1}$$

where s_i^* and $\hat{s}_j$ denote individual reasoning steps in the gold and predicted solutions. Each step s_i produces an intermediate result,

$$\mathrm{StepRes}(s_i) = a_i, \tag{2}$$

representing the numerical or symbolic value computed at that step.

To evaluate reasoning faithfulness, we compare these sequences both semantically and numerically. In addition, we employ *Dynamic Time Warping* (DTW) to capture the global structural alignment between step sequences. DTW provides an order-preserving but flexible alignment that accommodates insertions, deletions, or small reordering of steps while maintaining overall sequence coherence.

4.2 Reasoning Step Alignment

We assess the consistency between gold and predicted reasoning traces through two complementary components: semantic similarity and answer-level agreement, combined within a DTW-based alignment framework.

Semantic Similarity. Each step is encoded using a sentence encoder $\mathrm{Enc}(\cdot)$, and pairwise semantic similarity between gold and predicted steps is computed as

$$\mathrm{SS}(s_i^*, \hat{s}_j) = \mathrm{CosSim}(\mathrm{Enc}(s_i^*), \mathrm{Enc}(\hat{s}_j)), \quad (3)$$

where $\mathrm{CosSim}(\cdot,\cdot)$ denotes cosine similarity and $\mathrm{SS} \in [0,1]$.

Answer Match. For the intermediate results produced by each step, we evaluate numeric or symbolic consistency:

$$\mathrm{StepRes}(s_i^*) = a_i^*, \qquad \mathrm{StepRes}(\hat{s}_j) = \hat{a}_j.$$

We then define the answer-matching function:

$$\mathrm{AM}(s_i^*, \hat{s}_j) = \begin{cases} \mathbb{I}\left(\frac{|\hat{a}_j - a_i^*|}{|a_i^*|} \le \epsilon\right), \\ \text{if both are numeric,} \\ \mathbb{I}(\hat{a}_j = a_i^*), \text{ otherwise.} \end{cases} \quad (4)$$

where $\mathbb{I}(\cdot)$ is the indicator function, and $\epsilon = 0.05$ allows for up to 5% relative numeric deviation due to rounding or propagation errors.

Gated Step-Level Similarity. To ensure that a pair of steps is considered consistent only when both their semantics and results agree, we define a gated score:

$$\mathrm{Score}_{\mathrm{gate}}(i,j) = \mathrm{SS}(s_i^*, \hat{s}_j) \times \mathrm{AM}(s_i^*, \hat{s}_j). \quad (5)$$

This score forms the basis of the DTW alignment matrix.

Dynamic Sequence Alignment. To capture global reasoning consistency, we align the two step sequences using *Dynamic Time Warping* (DTW). DTW searches for an optimal monotonic alignment path between $(S^*, \hat{S})$ that minimizes cumulative cost while respecting step order. This formulation naturally handles local insertions, deletions, and step compressions, providing a structured measure of how closely the reasoning flows align.

Normalized DTW Alignment Score (Gate Mode). We transform the minimal DTW cost into a normalized similarity measure:

$$\mathrm{DTWNormGate}(S^*, \hat{S}) = 1 - \frac{\mathrm{Cost}_{\mathrm{DTW}}}{L_{\mathrm{path}}}, \quad (6)$$

where $\mathrm{Cost}_{\mathrm{DTW}}$ denotes the total alignment cost and L_{path} represents the length of the optimal alignment path. The resulting score lies in the range $[0,1]$, with higher values indicating stronger reasoning alignment between the gold and predicted solutions.

4.3 Final Answer Correctness

Beyond step-level reasoning alignment, we also assess the correctness of the final predicted outcome. Let s_m^* and $\hat{s}_n$ denote the last steps of the gold and predicted solutions, respectively. We define the **Final Answer Correctness (FAC)** metric as:

$$\mathrm{FAC}(S^*, \hat{S}) = \begin{cases} \mathbb{I}\left(\frac{|\hat{a}_n - a_m^*|}{|a_m^*|} \le \epsilon\right), \\ \text{if both are numeric,} \\ \mathbb{I}(\hat{a}_n = a_m^*), \text{otherwise,} \end{cases} \quad (7)$$

using the same tolerance $\epsilon = 0.05$ as before. FAC measures whether the model's final computation aligns with the correct end result, complementing the DTW-based metric that evaluates reasoning faithfulness throughout the entire solution sequence.

Finally, ChainEval equals DTWNormGate. It captures both local semantic and numeric agreement as well as global sequence-level coherence in the reasoning process. We verify that this metric best reflects the true reasoning quality by comparing it with human evaluations, as detailed in § 5.2.

5 Experiments and Results

5.1 Evaluated Models

We evaluate a total of **23 LLMs**, grouped into four categories according to their capability and relevance to financial reasoning. (1) **Frontier proprietary models** represent the current upper bound of performance, including `GPT-5, 4.1, 5-mini, 4.1-mini` (OpenAI, 2025a,b), `Claude Sonnet 4.5, 4, 3.7` (Anthropic, 2025b,c,a), `Gemini 2.5 Pro, Flash` (Comanici et al., 2025), and `DeepSeek V3.2, V3.1, R1` (Liu et al., 2024; Guo et al., 2025). (2) **Finance-specific models** include `Fino1` (Qian et al., 2025), `FinR1` (Liu et al.,

2025), and `WiroAI Finance-LLaMA`, `Qwen` (Abdullah Bezir, 2025), which are fine-tuned on financial corpora and reasoning templates to enhance domain accuracy. (3) **Math-enhanced models** include `WizardMath` (Luo et al., 2023), `MetaMath` (Yu et al., 2023), `Mathstral` (Mistral, 2024), and `Qwen2.5-Math` (Yang et al., 2024), which are trained on mathematical and symbolic datasets to improve quantitative reasoning. (4) **General-purpose open models** include `LLaMA 3.1` (Grattafiori et al., 2024) and `Qwen 2.5, 3` (Qwen, 2024, 2025), serving as strong domain-agnostic backbones for many specialized variants. This four-way taxonomy enables a structured comparison across reasoning paradigms, highlighting how model scale, domain specialization, and mathematical supervision jointly influence symbolic financial reasoning. Detailed configurations and model sources are provided in Appendix E.

5.2 ChainEval Validation

Before conducting large-scale experiments, we validated the proposed ChainEval through a controlled expert evaluation. We randomly sampled 20 instantiated questions from the dataset and generated answers using five models of different capacities and training paradigms, namely `GPT-5`, `GPT-4.1 mini`, `MetaMath`, `Fino1`, and `LLaMA 3.1`, to ensure a clear range of reasoning quality. This process produced 100 model-generated responses, which were then randomized and anonymized for human assessment. Financial experts independently evaluated each response with respect to the corresponding question and gold-standard reasoning trace. Each output was rated along two dimensions, *Reasoning Process Quality* and *Final Answer Accuracy*, on a five-point scale as described in Appendix F. We observed a strong association between these dimensions, with Spearman's ρ exceeding 0.94, confirming that coherent reasoning typically leads to accurate final answers. Consequently, subsequent analyses focus on *Reasoning Process Quality* as the primary evaluation dimension. We further computed several variants of ChainEval and compared their correlations with expert process scores against reference-based metrics such as ROUGE-2 and ROUGE-L (Lin, 2004). Among all variants, the DTWNormGate formulation showed the highest correlation with human judgments, capturing both semantic and numeric consistency across reasoning steps. Therefore, DTWNormGate is adopted as the principal evaluation metric in this study, with detailed ablation results reported in Appendix G.

5.3 Experimental Setup

We instantiate the benchmark by sampling 10 instances per symbolic template with distinct random seeds, yielding 58 topics $\times$ 5 templates $\times$ 10 instances = 2,900 test cases. All models are evaluated under a unified decoding configuration: temperature = 0.7, top-p = 0.95, and a maximum token limit of 4,096 unless these parameters are unavailable for a given model. Main results in Section 5.4.1 use a zero-shot setup with a standardized reasoning prompt:

```
Please answer the given question and provide a
step-by-step solution.
Use the format: Step 1: ..., Step 2: ..., ...
The question is: {q}
```

We use ChainEval as the primary evaluation metric, as it jointly measures final-answer correctness and alignment of intermediate reasoning steps. Model outputs are post-processed with regular expressions to extract the ordered list of reasoning steps, accommodating common variations such as "Step x:", "Step x", or "stepx". For comparison, we also report frequently used reference-based metrics (Xie et al., 2023c) to assess surface-level quality, including ROUGE-2 and ROUGE-L (Lin, 2004) and BERTScore (Zhang et al., 2020).

5.4 Results and Analysis

5.4.1 Overall Model Performance

As shown in Table 1, proprietary frontier models continue to set the overall performance ceiling across all evaluation metrics. Systems such as GPT-5, Claude 4.5, and Gemini 2.5 Pro consistently achieve the highest ChainEval, ROUGE, and BERTScore values, reflecting their superior reasoning, factual accuracy, and linguistic fluency. These results confirm the persistent performance gap between closed and open ecosystems, likely driven by larger model scales, broader pre-training corpora, and more extensive alignment pipelines. Despite strong progress in open-source development, none of the publicly available 7B–8B models yet match the frontier tier in aggregate capability. Nevertheless, fine-tuning substantially enhances the performance of open backbones, demonstrating that targeted adaptation remains an effective strategy for improving reasoning quality within resource-constrained settings. For example, Fin-R1 achieves a ChainEval of 57.34, comparable to much larger proprietary systems and notably

Model	Size	ChainEval ↑	ROUGE R_2 ↑	ROUGE R_L ↑	BERTScore ↑
Frontier Proprietary LLMs					
GPT-5	N/A	$57.07^{7.28}$	$\mathbf{28.84}^{12.30}$	$\mathbf{42.77}^{12.91}$	$\mathbf{88.77}^{2.39}$
GPT-4.1	N/A	$56.92^{6.12}$	$19.38^{8.91}$	$30.12^{10.57}$	$86.04^{2.09}$
GPT-5-mini	N/A	$\mathbf{57.38}^{7.12}$	$26.48^{11.99}$	$39.74^{12.83}$	$88.18^{2.35}$
GPT-4.1-mini	N/A	$57.24^{5.82}$	$18.67^{8.86}$	$29.05^{10.51}$	$86.05^{2.09}$
Claude Sonnet 4.5	N/A	$58.22^{6.43}$	$19.69^{7.77}$	$29.37^{8.81}$	$86.07^{2.00}$
Claude Sonnet 4	N/A	$58.18^{6.60}$	$19.87^{7.82}$	$29.50^{8.59}$	$86.38^{1.83}$
Claude Sonnet 3.7	N/A	$57.89^{6.22}$	$19.49^{7.77}$	$29.36^{8.56}$	$86.38^{1.82}$
Gemini 2.5 Pro	N/A	$\mathbf{58.65}^{6.87}$	$17.61^{7.07}$	$27.36^{8.28}$	$85.94^{1.88}$
Gemini 2.5 Flash	N/A	$58.01^{6.99}$	$18.98^{8.05}$	$29.22^{9.32}$	$86.34^{2.02}$
DeepSeek v3.2	N/A	$56.71^{6.25}$	$21.73^{10.32}$	$32.85^{11.31}$	$86.66^{2.15}$
DeepSeek v3.1	N/A	$56.76^{6.40}$	$\mathbf{21.72}^{10.29}$	$\mathbf{32.87}^{11.24}$	$\mathbf{86.68}^{2.14}$
DeepSeek R1	N/A	$53.75^{7.90}$	$8.67^{7.21}$	$12.93^{9.72}$	$84.39^{1.79}$
Finance Specific LLMs					
Fin-o1	8B	$39.34^{12.10}$	$3.47^{1.55}$	$6.35^{2.32}$	$83.55^{1.50}$
Fin-R1	7B	$\mathbf{57.34}^{5.62}$	$5.70^{2.44}$	$9.22^{3.33}$	$\mathbf{84.30}^{1.34}$
Finance-LlaMA	8B	$42.81^{9.33}$	$9.39^{4.69}$	$16.19^{5.84}$	$83.48^{2.09}$
Finance-Qwen	7B	$34.22^{10.69}$	$\mathbf{9.50}^{4.26}$	$\mathbf{16.46}^{5.44}$	$83.35^{1.70}$
Math Enhanced LLMs					
WizardMath	7B	$21.75^{15.45}$	$11.66^{6.57}$	$20.72^{7.83}$	$84.78^{2.36}$
MetaMath	7B	$6.09^{9.24}$	$11.45^{7.36}$	$21.08^{9.24}$	$84.86^{2.99}$
Mathstral	7B	$\mathbf{56.40}^{5.96}$	$\mathbf{16.79}^{7.82}$	$\mathbf{26.97}^{9.34}$	$\mathbf{86.13}^{2.18}$
Qwen2.5-Math	7B	$50.32^{12.64}$	$11.74^{5.87}$	$20.56^{7.61}$	$83.45^{1.85}$
General Purpose Open LLMs					
Llama 3.1 Instruct	8B	$55.88^{4.95}$	$4.61^{2.28}$	$8.09^{3.02}$	$83.35^{1.36}$
Qwen 2.5 Instruct	7B	$57.00^{5.35}$	$9.20^{4.51}$	$15.26^{5.85}$	$84.22^{1.78}$
Qwen 3	8B	$45.99^{10.84}$	$4.05^{1.69}$	$6.61^{2.14}$	$83.58^{1.24}$

Table 1: **Zero-shot performance across financial, mathematical, and general reasoning benchmarks.** Scores are reported as percentages, with standard deviation in superscript. Model size (*N/A*) denotes proprietary or undisclosed configurations. Within each model group, the best-performing system for each metric is highlighted in bold.

higher than typical open general models. Similarly, Mathstral attains 56.40 ChainEval, outperforming other math-enhanced and general-purpose models of the same scale. These results indicate that fine-tuning can meaningfully elevate smaller architectures toward frontier-level reasoning, bridging part of the performance divide through domain-specific supervision rather than additional parameters. However, the gains from fine-tuning depend strongly on the breadth and quality of adaptation data. Models fine-tuned on broad domains such as mathematics exhibit improved general reasoning and text coherence, while those tuned narrowly on financial content sometimes trade generality for domain specialization, as seen with Finance-LLaMA and Finance-Qwen. Overall, Table 1 illustrates a clear hierarchy: frontier proprietary models remain dominant, but fine-tuned open models can approach their performance by leveraging high-quality, domain-adaptive instruction data. This underscores fine-tuning as a pragmatic path toward competitive reasoning capability without access to massive proprietary resources.

5.4.2 Performance Over Domains

Figure 3 illustrates domain-level ChainEval performance across twelve financial categories, comparing the best proprietary and fine-tuned open models. The frontier system GPT-5 maintains consistently strong and balanced performance across nearly all domains, reflecting broad reasoning coverage and stable generalization. Gemini 2.5 Pro exhibits a similar trend but shows slight advantages in corporate finance, fintech, and financial markets, likely benefiting from broader exposure to business-oriented content during pretraining. Together, these results confirm that large-scale proprietary LLMs continue to offer the most uniform and reliable performance across diverse financial contexts.

In contrast, fine-tuned open models display domain-selective strengths that align with their adaptation objectives. Fin-R1 achieves competitive or superior accuracy in financial reporting, sustainable finance, and risk management, outperforming frontier models in several specialized areas. Mathstral performs strongly in quantitative domains such as financial ratios and investment analysis but

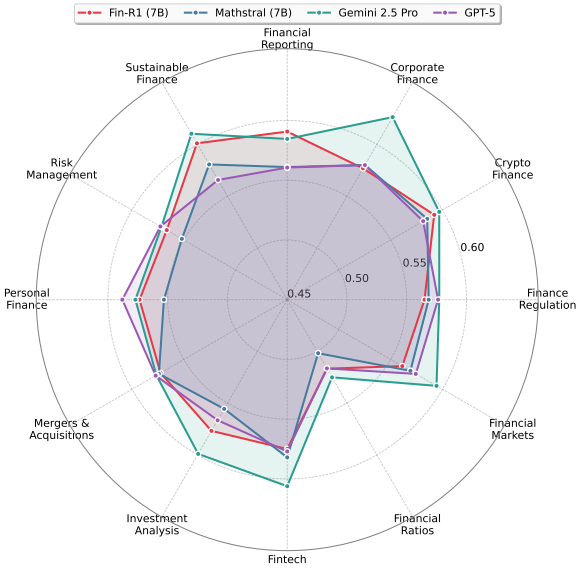

Figure 3: **Domain-level performance across financial domains.** Radar plot comparing the best-performing model from each category in Table 1, evaluated across twelve financial domains using ChainEval scores.

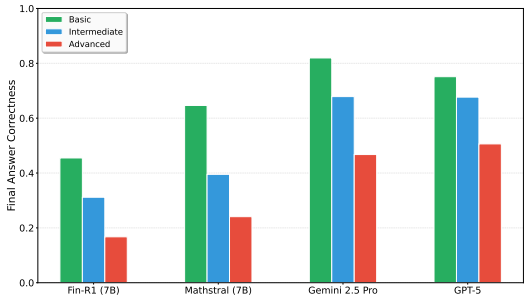

Figure 4: **ChainEval score across difficulty levels.**

trails in text-heavy categories like policy or corporate finance. These asymmetric patterns indicate that fine-tuning effectively amplifies reasoning in domain-relevant tasks but introduces variability across unrelated ones. Overall, fine-tuned open models demonstrate meaningful progress toward domain robustness, yet frontier proprietary systems remain more balanced and reliable across the entire financial spectrum.

5.4.3 Performance Over Difficulty Levels

To assess model robustness under increasing complexity of reasoning, we group the results into three predefined tiers (*Basic*, *Intermediate*, and *Advanced* (§ 3). Each difficulty level reflects an increase in reasoning and computational depth. Figure 4 presents model performance across three difficulty tiers measured by ChainEval scores. Across all tiers, frontier proprietary models demonstrate the highest overall correctness, with GPT-5 and Gemini 2.5 Pro showing strong resilience as task complexity increases. Their relatively small performance drop between Basic and Advanced levels highlights more stable reasoning and abstraction capabilities, likely reflecting the benefits of extensive instruction alignment and large, diverse training corpora. These trends reaffirm that frontier models maintain superior robustness and generalization across varying levels of reasoning difficulty.

In contrast, open and fine-tuned models exhibit sharper performance declines as difficulty increases but still show meaningful gains relative to their base architectures. Mathstral achieves competitive scores on Basic and Intermediate tasks, demonstrating that mathematical fine-tuning enhances structured reasoning, while Fin-R1 performs comparatively well on simpler finance-related queries but degrades more on Advanced ones. This suggests that fine-tuning improves domain reasoning and factual accuracy but offers limited transfer to high-complexity, multi-step tasks. Overall, the results indicate that while fine-tuned open models narrow the gap at lower difficulty levels, frontier proprietary systems retain a clear advantage in handling complex, compositional reasoning.

6 Conclusions and Future Work

We introduce FinChain, a symbolic benchmark for verifiable chain-of-thought financial reasoning. Built from a fine-grained taxonomy spanning 58 topics across 12 domains and three difficulty levels, it enables targeted evaluation across reasoning complexity and domain-specific challenges. To assess performance, we propose ChainEval, a dynamic alignment metric evaluating both final-answer correctness and intermediate-step consistency. Experiments show that while frontier LLMs achieve the highest accuracy, even the strongest systems struggle with complex symbolic reasoning. Fine-tuned open-source models such as Fin-R1 and Mathstral reduce but do not close this gap, underscoring the challenge of aligning fluency with verifiable reasoning. Future work will expand FinChain to multilingual and region-specific settings and explore how explicit reasoning traces can improve the trustworthiness of model-generated financial analyses, particularly in long-form question answering over real documents. This direction aims to bridge symbolic reasoning and factual verification (Xie et al., 2025), advancing more interpretable and reliable financial AI systems.

Limitations

We acknowledge several limitations in this work that we plan to address in future research.

First, our dataset is entirely synthetic, generated from symbolic templates. While this design enables fine-grained control, contamination-free generation, and automatic verifiability, it may lack the linguistic diversity and contextual richness of real-world financial texts. In future work, we plan to incorporate real-world financial documents, such as earnings reports, investor communications, or financial news, as seed sources for semi-structured or templated generation, allowing us to better simulate naturally occurring language while preserving symbolic grounding.

Second, the benchmark focuses narrowly on symbolic numerical reasoning and does not capture qualitative, contextual, or strategic aspects of financial decision-making, such as risk assessment, or market sentiment. To address this, we aim to complement our benchmark with additional tasks that evaluate models on these higher-level reasoning dimensions, potentially through multi-modal inputs (e.g., combining text with charts or scenarios) or interactive decision-making simulations.

Third, FinChain is restricted to English and U.S. centric financial conventions, such as currency formats, and investment norms, which limits its generalizability to multilingual and regional financial contexts. In future extensions, we plan to expand coverage to other languages and financial systems, through collaboration with domain experts and native speakers to localize templates and ensure culturally grounded reasoning tasks.

Fourth, we use regular expressions for parsing responses of LLMs, but it is not ideal. That is why ChainEval may be computed on unrelated parts of the text. In future we plan to migrate to LLM-based parsing of responses and inclusion of length penalty, which will help to penalize models for overly verbose steps.

Ethical Statement and Broad Impact

This work uses only synthetic data generated through templated code and language model outputs. No private, sensitive, or copyrighted content was used. Our benchmark is designed for transparency and reproducibility in financial AI. However, caution should be taken when deploying LLMs in real-world financial decision-making, especially where symbolic correctness and regulatory compliance are critical. We believe FinChain will support research toward more interpretable, verifiable, and safe reasoning systems in high-stakes domains.

Data License The FinChain dataset and accompanying code will be released under the MIT License.

References

Cengiz Asmazoğlu Abdullah Bezir, Furkan Burhan Türkay. 2025. Wiroai/wiroai-finance-qwen-7b. *Hugging Face Hub*.

Anthropic. 2025a. Claude 3.7 sonnet system card. https://assets.anthropic.com/m/785e231869ea8b3b/original/claude-3-7-sonnet-system-card.pdf.

Anthropic. 2025b. Claude sonnet 4.5 system card. https://assets.anthropic.com/m/12f214efcc2f457a/original/Claude-Sonnet-4-5-System-Card.pdf.

Anthropic. 2025c. System card: Claude opus 4 & claude sonnet 4. https://www-cdn.anthropic.com/6be99a52cb68eb70eb9572b4cafad13df32ed995.pdf.

Dogu Araci. 2019. Finbert: Financial sentiment analysis with pre-trained language models. *ArXiv preprint*, abs/1908.10063.

Zvi Bodie, Robert C. Merton, and Richard T. Thakor. 2025. *Principles of Finance*. Cambridge University Press.

Jian Chen, Peilin Zhou, Yining Hua, Loh Xin, Kehui Chen, Ziyuan Li, Bing Zhu, and Junwei Liang. 2024a. FinTextQA: A dataset for long-form financial question answering. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6025–6047, Bangkok, Thailand. Association for Computational Linguistics.

Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2021. FinQA: A dataset of numerical reasoning over financial data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3697–3711, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah, and William Yang Wang. 2022. ConvFinQA: Exploring the chain of numerical reasoning in conversational finance question answering. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 6279–6292, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Zhiyu Zoey Chen, Jing Ma, Xinlu Zhang, Nan Hao, An Yan, Armineh Nourbakhsh, Xianjun Yang, Julian McAuley, Linda Petzold, and William Yang Wang. 2024b. A survey on large language models for critical societal domains: Finance, healthcare, and law. *ArXiv preprint*, abs/2405.01769.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, et al. 2021. Training verifiers to solve math word problems. *ArXiv preprint*, abs/2110.14168.

Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.

Duanyu Feng, Yongfu Dai, Jimin Huang, Yifang Zhang, Qianqian Xie, Weiguang Han, Zhengyu Chen, Alejandro Lopez-Lira, and Hao Wang. 2023. Empowering many, biasing a few: Generalist credit scoring through large language models. *ArXiv preprint*, abs/2310.00566.

Olga Golovneva, Moya Chen, Spencer Poff, Martin Corredor, Luke Zettlemoyer, Maryam Fazel-Zarandi, and Asli Celikyilmaz. 2023. ROSCOE: A suite of metrics for scoring step-by-step reasoning. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *ArXiv preprint*, abs/2407.21783.

Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

Oudom Hean, Utsha Saha, and Binita Saha. 2025. Can ai help with your personal finances? *Applied Economics*, pages 1–9.

Simerjot Kaur, Charese Smiley, Akshat Gupta, Joy Sain, Dongsheng Wang, Suchetha Siddagangappa, Toyin Aguda, and Sameena Shah. 2023. Refind: Relation extraction financial dataset. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR 2023, Taipei, Taiwan, July 23-27, 2023*, pages 3054–3063. ACM.

Rik Koncel-Kedziorski, Michael Krumdick, Viet Lai, Varshini Reddy, Charles Lovering, and Chris Tanner. 2023. Bizbench: A quantitative reasoning benchmark for business and finance. *ArXiv preprint*, abs/2311.06602.

Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.

Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.

Xiao-Yang Liu, Guoxuan Wang, Hongyang Yang, and Daochen Zha. 2023. Fingpt: Democratizing internet-scale data for financial large language models. *ArXiv preprint*, abs/2307.10485.

Zhaowei Liu, Xin Guo, Fangqi Lou, Lingfeng Zeng, Jinyi Niu, Zixuan Wang, Jiajie Xu, Weige Cai, Ziwei Yang, Xueqian Zhao, et al. 2025. Fin-r1: A large language model for financial reasoning through reinforcement learning. *ArXiv preprint*, abs/2503.16252.

Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. 2023. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *ArXiv preprint*, abs/2308.09583.

Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. 2023. Faithful chain-of-thought reasoning. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 305–329, Nusa Dua, Bali. Association for Computational Linguistics.

Iman Mirzadeh, Keivan Alizadeh, Hooman Shahrokhi, Oncel Tuzel, Samy Bengio, and Mehrdad Farajtabar. 2024. Gsm-symbolic: Understanding the limitations of mathematical reasoning in large language models. *ArXiv preprint*, abs/2410.05229.

AI Mistral. 2024. Mathstral. 7B parameter model for mathematical reasoning, released under Apache 2.0 license.

Rajdeep Mukherjee, Abhinav Bohra, Akash Banerjee, Soumya Sharma, Manjunath Hegde, Afreen Shaikh, Shivani Shrivastava, Koustuv Dasgupta, Niloy Ganguly, Saptarshi Ghosh, and Pawan Goyal. 2022. ECTSum: A new benchmark dataset for bullet point summarization of long earnings call transcripts. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 10893–10906, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Yuqi Nie, Yaxuan Kong, Xiaowen Dong, John M Mulvey, H Vincent Poor, Qingsong Wen, and Stefan Zohren. 2024. A survey of large language models for financial applications: Progress, prospects and challenges. *ArXiv preprint*, abs/2406.11903.

OpenAI. 2025a. Gpt-5 system card. https://cdn.openai.com/gpt-5-system-card.pdf. System card describing GPT-5 variants (thinking/main/mini/-nano).

OpenAI. 2025b. Introducing gpt-4.1 in the api. https://openai.com/index/gpt-4-1/. Product research post introducing GPT-4.1, 4.1 mini, and 4.1 nano.

Yulong Pei, Amarachi Mbakwe, Akshat Gupta, Salwa Alamir, Hanxuan Lin, Xiaomo Liu, and Sameena Shah. 2022. TweetFinSent: A dataset of stock sentiments on Twitter. In *Proceedings of the Fourth Workshop on Financial Technology and Natural Language Processing (FinNLP)*, pages 37–47, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.

Lingfei Qian, Weipeng Zhou, Yan Wang, Xueqing Peng, Jimin Huang, and Qianqian Xie. 2025. Fino1: On the transferability of reasoning enhanced llms to finance. *ArXiv preprint*, abs/2502.08127.

Team Qwen. 2024. Qwen2.5 technical report. *ArXiv preprint*, abs/2412.15115.

Team Qwen. 2025. Qwen3.

Agam Shah, Ruchit Vithani, Abhinav Gullapalli, and Sudheer Chava. 2023. Finer: Financial named entity recognition dataset and weak-supervision model. *ArXiv preprint*, abs/2302.11157.

Raj Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Raman, Charese Smiley, Jiaao Chen, and Diyi Yang. 2022. When FLUE meets FLANG: Benchmarks and large pretrained language model for financial domain. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2335, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Eugene Sy, Tzu-Cheng Peng, Shih-Hsuan Huang, Heng-Yu Lin, and Yung-Chun Chang. 2023. Fine-grained argument understanding with BERT ensemble techniques: A deep dive into financial sentiment analysis. In *Proceedings of the 35th Conference on Computational Linguistics and Speech Processing (ROCLING 2023)*, pages 242–249, Taipei City, Taiwan. The Association for Computational Linguistics and Chinese Language Processing (ACLCLP).

Yixuan Tang and Yi Yang. 2025. Finmteb: Finance massive text embedding benchmark. *ArXiv preprint*, abs/2502.10990.

Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. *ArXiv preprint*, abs/2303.17564.

Qianqian Xie, Weiguang Han, Zhengyu Chen, Ruoyu Xiang, Xiao Zhang, Yueru He, Mengxi Xiao, Dong Li, Yongfu Dai, Duanyu Feng, et al. 2024. The finben: An holistic financial benchmark for large language models. *ArXiv preprint*, abs/2402.12659.

Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. 2023a. Pixiu: A large language model, instruction data and evaluation benchmark for finance. *ArXiv preprint*, abs/2306.05443.

Zhuohan Xie, Trevor Cohn, and Jey Han Lau. 2023b. The next chapter: A study of large language models in storytelling. In *Proceedings of the 16th International Natural Language Generation Conference*, pages 323–351, Prague, Czechia. Association for Computational Linguistics.

Zhuohan Xie, Miao Li, Trevor Cohn, and Jey Lau. 2023c. DeltaScore: Fine-grained story evaluation with perturbations. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 5317–5331, Singapore. Association for Computational Linguistics.

Zhuohan Xie, Rui Xing, Yuxia Wang, Jiahui Geng, Hasan Iqbal, Dhruv Sahnan, Iryna Gurevych, and Preslav Nakov. 2025. FIRE: Fact-checking with iterative retrieval and verification. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2901–2914, Albuquerque, New Mexico. Association for Computational Linguistics.

An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, Keming Lu, Mingfeng Xue, Runji Lin, Tianyu Liu, Xingzhang Ren, and Zhenru Zhang. 2024. Qwen2.5-math technical report: Toward mathematical expert model via self-improvement. *ArXiv preprint*, abs/2409.12122.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023. Metamath: Bootstrap your own mathematical questions for large language models. *ArXiv preprint*, abs/2309.12284.

Tongxin Yuan, Zhiwei He, Lingzhong Dong, Yiming Wang, Ruijie Zhao, Tian Xia, Lizhen Xu, Binglin Zhou, Fangqi Li, Zhuosheng Zhang, Rui Wang, and Gongshen Liu. 2024. R-judge: Benchmarking safety risk awareness for LLM agents. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 1467–1490, Miami, Florida, USA. Association for Computational Linguistics.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. Bertscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, Yifan Du, Chen Yang, Yushuo Chen, Zhipeng Chen, Jinhao Jiang, Ruiyang Ren, Yifan Li, Xinyu Tang, Zikang Liu,

Peiyu Liu, Jian-Yun Nie, and Ji-Rong Wen. 2023a. A survey of large language models. *ArXiv preprint*, abs/2303.18223.

Yilun Zhao, Yunxiang Li, Chenying Li, and Rui Zhang. 2022. Multihiertt: Numerical reasoning over multi hierarchical tabular and textual data. *arXiv preprint arXiv:2206.01347*.

Yilun Zhao, Hongjun Liu, Yitao Long, Rui Zhang, Chen Zhao, and Arman Cohan. 2023b. Financemath: Knowledge-intensive math reasoning in finance domains. *arXiv preprint arXiv:2311.09797*.

Yilun Zhao, Yitao Long, Hongjun Liu, Ryo Kamoi, Linyong Nan, Lyuhao Chen, Yixin Liu, Xiangru Tang, Rui Zhang, and Arman Cohan. 2023c. Docmath-eval: Evaluating math reasoning capabilities of llms in understanding long and specialized documents. *arXiv preprint arXiv:2311.09805*.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. TAT-QA: A question answering benchmark on a hybrid of tabular and textual content in finance. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 3277–3287, Online. Association for Computational Linguistics.

A Financial-Symbolic: Template Creation Prompt

To construct symbolic financial reasoning benchmarks analogous to GSM-Symbolic, we design a structured prompt that guides the generation of executable financial templates. These templates support variable-based instantiation, symbolic step-wise supervision, and controlled perturbations for robustness evaluation. Below, we present the prompt used for template construction.

System Instruction: You are a financial NLP expert developing symbolic reasoning datasets. Your task is to construct symbolic templates for financial reasoning problems. Each template should support (i) controlled generation of diverse question instances, (ii) executable reasoning traces for automatic verification, and (iii) systematic variation in surface and logical complexity.

Please follow the steps below:

1. **Identify a financial reasoning task:** For example, compound interest, IRR, ROI, NPV, breakeven analysis, loan amortization, etc.

2. **Write a natural language question template:** Formulate the question using variable placeholders instead of fixed values. For instance, use `{principal}`, `{rate}`, `{years}`, etc.

3. **Define variables and constraints:** Specify the domain (e.g., numerical range or categorical values) for each variable. Add algebraic constraints to ensure the question is solvable and the answer valid.

4. **Write a symbolic solution trace:** Provide a step-by-step solution using the variables. Ensure the reasoning chain is executable in Python for automatic evaluation.

5. **Vary difficulty levels:** For each task, generate 10 templates with increasing complexity. Longer and more compositional reasoning chains should correspond to harder levels.

B Financial Expert Demography

To ensure the reliability and domain robustness of our benchmark, all annotations were conducted by a diverse team of financial experts and advanced students with strong quantitative and economic backgrounds. The annotators collectively represent three major categories: (1) industry professionals in quantitative research and financial engineering,

Model	Organization	Size	Backbone	Source
Frontier Proprietary LLMs				
GPT-5	OpenAI	Undisclosed	–	gpt-5-2025-08-07
GPT-4.1	OpenAI	Undisclosed	–	gpt-4.1-2025-04-14
GPT-5 mini	OpenAI	Undisclosed	–	gpt-5-mini-2025-08-07
GPT-4.1 mini	OpenAI	Undisclosed	–	gpt-4.1-mini-2025-04-14
Claude Sonnet 4.5	Anthropic	Undisclosed	–	claude-sonnet-4-5-20250929
Claude Sonnet 4	Anthropic	Undisclosed	–	claude-sonnet-4-20250514
Claude Sonnet 3.7	Anthropic	Undisclosed	–	claude-3-7-sonnet-20250219
Gemini 2.5 Pro	Google	Undisclosed	–	Last Update: June 2025
Gemini 2.5 Flash	Google	Undisclosed	–	Last Update: June 2025
DeepSeek V3.2	DeepSeek	Undisclosed	–	Last Update: Sep 29 2025
DeepSeek V3.1	DeepSeek	Undisclosed	–	Last Update: Sep 22 2025
DeepSeek R1	DeepSeek	Undisclosed	–	Last Update: Jan 20 2025
Finance Specific LLMs				
Fino1	TheFinAI	8B	meta-llama/Llama-3.1-8B	TheFinAI/Fin-o1-8B
Fin-R1	SUFE-AIFLM-Lab	7B	Qwen/Qwen2.5-7B-Instruct	SUFE-AIFLM-Lab/Fin-R1
Finance-LlaMA	Wiro AI	8B	deepseek-ai/DeepSeek-R1-Distill-Llama-8B	WiroAI/WiroAI-Finance-Llama-8B
Finance-Qwen	Wiro AI	7B	Qwen/Qwen2.5-7B	WiroAI/WiroAI-Finance-Qwen-7B
Math Enhanced LLMs				
WizardMath	WizardLM Team	7B	mistralai/Mistral-7B-v0.1	WizardLMTeam/WizardMath-7B-V1.1
MetaMath	MetaMath Project	7B	EleutherAI/llemma 7b	meta-math/MetaMath-7B-V1.0
Mathstral	Mistral AI	7B	mistralai/Mistral-7B-v0.1	mistralai/Mathstral-7B-v0.1
Qwen2.5-Math	Qwen Team	7B	Qwen/Qwen2.5-7B	Qwen/Qwen2.5-Math-7B-Instruct
General Purpose Open LLMs				
Llama 3.1	Meta	8B	-	meta-llama/Llama-3.1-8B
Qwen 2.5	Qwen Team	7B	-	Qwen/Qwen2.5-7B-Instruct
Qwen 3	Qwen Team	8B	-	Qwen/Qwen3-8B

Table 2: Details of the organization and model source (i.e. model version for proprietary models, and HuggingFace model name for open-source models) for the LLMs evaluated in FinChain.

(2) postgraduate students specializing in finance, economics, and auditing, and (3) experienced annotators trained in data labeling and financial analysis.

Several annotators have extensive industry experience across financial technology, quantitative research, and trading, with prior roles in investment banks, hedge funds, and fintech companies. Others are graduate students conducting research in finance, economics, and auditing, contributing academic rigor and theoretical grounding. Together, they bring complementary expertise that enhances both the practical and analytical aspects of our benchmark construction.

Summary. Our benchmark construction relies on a team of ten highly qualified annotators, including three industry professionals with prior experience in quantitative research or trading, and seven academic annotators who are graduate students in finance, economics, and auditing. This balanced composition, encompassing strong and diverse backgrounds in computer science, mathematics, statistics, and finance, ensures both professional authenticity and academic depth. Their combined expertise provides a robust foundation for high-quality, domain-consistent annotations, contributing to the overall reliability of FinChain. The following are the details for each of them.

Annotator A: Currently pursuing a Ph.D. at a leading university in Asia, this annotator previously worked as a quantitative researcher at a fintech company, with experience across multiple financial markets including domestic equities, U.S. equities, Hong Kong equities, and cryptocurrencies. Their research focused on financial data generation, risk modeling, and trading strategies. They have also served as a research lead in risk management at a cryptocurrency investment fund. This blend of academic research and cross-market industry practice enhances the robustness and domain relevance of the benchmark annotations.

Annotator B: A Master's student at a leading university with a strong undergraduate background in finance. They previously interned in the equity financing division of a major securities firm, contributing practical insights into capital markets and investment banking.

Annotator C: A Master's student at a top institution, holding a bachelor's degree in economics. Their training bridges theoretical economics and applied policy research, enriching the annotation process with domain-specific understanding.

Annotator D: Holds a bachelor's degree in economics and has received graduate admission offers

from top international institutions. Their interdisciplinary background strengthens the dataset's coverage of trade and international finance contexts.

Annotator E: Holds a bachelor's degree in economics, providing a solid foundation in macroeconomic theory and financial principles that supports reliable annotation and consistency across financial texts.

Annotator F: A Master's student at a well-known university specializing in auditing and intelligent systems, with a research focus on large language model evaluation and its applications in auditing. Their familiarity with both auditing and financial concepts supports the annotation of financial news and auditing benchmarks from a research-oriented perspective.

Annotator G: A Master's student at a university recognized for its auditing and financial programs, with strong grounding in auditing, financial analysis, and data quality control. Their prior participation in annotation projects ensures consistent standards for annotation accuracy.

Annotator H: A quantitative analyst with an MSc-equivalent degree in financial technology from a top UK university. They have prior experience at major global financial institutions, focusing on stochastic modeling, risk management, and process automation. They also contribute to research on large language models in finance and are advancing toward professional certification in investment analysis.

Annotator I: A quantitative researcher at a global investment firm with prior experience at quantitative research and technology companies. Their work spans cross-asset systematic strategies, portfolio optimization, and machine learning applications in trading. They also serve as a teaching assistant for a postgraduate course on systematic trading strategies.

Annotator J: A quantitative trading analyst focused on equity derivatives, holding a postgraduate degree in financial engineering and risk management from a top European university and a bachelor's degree from a globally recognized institution. Their professional experience includes roles at several financial institutions across asset management, banking, and fintech, covering alpha-signal development, portfolio optimization, and derivatives trading.

C Annotation Platform

We developed a custom annotation platform to evaluate the correctness of Python templates that generate financial questions and solutions. Each template corresponds to a financial scenario (e.g., investment analysis, compound interest, deposits, or ratio calculations). Annotators are instructed to review the code and determine whether both the financial framework and its implementation are correct, and whether the output representation (e.g., units, rounding) complies with the annotation policy.

The annotation task requires a binary verdict: *Correct* or *Incorrect*. Templates labeled as *Correct* need no modifications, though annotators may optionally provide comments. Templates labeled as *Incorrect* must be associated with one or more issue tags, accompanied by a minimal code correction and a brief explanation.

We defined two verdict categories. A template is considered *Correct* when its financial framework, calculations, and representation fully conform to the policy. It is marked as *Incorrect* if any substantive flaw is present in framework selection, mathematical logic, representation, robustness, or clarity. To facilitate consistent labeling, we introduced five issue tags:

- **Formula Choice Error**: An incorrect financial framework or formula is applied (e.g., simple vs. compound interest).
- **Math/Logic Error**: Arithmetic or algorithmic errors within the chosen formula (e.g., $r \times n$ instead of r/n).
- **Representation Error**: Inconsistent or incorrect handling of numbers, units, or rounding (e.g., annual vs. monthly mismatch).
- **Robustness Error**: Failures on boundary or extreme inputs (e.g., division by zero, negative values).
- **Clarity Issue**: Ambiguous variable names or comments that hinder auditability, even if the numerical results are correct.

To further support annotators, the platform provides curated reference cases across five templates within a single finance topic: compound interest. Each case includes (1) a question example, (2) a potential error type aligned with one of the defined

Issue Type	Count	Proportion (%)
Representation Error	12	41.4
Clarity Issue	9	31.0
Formula Choice Error	5	17.2
Math/Logic Error	3	10.3
Robustness Error	2	6.9
Total Tagged Templates	**29**	**100.0**

Table 3: Distribution of issue types among annotated templates.

issue tags, (3) a bad solution illustrating the error, and (4) a minimal code fix. Figure 5 shows two representative cases: a *Formula Choice Error*, where simple interest is incorrectly applied in a compound interest setting, and a *Math/Logic Error*, where the exponent is omitted. Such examples provide concrete guidance for annotators, ensuring consistency and reliability.

After reviewing these reference cases, annotators proceed to the main annotation interface, where they evaluate unseen templates (Figure 6). For each template, annotators must issue a binary verdict, select one or more issue tags if applicable, and provide a minimal code correction with a short justification. The interface presents the Python template and its generated question on the left, while the right panel allows annotators to record their verdict, choose tags, and edit the code directly. This design mirrors realistic auditing conditions and ensures that annotations capture both error identification and corrective reasoning.

D Annotated Template Issue Statistics

Out of 290 templates, 29 (10%) were tagged as containing errors during the annotation process. Table 3 summarizes the distribution of issue types. Most problems stem from representation and clarity errors, followed by formula selection, logical inconsistencies, and robustness issues.

E Model Detail Information

We provide the details of the evaluated models in Table 2.

F Review Rubrics

To ensure fair and interpretable human evaluation, each model response is assessed along two complementary dimensions: *Reasoning Process Quality* and *Final Answer Accuracy*. For each question, reviewers are provided with the question itself, the standard reference answer, and the generated responses from different models. They independently assign scores on a 1–5 scale for each dimension, following the detailed rubrics below.

F.1 Reasoning Process Quality

This dimension evaluates how clearly, logically, and correctly the model articulates its reasoning steps leading to the final answer. High-quality reasoning should demonstrate coherent logical flow, factual correctness, and consistency with valid domain principles.

- **1 (Unacceptable):** Illogical, incoherent, or irrelevant reasoning; missing steps or severe conceptual errors.
- **2 (Poor):** Some reasoning attempt but with major factual or procedural flaws; inconsistent or unclear chain of thought.
- **3 (Fair):** Partial understanding with mixed correct and incorrect reasoning; superficial or incomplete explanation.
- **4 (Good):** Mostly correct and coherent reasoning with minor inaccuracies or unclear phrasing; logical flow generally sound.
- **5 (Excellent):** Clear, well-structured, and logically consistent reasoning throughout; fully correct and well-justified steps.

F.2 Final Answer Accuracy

This dimension evaluates the correctness and completeness of the model's final output relative to the reference solution. Reviewers compare each model's final answer with the standard answer to determine whether the model's conclusion is correct and sufficiently supported.

- **1 (Unacceptable):** Completely incorrect or missing answer; no alignment with the reference solution.
- **2 (Poor):** Largely incorrect due to major conceptual or computational errors.
- **3 (Fair):** Partially correct; captures some relevant elements but omits or distorts key aspects of the correct solution.
- **4 (Good):** Largely correct and complete with only minor inaccuracies that do not affect the main result.

Simple CI Calculation (annual compounding): Template Questions & Tagged Error Cases

Question example

```
{Investor} invested $P in Project X. The investment grows at an annual interest rate of r% compounded annually over t years. Calculate the compound interest.
```

Cases

- **Uses simple interest inside a CI template**
  Tag: `Formula Choice Error`
  Why: Selected simple-interest framework (1+r*t) instead of compounding (1+r)^t.

Bad solution:

```
A = P * (1 + r/100 * t)
CI = A - P
```

Minimal fix:

```
A = P * (1 + r/100) ** t
CI = A - P
```

- **Exponent omitted**
  Tag: `Math/Logic Error`
  Why: Forgot to raise to the power t.

Bad solution:

```
A = P * (1 + r/100)
CI = A - P
```

Minimal fix:

```
A = P * (1 + r/100) ** t
```

Figure 5: Reference examples for compound interest templates, illustrating typical annotation cases with error tags, flawed solutions, and minimal fixes.

- **5 (Excellent):** Fully correct, precise, and complete; matches the reference solution exactly or with an equivalent formulation.

G Metric Evaluation and Ablations

To better understand the behavior of our proposed metric, we conducted a series of ablation experiments and comparative analyses. All quantitative results reported in this section are benchmarked against expert human evaluations of reasoning quality (see Appendix B for expert details).

G.1 Ablations of the DTW-Based Metric

Our main evaluation metric, the **Normalized DTW Alignment Score (Gate Mode)**, measures both local semantic-numeric agreement and global sequence-level alignment between predicted and gold reasoning traces. To assess its robustness and the effect of its design choices, we considered several variants:

- **DTW Gate Mode.** This is the primary formulation used in the paper. Semantic similarity and numeric agreement are combined multiplicatively, i.e., $\mathrm{Score}_{\mathrm{gate}}(i,j) = \mathrm{SS}(s_i^*, \hat{s}_j) \times \mathrm{AM}(s_i^*, \hat{s}_j)$. This "gating" ensures that steps only contribute when both semantic meaning and intermediate results align.

- **DTW Soft Mode.** A more permissive variant that blends semantic and numeric agreement through a weighted combination: $\mathrm{Score}_{\mathrm{soft}}(i,j) = \alpha\,\mathrm{SS}(s_i^*, \hat{s}_j) + \beta\,\mathrm{AM}(s_i^*, \hat{s}_j)$, with weights $\alpha = 0.85$ and $\beta = 0.15$. This "soft" formulation captures cases where partial numeric agreement still reflects correct reasoning, providing smoother sensitivity to small deviations. In other words, while the Gated version will assign 0 to a sequence of aligning reasoning steps, which resulted in a wrong answer (which can be a case if an LLM fails mathematics behind the solution), Soft version will still give a higher score.

- **DTW Precision, Recall, and F1.** In addition to the normalized alignment score, we derive DTW-based precision, recall, and F1 measures that quantify step-level coverage under the

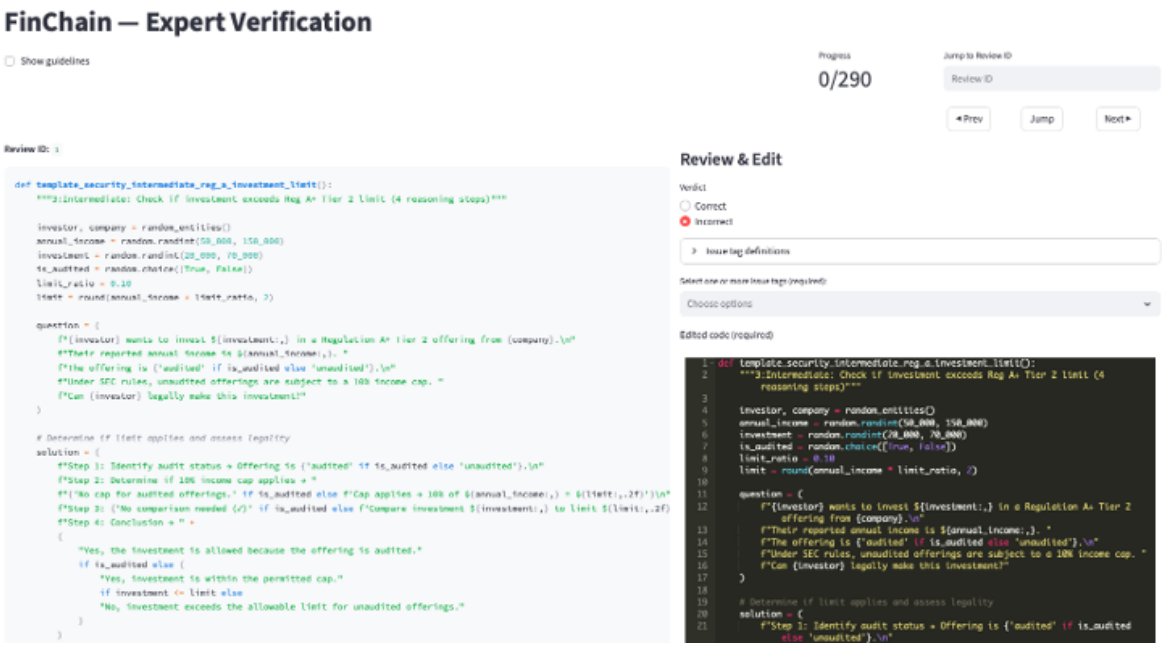

Figure 6: Expert annotation interface. Annotators review each template, assign a verdict, select issue tags, and provide minimal code corrections.

Metric	Spearman ρ
DTWNormGate	**0.640**
DTW Precision (Soft)	0.625
DTW Precision (Gate)	0.622
DTW F1 (Soft)	0.619
DTW F1 (Gate)	0.618
Step Precision (non-DTW)	0.604
DTWNormSoft	0.592
DTW Recall (Gate)	0.573
Step Recall (non-DTW)	0.570
DTW Avg. Path Score (Gate)	0.529
DTW Avg. Path Score (Soft)	0.526
DTW Recall (Soft)	0.512
ROUGE-2	0.469
ROUGE-L	0.434

Table 4: Spearman correlation with expert process scores across all metrics.

DTW alignment path. These provide a finer breakdown of reasoning alignment quality.

G.2 Comparative Evaluation

We also evaluated a range of traditional text-similarity and reasoning metrics, including ROUGE-2, ROUGE-L, step-level precision and recall (marked as 'non-DTW' in the table). Each metric was correlated with expert-assigned *Reasoning Process Quality* score. Table 4 summarizes the top Spearman correlations with expert process judgments.

G.3 Discussion

The DTW-based variants consistently achieve the highest correlation with expert judgments, with the **Normalized DTW Alignment Score (Gate Mode)** emerging as the most reliable indicator of reasoning faithfulness. The "Soft" variant yields slightly lower but still strong correlations, suggesting that the gating formulation better captures strict consistency, while the soft variant provides smoother sensitivity to near-correct reasoning. Compared to traditional metrics such as ROUGE or simple step-level precision and recall, DTW captures not only semantic similarity but also the structural coherence and numerical consistency of reasoning chains. These results highlight the usefulness of our proposed metric.