

# UniFinEval: Towards Unified Evaluation of Financial Multimodal Models across Text, Images and Videos

Anonymous ACL submission

## Abstract

Multimodal large language models are playing an increasingly significant role in empowering the financial domain, however, the challenges they face, such as multimodal and high-density information and cross-modal multi-hop reasoning, go beyond the evaluation scope of existing multimodal benchmarks. To address this gap, we propose UniFinEval, the first unified multimodal benchmark designed for high-information-density financial environments, covering text, images, and videos. UniFinEval systematically constructs five core financial scenarios grounded in real-world financial systems: Financial Statement Auditing, Company Fundamental Reasoning, Industry Trend Insights, Financial Risk Sensing, and Asset Allocation Analysis. We manually construct a high-quality dataset consisting of 3,767 question-answer pairs in both Chinese and English and systematically evaluate 10 mainstream MLLMs under Zero-Shot and CoT settings. Results show that Gemini-3-pro-preview achieves the best overall performance, yet still exhibits a substantial gap compared to financial experts. Further error analysis reveals systematic deficiencies in current models. UniFinEval aims to provide a systematic assessment of MLLMs' capabilities in fine-grained, high-information-density financial environments, thereby enhancing the robustness of MLLMs applications in real-world financial scenarios. Data and code are available at <https://anonymous.4open.science/r/anonym4B75>.

## 1 Introduction

In the current era of rapid multimodal large language models (MLLMs) development and explosive information growth, the demand for utilizing large language models (LLMs) to process information across various fields is becoming increasingly urgent (Rahman et al., 2025; Li et al., 2025c; Zhang

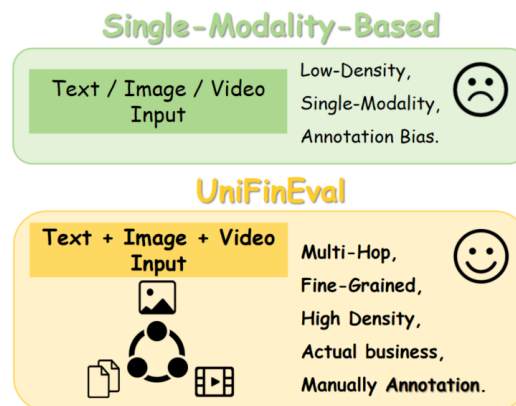


Figure 1: UniFinEval is manually constructed and supports full-modality inputs including text, images, and videos. It is equipped with cross-modal reasoning capabilities and features high information density while closely aligning with real financial business practices.

et al., 2025a; Dennstädt et al., 2025; Xia et al., 2024; Das et al., 2025). In real-world financial scenarios, models are often required to simulate the role of analysts, simultaneously processing voluminous financial and research reports, understanding implicit correspondences between charts and text, and conducting continuous analysis by incorporating financial analysis videos. However, there is a significant misalignment between existing multimodal financial benchmarks and these real-world demands. As illustrated in Figure 1 these limitations are primarily manifested in the following aspects: On one hand, existing benchmarks are limited to single-modality evaluations, such as FinVQA (Chen et al., 2021), Fin-Fact (Rangapur et al., 2025), and MMMU (Yue et al., 2024), and have not yet extended to broader multi-modal tasks. This modality-restricted approach leads to inconsistencies in characterizing model capabilities and limits the reference value of evaluation results for real-world financial applications. On the other hand, the high-density and noisy information environment of

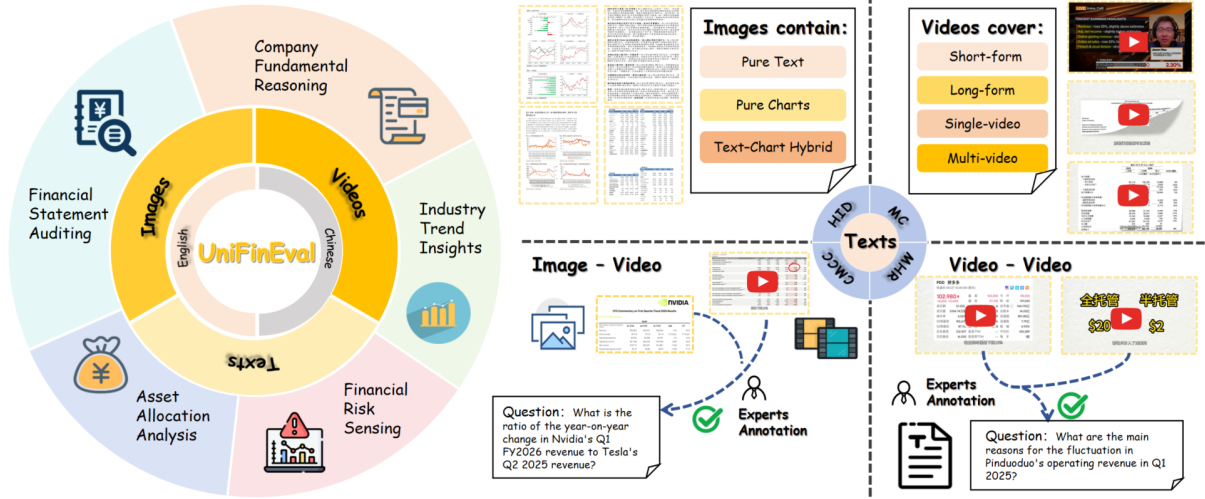


Figure 2: UniFinEval covers five major financial scenarios and constructs datasets spanning text, images, videos, as well as multiple cross-modal combinations. It features high-information-density and manually construct data, together with dedicated designs for cross-modal consistency checking and multi-Hop reasoning, providing comprehensive support for MLLMs evaluation in financial domains.

065 real financial markets places higher demands on  
 066 models' fine-grained analysis capabilities. How-  
 067 ever, existing studies mostly use LLMs to construct  
 068 datasets based on simplified or truncated data (Luo  
 069 et al., 2025; Liu et al., 2025d). This not only makes  
 070 it difficult to evaluate the comprehensive perfor-  
 071 mance of models in high-information-density en-  
 072 vironments but also introduces potential annota-  
 073 tion errors, thereby affecting the reliability of the  
 074 evaluation conclusions. These bias makes it dif-  
 075 ficult for a model's performance on benchmarks  
 076 to reflect its capability boundaries in real financial  
 077 operations. Consequently, this may lead to unsta-  
 078 ble analysis conclusions, exposure to compliance  
 079 risks, and even severe financial losses in practical  
 080 applications (Lee et al., 2024).

081 To fill these gaps, we propose UniFinEval,  
 082 the first financial unified benchmark designed for  
 083 high-information-density financial environments,  
 084 integrating text, images, and videos modalities.  
 085 UniFinEval grounded in real-world financial busi-  
 086 ness workflows and systematically covers five core  
 087 scenarios: Financial Statement Auditing, Company  
 088 Fundamental Reasoning, Industry Trend Insights,  
 089 Financial Risk Sensing, and Asset Allocation Anal-  
 090 ysis, completely characterizing complete cognitive  
 091 loop from multimodal perception and cross-modal  
 092 alignment to high-level decision making. All sam-  
 093 ples in this benchmark were manually constructed  
 094 by financial experts, constructing a dataset of 3,767  
 095 high-quality Q&A pairs aimed at assessing the ca-  
 096 pability boundaries of MLLMs in real financial

097 scenarios. The main contributions of this paper are  
 098 summarized as follows:

099 (1) We propose UniFinEval, the first multimodal  
 100 unified financial evaluation benchmark. By con-  
 101 structing **manually curated** question-answer pairs  
 102 that deeply integrate text, images, and videos, we  
 103 provide a unified paradigm for analyzing the pro-  
 104 found capabilities of MLLMs within the financial  
 105 domain. The structure is illustrated in Figure 2.

106 (2) We meticulously design multi-hop reasoning  
 107 questions from the unique perspective of multi-  
 108 source information fusion within the financial do-  
 109 main, thereby posing a rigorous challenge to the  
 110 actual performance of MLLMs.

111 (3) We conducted a targeted error analysis aimed  
 112 at optimizing the capabilities of MLLMs within  
 113 high-information-density complex application sce-  
 114 narios in the financial domain, providing practical  
 115 and feasible directions for improvement.

116 The organization of this paper is as follows: Sec-  
 117 tion 2 reviews the research on financial MLLMs  
 118 and related multimodal benchmarks. Section 3  
 119 details the construction process of UniFinEval, in-  
 120 cluding data collection, question design, and qual-  
 121 ity control processes. Sections 4 and 5 present our  
 122 experimental settings and results under different  
 123 difficulty levels, followed by error analysis. Fi-  
 124 nally, Section 6 summarizes the work of this paper  
 125 and discusses potential future directions in the field  
 126 of multimodal financial intelligence.

127  
128  
129  
130  
131  
132  
133  
134  
135  
136  
137  
138  
139  
140  
141  
142  
143  
144  
145  
146  
147  
148  
149  
150  
151  
152  
153  
154  
155  
  
156  
157  
  
158  
159  
160  
161  
162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
174  
175  
176

## 2 Related Work

### 2.1 Financial Applications Analysis

Benefiting from breakthrough advancements in general Large Language Models (LLMs), research on intelligence in the financial vertical domain has made significant progress. Early work mainly focused on tasks such as financial sentiment analysis (Delgadillo et al., 2024; Kim et al., 2024; Iacovides et al., 2024), financial time series analysis (Li et al., 2024b,c; Wang et al., 2024b; Mai, 2024; Cao et al., 2024), financial text understanding (Masry and Hajian, 2024; Wilson et al., 2024), and decision support (Yu et al., 2024c,b; Liu et al., 2025c; Xiao et al., 2025). However, real-world financial information inherently possesses highly multimodal characteristics. To further enhance the deployment capabilities of models in real scenarios, the focus of related research has gradually expanded from a single text modality to a multimodal perspective. FinVis-GPT (Wang et al., 2023), Finzero (Wang et al., 2025b), and FinTral (Bhatia et al., 2024) achieve deep analysis of financial charts. FinAgent (Zhang et al., 2024a) created a multimodal agent capable of autonomous trading decisions, MM-DREX (Chen et al., 2025b) utilizes multimodal information to implement dynamic routing for expert trading; and AT-FinGPT (Liu et al., 2025b) utilizes voice features to assist in financial risk prediction.

### 2.2 Multimodal Financial Benchmark Analysis

Compared to the rapid expansion of model capabilities, existing financial multimodal evaluation systems appear significantly lagged. On one hand, traditional financial benchmarks are confined to the text modality (Zhu et al., 2024; Nie et al., 2024; Zhao et al., 2024; Wang et al., 2024a; Reddy et al., 2024; Chen et al., 2024, 2025a; Guo et al., 2025; Liu et al., 2025a; Li et al., 2025a; Matlin et al., 2025), while general multimodal benchmarks lack domain-specific knowledge depth (Liu et al., 2024; Li et al., 2024a; Zhang et al., 2024b; Yu et al., 2024a); neither can accurately measure the financial multimodal capabilities of financial large models. On the other hand, although current financial multimodal evaluation benchmarks have made progress in specific tasks, they still exhibit significant limitations. Specifically, benchmarks such as MME-Finance (Gan et al., 2024), FinChart-Bench (Shu et al., 2025), CFBenchmark-MM (Li

et al., 2025b), and MultiFinBen (Peng et al., 2025) mainly focus on static chart understanding, lacking consideration for dynamic time-series information such as financial videos. Data for works like FinMR (Deng et al., 2025), FAMMA (Xue et al., 2024), and XFinBench (Zhang et al., 2025b) are mostly derived from textbooks or standardized exams, creating a large gap with real financial environments full of noise and unstructured information. Works like FinMMR (Tang et al., 2025) and FinMultiTime (Xu et al., 2025) utilize large model annotation paradigms, potentially introducing model bias and hallucination risks. Furthermore, VideoConviction (Galarnyk et al., 2025) and Fincap (Sukhani et al., 2025), among the few benchmarks involving financial videos, neither evaluate financial business competence nor assess joint reasoning over video and text-image information.

## 3 UniFinEval

### 3.1 Overview

To evaluate the performance and failure modes of MLLMs in real-world financial business environments, we propose **UniFinEval**, a benchmark designed around complex and information-dense financial scenarios. Guided by financial experts and grounded in authentic business practices, UniFinEval identifies five representative and critical financial scenarios: Financial Statement Auditing (FSA), Company Fundamental Reasoning (CFR), Industry Trend Insights (ITI), Financial Risk Sensing (FRS), and Asset Allocation Analysis (AAA).

These scenarios cover a broad spectrum of financial reasoning tasks, ranging from fine-grained information verification and consistency checking to cross-firm and cross-cycle analysis, and ultimately high-level risk control and decision-making. Such a progression reflects increasing requirements on models' information perception, cross-modal alignment, multi-hop reasoning, and decision robustness. In addition, UniFinEval incorporates a series of task-specific perturbations aligned with real-world financial business settings into each scenario. These perturbations enable the evaluation of model behavior under non-ideal visual and textual conditions commonly encountered in practical financial applications; details are provided in Appendix A.4.

Centered on these business scenarios, UniFinEval is guided throughout by financial experts to construct a data system featuring

177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
190  
191  
192  
193  
194  
  
195  
196  
  
197  
198  
199  
200  
201  
202  
203  
204  
205  
206  
207  
208  
209  
210  
211  
212  
213  
214  
215  
216  
217  
218  
219  
220  
221  
222  
223  
224  
225  
226

high information density and deep multi-modal fusion. The benchmark facilitates multi-hop integration across three foundational modalities—text, images, and video—by systematically incorporating complex cross-modal combinations, including text–image, text–video, image–video, and text–image–video, as illustrated in Figure 3. UniFinEval authentically replicates the parallel information structure of professional financial workflows. Through a rigorous multi-layer quality control process, we construct a dataset comprising 3,767 high-quality bilingual (Chinese and English) Q&A pairs. A detailed breakdown of data distribution across the five financial scenarios is provided in Table 3 in the Appendix. Regarding task design, the benchmark incorporates both *single-turn* and *multi-turn* Q&A mechanisms and extensively integrates cross-modal multi-hop reasoning requirements. Such a design enables the assessment of the model’s comprehensive capabilities in immediate understanding, cross-turn context integration, and complex logic chain construction, providing a solid foundation for systematically evaluating the performance of MLLMs in real financial environments.

### 3.2 Scenario Construction

The scenario construction of UniFinEval is not centered on isolated tasks or individual model capabilities, but is designed to align with real-world financial business processes. In practical financial settings, performance on a single capability is often insufficient to support reliable deployment; instead, model stability and consistency across multiple business stages are critical considerations for real applications. Motivated by this observation, UniFinEval establishes five interconnected and progressively hierarchical financial business scenarios. As the scenarios advance, the object of analysis expands from localized information processing to global judgment, while data forms evolve from static inputs to dynamic multi-modal combinations. Accordingly, the demands placed on the model increase across stages, and performance at each step provides insight into the challenges and feasibility of deploying MLLMs in realistic financial environments.

**Financial Statement Auditing:** As the primary stage of financial analysis, this scenario focuses on verifying the accuracy and consistency of financial report information, with the goal of establishing a reliable data foundation for subse-

quent end-to-end analysis. To address the practical challenge faced by practitioners when verifying information in complex visual environments, we design high–information-density images that integrate textual content, charts, and their coupled representations. Authentic layouts and redundant information are deliberately preserved to reflect real-world business settings. Models are required to complete key information extraction and verification through both single-turn and multi-turn Q&A, enabling the evaluation of their ability to precisely locate and validate critical financial facts under conditions of high information density, complex layouts, and redundant interference.

**Company Fundamental Reasoning:** Following the FSA scenario, the CFR scenario focuses on analyzing corporate operating conditions and intrinsic value. To align with real-world practice in which practitioners examine financial reports alongside research reports, the data in this scenario continue to use text and charts as primary information carriers, while placing greater emphasis on financial variables and calculation bases distributed across heterogeneous modalities. The task design centers on the derivation of core financial indicators, requiring the model to extract relevant parameters from multi-source information and perform rigorous multi-step calculations. By incorporating high-difficulty cross-modal multi-hop reasoning, the scenario effectively differentiates basic information retrieval ability from deeper financial mathematical and logical reasoning capabilities.

**Industry Trend Insights:** The ITI scenario elevates the analytical perspective from individual enterprises to the industry level, focusing on cross-enterprise and cross-period analysis to evaluate the model’s ability to integrate multi-source information and infer industry logic. To reflect real-world demands for cross–data-source industry research, the data construction integrates multi-period financial reports, industry research reports, and macroeconomic data. A broader informational context is constructed using text and trend charts as primary carriers. The tasks center on industry trend assessment and cycle summarization, requiring complex cross-enterprise comparative reasoning and multi-hop inference. Through single-turn and multi-turn interactions, the scenario simulates how practitioners extract structured industry insights from fragmented information.

**Financial Risk Sensing:** The scenario focuses on the multi-dimensional identification and

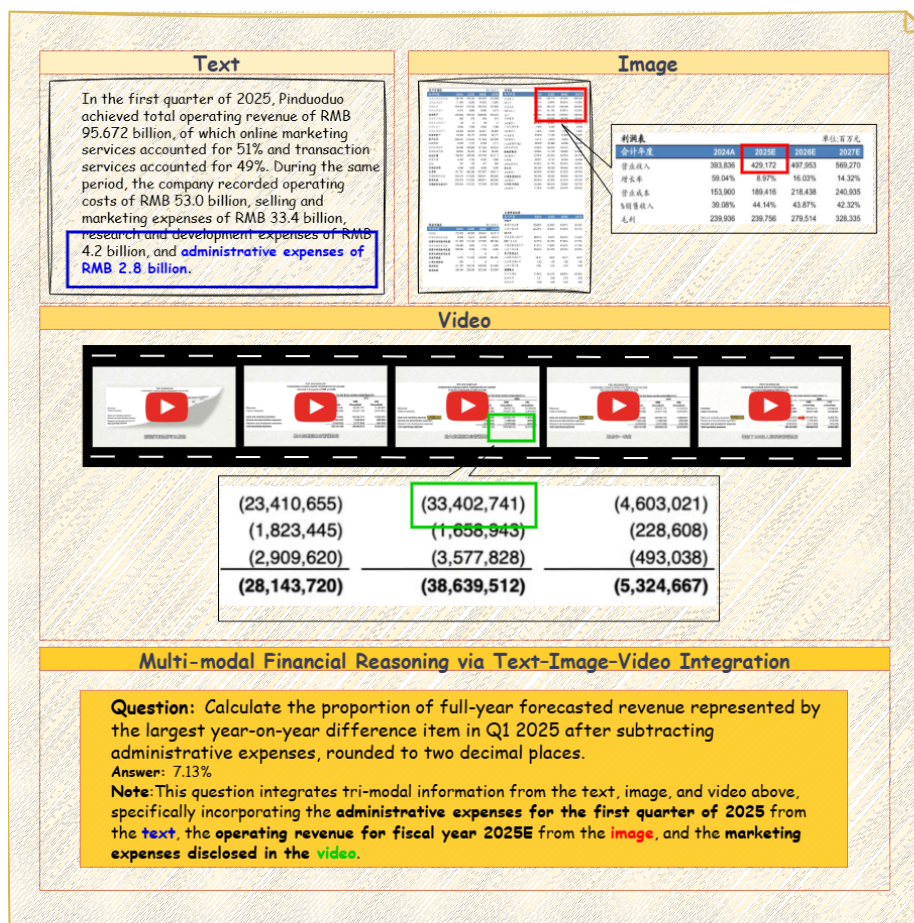


Figure 3: An example of a cross-modal multi-hop question in UniFinEval. The answer is derived from the acquisition and integration of key information from the presented text, images, and videos.

integration of potential risk signals, serving as a critical component in safeguarding investment decisions within financial analysis. It emphasizes evaluating the model’s robustness and risk perception capabilities under conditions of dynamic multi-modal information fusion. Since risk signals are often implicit and embedded in unstructured and time-varying information, the scenario introduces financial analysis videos, constructing a *text–image–video* multi-modal setting. Dynamic viewpoints presented in videos are explicitly aligned with written quantitative data, restoring the multi-source cross-validation requirements characteristic of real-world risk analysis. The tasks center on risk signal identification and cross-modal consistency verification, incorporating cross-modal multi-hop reasoning. This design requires the model to jointly assess heterogeneous information sources and accurately capture latent downside risk signals.

**Asset Allocation Analysis:** As the final decision-making stage in financial workflows, As-

set Allocation Analysis requires synthesizing insights from all preceding analytical stages to formulate executable strategies under multiple constraints. In this scenario, UniFinEval constructs the most complex input structure, integrating multi-modal data from all prior scenarios while further introducing mixed *text–image–video* inputs and explicit real-world constraints, thereby closely approximating the information boundaries of realistic decision-making processes. The task design follows the complete asset allocation pipeline and primarily adopts multi-turn Q&A interactions, requiring the model to iteratively integrate prior analytical results to produce logically consistent allocation strategies. This scenario ultimately evaluates the model’s comprehensive decision-making ability and its capacity to manage trade-offs under high information density and complex constraints.

*Overall*, through the progressive construction of the five scenarios described above, UniFinEval establishes a comprehensive evaluation framework

that spans from static high-density perception, to dynamic cross-modal reasoning, and ultimately to expert-level decision support. As the scenarios advance, data modalities become increasingly diverse and task structures grow more complex, enabling fine-grained and interpretable quantitative assessment of model capabilities across different cognitive levels. This framework provides a reliable reference for evaluating the practical applicability of MLLMs in real-world financial business environments.

### 3.3 Quality Control

Given the instability and potential bias of current multimodal models in complex financial tasks, as well as the financial industry’s stringent requirements for precision, interpretability, and logical consistency, UniFinEval adopts a fully expert-driven manual construction strategy for question generation and validation. To eliminate any form of model-induced bias, no LLMs are involved at any stage of Q&A creation, ensuring that all evaluation tasks strictly adhere to real-world financial business logic and regulatory compliance standards.

The dataset construction is carried out by a dedicated team of ten senior financial experts, including PhD students from top-tier finance and economics institutions and experienced practitioners from leading securities firms and financial institutions. All experts hold professional certifications such as CFA or CPA, with an average of more than five years of frontline industry or research experience. This dual composition of academic and industry expertise ensures both theoretical rigor and practical relevance in scenario design.

Question construction follows a rigorous four-stage quality control pipeline. First, during data filtering, automated scripts combined with manual inspection are used to select multimodal financial materials with high information density and strong business relevance. Second, domain experts independently formulate questions and corresponding standard answers based on the curated data. Third, all annotations undergo cross-validation by other experts to identify logical inconsistencies, ambiguity, or deviations from realistic business practices. Finally, a unified audit is conducted to ensure that the resulting tasks faithfully reflect the complexity, decision constraints, and reasoning processes encountered in real financial environments. Detailed descriptions of the quality control procedures are provided in Appendix A.3.

## 4 Experiments Settings

### 4.1 Baseline Models

We tested 10 mainstream MLLMs. Closed-source models were accessed via their respective APIs, while open-source models were deployed locally. All inference tasks were run on 8×NVIDIA A800 GPUs, using vLLM for efficient local deployment and inference (Llama used LMDeploy). The evaluation encompassed 4 closed-source models, including GPT-5.1 (OpenAI, 2025), Gemini-3-pro-preview (Google, 2025), Grok-4.1 (xAI, 2025), and Claude-Sonnet-4.5 (Anthropic, 2025), alongside 6 open-source models from multiple mainstream MLLMs, including Qwen3-VL-235B-A22B-Thinking, Qwen3-VL-32B-Thinking (Yang et al., 2025), InternVL3.5-241B-A28B, InternVL3.5-30B-A3B (Wang et al., 2025a), MiniCPM-V-4.5 (Yu et al., 2025), and Llama-3.2-11B-Vision-Instruct (MetaAI, 2025). For more details on the models, please see Appendix B.

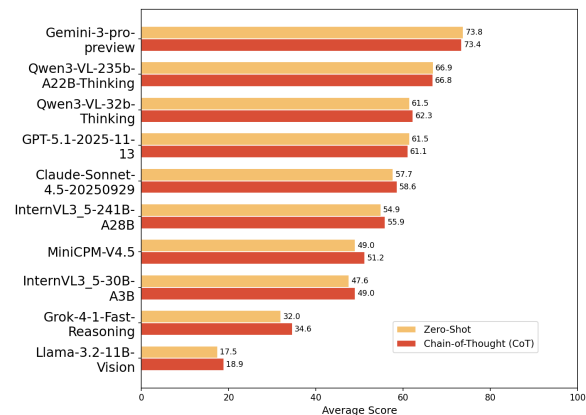


Figure 4: As evident from the visualization of the result comparisons, the performance of the vast majority of models achieved a slight improvement under the CoT evaluation setting, though the overall magnitude of this enhancement remains relatively limited.

### 4.2 Evaluation Methods

We employ two core evaluation settings: Zero-Shot and Zero-Shot CoT (abbreviated as CoT). To ensure robust assessment despite the limitations of rule-based matching in complex financial reasoning, we also integrate Qwen-Max (Yang et al., 2025) to standardize output extraction and evaluation. It should be noted that a manual inspection of 30% of the judging results from Qwen-Max revealed an average error rate of less than 1%. Additionally, Accuracy is adopted as the core metric to ensure an objective and scalable evaluation.

Table 1: Performance of in Zero-shot and CoT settings on UniFinEval(%). FSA stands for Financial Statement Auditing, CFR denotes Company Fundamental Reasoning, ITI refers to Industry Trend Insights, FRS represents Financial Risk Sensing, and AAA corresponds to Asset Allocation Analysis. The highest values in each column are highlighted with a **blue background**, while the second-best results are underlined.

Model	FSA		CFR		ITI		FRS		AAA		Average	
	Zero-Shot	CoT	Zero-Shot	CoT	Zero-Shot	CoT	Zero-Shot	CoT	Zero-Shot	CoT	Zero-Shot	CoT
Gemini-3-pro-preview	<b>83.5</b>	<b>83.8</b>	<b>82.2</b>	<b>82.8</b>	<b>73.3</b>	<b>74.7</b>	<b>68.8</b>	<b>70.1</b>	<b>61.1</b>	<b>55.4</b>	<b>73.8</b>	<b>73.4</b>
Qwen3-VL-235B-A22B-Thinking	<u>80.2</u>	<u>81.3</u>	<u>78.9</u>	<u>74.9</u>	<u>69.4</u>	<u>64.6</u>	<u>62.9</u>	<u>62.7</u>	43.3	<u>50.3</u>	<u>66.9</u>	<u>66.8</u>
Qwen3-VL-32B-Thinking	75.1	76.2	71.0	70.3	65.6	<u>65.2</u>	54.8	56.6	40.8	43.3	61.5	62.3
GPT-5.1	76.9	77.8	67.1	65.0	65.8	60.4	50.0	54.1	<u>47.8</u>	48.4	61.5	61.1
Claude-Sonnet-4.5	70.8	71.9	65.4	68.2	61.7	61.4	50.0	50.6	40.8	42.0	57.7	58.6
InternVL3.5-241B-A28B	69.0	70.6	66.2	68.7	63.8	63.8	37.1	36.2	38.2	40.1	54.9	55.9
MiniCPM-V-4.5	65.9	66.2	62.3	64.1	53.2	57.9	30.6	38.0	33.1	29.9	49.0	51.2
InternVL3.5-30B-A3B	61.5	61.7	64.7	59.9	50.0	52.7	33.9	35.8	28.0	34.4	47.6	49.0
Grok-4.1-Fast-Reasoning	50.3	52.5	43.1	44.1	32.5	34.9	16.1	19.3	17.8	22.3	32.0	34.6
Llama-3.2-11B-Vision	22.2	23.1	20.9	23.7	19.0	21.4	14.1	15.7	11.5	10.8	17.5	18.9
Expert	97.5		95.3		90.1		88.5		85.2		91.3	

## 5 Results

### 5.1 Main Result

Table 1 presents the specific performance and average results of each participating model across various tasks under both Zero-Shot and CoT settings, while Figure 4 provides a more intuitive visualization for comparison. Under the Zero-Shot setting, Gemini-3-pro-preview demonstrates a stable and comprehensive performance advantage across all tasks, ranking first with an average accuracy of 73.8%. Qwen3-VL-235B-Thinking follows closely with an average accuracy of 66.9%, a gap of only 6.9%. This result suggests that the technical divide between current large-scale open-source models and closed-source models is progressively narrowing. Qwen3-VL-32B-Thinking and GPT-5.1 constitute the second tier; both perform robustly in perception-based tasks and certain reasoning tasks but still face significant bottlenecks in high-level decision-making. Notably, these two models achieve identical average accuracies, not only surpassing Claude-Sonnet-4.5 and InternVL3.5-241B-A28B but also demonstrating superior performance across various sub-tasks, highlighting their relatively stronger comprehensive competitiveness. In contrast, the remaining models exhibit a clear performance gap, a stratification effect that is further amplified across different financial tasks. Model performance shows a significant gradient decay as task cognitive complexity increases. In perception-oriented tasks such as FSA and CFR, the performance gap between top-tier models and human experts is narrow, with models able to consistently identify most key information within complex research reports and charts. However, as tasks

shift from explicit information recognition to the repeated verification of cross-modal information, all models experience a marked decline in performance. Even the top-performing Gemini-3-pro-preview achieves only 73.3% accuracy in ITI scenarios, significantly lower than the 90.1% achieved by human experts, exposing deficiencies in building consistent semantic mappings between different financial modalities. In the FRS task, which incorporate video modalities, most models fail to achieve a performance breakthrough, indicating a clear weakness in their ability to model logic across the temporal dimension.

When task complexity further escalates to AAA tasks, the performance shortfalls of all models are fully exposed: Gemini-3-pro-preview achieves an accuracy of only 61.1%, while the performance of other models is even more inferior, creating a significant gap compared to the 85.2% accuracy of human experts. This result clearly illustrates that although some models demonstrate strong capabilities in preliminary perception and reasoning tasks, they struggle to maintain long-term, stable logical consistency in complex financial multimodal scenarios with high information density. In contrast, human experts maintain a significant advantage across all task scenarios, and this advantage continues to expand as tasks progress from the perception level to the decision-making level. This profoundly reveals the critical capability gap that exists between current MLLMs and real-world financial experts.

From the perspective of comparison with the CoT setting, most models achieve a certain degree of performance improvement, though the magnitude of this improvement remains relatively limited.

This phenomenon suggests that the vast majority of current MLLMs are already capable of adapting to explicit reasoning mechanisms; even without embedded explicit reasoning instructions, the comprehensive performance of these models can be relatively fully unleashed.

### 5.2 Error Analysis

We conducted a systematic error analysis by randomly sampling approximately 50% of all incorrect predictions generated by the evaluated MLLMs. The overall error types are categorized into five core dimensions: Perception, Hallucination, Knowledge, Cross-modal, and Computation. These correspond to specific errors including Financial Image Perception and Data Interpretation (FIPDI), Inconsistent Financial Reasoning and Hallucinations (IFRH), Financial Knowledge Reasoning and Domain-Specific Understanding (FKRDU), Cross-modal Data Integration and Alignment (CDIA), and Financial Computation and Numerical Analysis (FCNA).

As illustrated in Figure 5, several models still exhibit significant deficiencies in numerical computation; specifically, Qwen3-VL-32B-Thinking shows a markedly higher proportion of errors in computational tasks compared to other evaluated models. Aside from the prominent hallucination issues observed in Llama-3.2-11B-Vision and the substantial share of financial knowledge reasoning errors in Qwen3-235B-VL-A22B-Thinking, the errors of the vast majority of models are concentrated in two dimensions: image content perception and cross-modal data alignment. This phenomenon clearly demonstrates that current MLLMs still possess obvious capability shortfalls when confronting the high information density and time-sensitive complexities of financial market environments. They remain unable to fully adapt to multi-hop reasoning tasks with fine-grained requirements, failing to satisfy the rigorous demands for precision and stability inherent in the financial domain. Given that the overall performance of Grok-4.1 fell short of expectations, we conducted a targeted error analysis in Figure 21 in Appendix B.1 to substantiate the finding that some general-purpose MLLMs exhibit weak adaptability in the financial domain. Simultaneously, Appendix B.1 provides supplementary case studies of representative error types to further support our research conclusions.

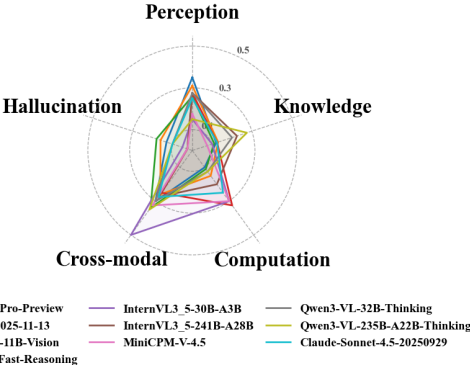


Figure 5: The radar chart summarizes the relative proportions of five major categories of errors observed in incorrect model predictions. Each axis reflects the proportion of a specific error type among all erroneous cases for a given model, highlighting differences in error concentration and reasoning weaknesses across models.

### 6 Conclusion

We proposed UniFinEval, a high-information-density multimodal benchmark oriented towards real financial business scenarios, aimed at systematically characterizing the capability boundaries of MLLMs in financial scenarios. UniFinEval centers on five scenarios: Financial Statement Auditing, Company Fundamental Reasoning, Industry Trend Insights, Financial Risk Sensing, and Asset Allocation Analysis, covering the complete financial analysis link from basic perception to expert-level decision support. The benchmark was constructed entirely manually by financial domain experts; all questions align strictly with real business logic and explicitly introduce cross-modal consistency judgment and multi-hop information integration, thereby achieving performance evaluation of actual business capabilities. Experimental results show that current MLLMs possess strong capabilities in information extraction and fact recognition tasks, but their performance remains significantly limited and lags notably behind financial experts when performing cross-modal semantic alignment, subjective judgment understanding, and complex financial decision-making under high-information-density conditions. Error analysis further revealed six major shortcomings of MLLMs in real financial scenarios. Overall, UniFinEval contributes a unified and practical evaluation framework for systematically assessing the real-world effectiveness of MLLMs in financial scenarios.

## 606 Limitations

607 Although UniFinEval strives to align with real fi-  
608 nancial business processes in scenario design and  
609 capability characterization, and has made break-  
610 through progress in evaluating MLLMs in the fi-  
611 nancial domain, some limitations remain. First  
612 although the benchmark systematically introduces  
613 cross-modal consistency, multi-hop reasoning, and  
614 multi-turn information integration requirements in  
615 various scenarios, current tasks are still dominated  
616 by offline Q&A forms, not yet fully simulating the  
617 complex closed loop of long-duration interaction  
618 and dynamic decision feedback in real analysis pro-  
619 cesses. Second, evaluation results mainly focus  
620 on the correctness and consistency of model out-  
621 puts, leaving room for deepening the analysis of  
622 the interpretability of model reasoning paths.

## 623 References

624 Anthropic. 2025. Claude-sonnet-4.5. [https://www.](https://www.anthropic.com/claude/sonnet)  
625 [anthropic.com/claude/sonnet](https://www.anthropic.com/claude/sonnet).

626 Gagan Bhatia, El Moatez Billah Nagoudi, Hasan Cavu-  
627 soglu, and Muhammad Abdul-Mageed. 2024. Fintral:  
628 A family of gpt-4 level multimodal financial large lan-  
629 guage models. *arXiv preprint arXiv:2402.10986*.

630 Yupeng Cao, Zhi Chen, Prashant Kumar, Qingyun Pei,  
631 Yangyang Yu, Haohang Li, Fabrizio Dimino, Lorenzo  
632 Ausiello, KP Subbalakshmi, and Papa Momar Ndi-  
633 aye. 2024. Risklabs: predicting financial risk using  
634 large language model based on multimodal and multi-  
635 sources data. *arXiv preprint arXiv:2404.07452*.

636 Jialin Chen, Aosong Feng, Ziyu Zhao, Juan Garza,  
637 Gaukhar Nurbek, Cheng Qin, Ali Maatouk, Lean-  
638 dros Tassioulas, Yifeng Gao, and Rex Ying. 2025a.  
639 Mtbench: A multimodal time series benchmark for  
640 temporal reasoning and question answering. *arXiv*  
641 *preprint arXiv:2503.16858*.

642 Jian Chen, Peilin Zhou, Yining Hua, Yingxin Loh,  
643 Kehui Chen, Ziyuan Li, Bing Zhu, and Junwei  
644 Liang. 2024. Fintextqa: A dataset for long-  
645 form financial question answering. *arXiv preprint*  
646 *arXiv:2405.09980*.

647 Yang Chen, Yueheng Jiang, Zhaozhao Ma, Yuchen Cao,  
648 Jacky Keung, Kun Kuang, Leilei Gan, Yiquan Wu,  
649 and Fei Wu. 2025b. Mm-drex: Multimodal-driven  
650 dynamic routing of llm experts for financial trading.  
651 *arXiv preprint arXiv:2509.05080*.

652 Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena  
653 Shah, Iana Borova, Dylan Langdon, Reema Moussa,  
654 Matt Beane, Ting-Hao Huang, Bryan Routledge,  
655 and 1 others. 2021. Finqa: A dataset of numerical  
656 reasoning over financial data. *arXiv preprint*  
657 *arXiv:2109.00122*.

Sarmistha Das, R. E. Zera Marveen Lyngkhoi, Sriparna  
658 Saha, and Alka Maurya. 2025. [Unlocking financial](#)  
659 [insights: An advanced multimodal summarization](#)  
660 [with multimodal output framework for financial advi-](#)  
661 [sory videos](#). In *Proceedings of the 33rd ACM Inter-*  
662 *national Conference on Multimedia, MM '25*, page  
663 11976–11985, New York, NY, USA. Association for  
664 Computing Machinery. 665

Josiel Delgadillo, Johnson Kinyua, and Charles  
666 Mutigwe. 2024. Finsosent: Advancing financial mar-  
667 ket sentiment analysis through pretrained large lan-  
668 guage models. *Big Data and Cognitive Computing*,  
669 8(8):87. 670

Shuangyan Deng, Haizhou Peng, Jiachen Xu, Rui Mao,  
671 Ciprian Doru Giurcaneanu, and Jiamou Liu. 2025.  
672 Finmr: A knowledge-intensive multimodal bench-  
673 mark for advanced financial reasoning. In *Proceed-*  
674 *ings of the 6th ACM International Conference on AI*  
675 *in Finance*, pages 168–176. 676

Fabio Dennstädt, Luc Lerch, Max Schmerder, Nikola  
677 Cihoric, Grazia Maria Cereghetti, Roberto Gaio, Har-  
678 ald Bonel, Irina Filchenko, Janna Hastings, Florian  
679 Dammann, and 1 others. 2025. A comparative per-  
680 formance analysis of regular expressions and an llm-  
681 based approach to extract the bi-rads score from radi-  
682 ological reports. *medRxiv*, pages 2025–06. 683

Michael Galarnyk, Veer Kejriwal, Agam Shah, Yash  
684 Bhardwaj, Nicholas Watney Meyer, Anand Krishnan,  
685 and Sudheer Chava. 2025. Videoconviction: A mul-  
686 timodal benchmark for human conviction and stock  
687 market recommendations. In *Proceedings of the 31st*  
688 *ACM SIGKDD Conference on Knowledge Discovery*  
689 *and Data Mining V. 2*, pages 5447–5458. 690

Ziliang Gan, Yu Lu, Dong Zhang, Haohan Li, Che Liu,  
691 Jian Liu, Ji Liu, Haipang Wu, Chaoyou Fu, Zenglin  
692 Xu, and 1 others. 2024. Mme-finance: A multimodal  
693 finance benchmark for expert-level understanding  
694 and reasoning. *arXiv preprint arXiv:2411.03314*. 695

Google. 2025. Gemini-3-pro-preview. [https://](https://gemini.google.com/)  
696 [gemini.google.com/](https://gemini.google.com/). 697

Xin Guo, Haotian Xia, Zhaowei Liu, Hanyang Cao, Zhi  
698 Yang, Zhiqiang Liu, Sizhe Wang, Jinyi Niu, Chuqi  
699 Wang, Yanhui Wang, Xiaolong Liang, Xiaoming  
700 Huang, Bing Zhu, Zhongyu Wei, Yun Chen, Weining  
701 Shen, and Liwen Zhang. 2025. [FinEval: A Chinese fi-](#)  
702 [nancial domain knowledge evaluation benchmark for](#)  
703 [large language models](#). In *Proceedings of the 2025*  
704 *Conference of the Nations of the Americas Chap-*  
705 *ter of the Association for Computational Linguistics:*  
706 *Human Language Technologies (Volume 1: Long Pa-*  
707 *pers)*, pages 6258–6292, Albuquerque, New Mexico.  
708 Association for Computational Linguistics. 709

Giorgos Iacovides, Thanos Konstantinidis, Mingxue Xu,  
710 and Danilo Mandic. 2024. Finllama: Llm-based fi-  
711 nancial sentiment analysis for algorithmic trading. In  
712 *Proceedings of the 5th ACM International Confer-*  
713 *ence on AI in Finance*, pages 134–141. 714



828	Llm-based data science agents: A survey of capabilities, challenges, and future directions. <i>arXiv preprint arXiv:2510.04023</i> .	Ezhilan Wilson, Anshul Saxena, Jayant Mahajan, Lekha Panikulangara, Shruti Kulkarni, and Pritty Jain. 2024. Fin2sum: advancing ai-driven financial text summarization with llms. In <i>2024 International Conference on Trends in Quantum Computing and Emerging Business Technologies</i> , pages 1–5. IEEE.	882
829			883
830			884
831	Aman Rangapur, Haoran Wang, Ling Jian, and Kai Shu. 2025. Fin-fact: A benchmark dataset for multimodal financial fact-checking and explanation generation. In <i>Companion Proceedings of the ACM on Web Conference 2025</i> , pages 785–788.	xAI. 2025. Grok 4.1 fast and agent tools api. <a href="https://x.ai/grok">https://x.ai/grok</a> .	885
832			886
833			887
834			888
835			889
836	Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai, Michael Krumdick, Charles Lovering, and Chris Tanner. 2024. Docfinqa: A long-context financial reasoning dataset. <i>arXiv preprint arXiv:2401.06915</i> .	Haotian Xia, Zhengbang Yang, Junbo Zou, Rhys Tracy, Yuqing Wang, Chi Lu, Christopher Lai, Yanjun He, Xun Shao, Zhuoqing Xie, and 1 others. 2024. Sportu: A comprehensive sports understanding benchmark for multimodal large language models. <i>arXiv preprint arXiv:2410.08474</i> .	890
837			891
838			892
839			893
840	Dong Shu, Haoyang Yuan, Yuchen Wang, Yanguang Liu, Huopu Zhang, Haiyan Zhao, and Mengnan Du. 2025. Finchart-bench: Benchmarking financial chart comprehension in vision-language models. <i>arXiv preprint arXiv:2507.14823</i> .	Yijia Xiao, Edward Sun, Tong Chen, Fang Wu, Di Luo, and Wei Wang. 2025. Trading-r1: Financial trading with llm reasoning via reinforcement learning. <i>arXiv preprint arXiv:2509.11420</i> .	894
841			895
842			896
843			897
844			898
845	Siddhant Sukhani, Yash Bhardwaj, Riya Bhadani, Veer Kejriwal, Michael Galarnyk, and Sudheer Chava. 2025. Fincap: Topic-aligned captions for short-form financial youtube videos. <i>arXiv preprint arXiv:2509.25745</i> .	Wenyan Xu, Dawei Xiang, Yue Liu, Xiyu Wang, Yanxiang Ma, Liang Zhang, Shu Hu, Chang Xu, and Jiaheng Zhang. 2025. Finmultitime: A four-modal bilingual dataset for financial time-series analysis. <i>arXiv preprint arXiv:2506.05019</i> .	900
846			901
847			902
848			903
849			904
850	Zichen Tang, Jiacheng Liu, Zhongjun Yang, Rongjin Li, Zihua Rong, Haoyang He, Zhuodi Hao, Xinyang Hu, Kun Ji, Ziyang Ma, and 1 others. 2025. Finmmr: make financial numerical reasoning more multimodal, comprehensive, and challenging. In <i>Proceedings of the IEEE/CVF International Conference on Computer Vision</i> , pages 3245–3257.	Siqiao Xue, Tingting Chen, Fan Zhou, Qingyang Dai, Zhixuan Chu, and Hongyuan Mei. 2024. Famma: A benchmark for financial domain multilingual multimodal question answering. <i>arXiv preprint arXiv:2410.04526</i> .	905
851			906
852			907
853			908
854			909
855			910
856			911
857	Haochen Wang, Kai Hu, Haoyu Dong, and Liangcai Gao. 2024a. Doctabqa: Answering questions from long documents using tables. In <i>International Conference on Document Analysis and Recognition</i> , pages 470–487. Springer.	An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. <i>arXiv preprint arXiv:2505.09388</i> .	912
858			913
859			914
860			915
861			916
862	Saizhuo Wang, Hang Yuan, Lionel M Ni, and Jian Guo. 2024b. Quantagent: Seeking holy grail in trading by self-improving large language model. <i>arXiv preprint arXiv:2402.03755</i> .	Tianyu Yu, Zefan Wang, Chongyi Wang, Fuwei Huang, Wenshuo Ma, Zhihui He, Tianchi Cai, Weize Chen, Yuxiang Huang, Yuanqian Zhao, and 1 others. 2025. Minicpm-v 4.5: Cooking efficient mllms via architecture, data, and training recipe. <i>arXiv preprint arXiv:2509.18154</i> .	917
863			918
864			919
865			920
866	Weiyun Wang, Zhangwei Gao, Lixin Gu, Hengjun Pu, Long Cui, Xingguang Wei, Zhaoyang Liu, Linglin Jing, Shenglong Ye, Jie Shao, and 1 others. 2025a. Internv13. 5: Advancing open-source multimodal models in versatility, reasoning, and efficiency. <i>arXiv preprint arXiv:2508.18265</i> .	Weihaoyu Yu, Zhengyuan Yang, Lingfeng Ren, Linjie Li, Jianfeng Wang, Kevin Lin, Chung-Ching Lin, Zicheng Liu, Lijuan Wang, and Xinchao Wang. 2024a. Mm-vet v2: A challenging benchmark to evaluate large multimodal models for integrated capabilities. <i>arXiv preprint arXiv:2408.00765</i> .	921
867			922
868			923
869			924
870			925
871			926
872	Yanlong Wang, Jian Xu, Fei Ma, Hongkang Zhang, Hang Yu, Tiantian Gao, Yu Wang, Haochen You, Shao-Lun Huang, Danny Dongning Sun, and 1 others. 2025b. Finzero: Launching multi-modal financial time series forecast with large reasoning model. <i>arXiv preprint arXiv:2509.08742</i> .	Yangyang Yu, Haohang Li, Zhi Chen, Yuechen Jiang, Yang Li, Denghui Zhang, Rong Liu, Jordan W Suchow, and Khaldoun Khashanah. 2024b. Finmem: A performance-enhanced llm trading agent with layered memory and character design. In <i>Proceedings of the AAAI Symposium Series</i> , volume 3, pages 595–597.	927
873			928
874			929
875			930
876			931
877			932
878	Ziao Wang, Yuhang Li, Junda Wu, Jaehyeon Soon, and Xiaofeng Zhang. 2023. Finvis-gpt: A multimodal large language model for financial chart analysis. <i>arXiv preprint arXiv:2308.01430</i> .	Yangyang Yu, Zhiyuan Yao, Haohang Li, Zhiyang Deng, Yupeng Cao, Zhi Chen, Jordan W Suchow, Rong Liu, Zhenyu Cui, Zhaozhuo Xu, and 1 others. 2024c. Fincon: A synthesized llm multi-agent system with conceptual verbal reinforcement for enhanced financial decision making. <i>arXiv preprint arXiv:2407.06567</i> .	933
879			934
880			935
881			936
			937
			938

939	Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng,
940	Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang,
941	Weiming Ren, Yuxuan Sun, and 1 others. 2024.
942	Mmmu: A massive multi-discipline multimodal un-
943	derstanding and reasoning benchmark for expert agi.
944	In <i>Proceedings of the IEEE/CVF Conference on Com-</i>
945	<i>puter Vision and Pattern Recognition</i> , pages 9556–
946	9567.
947	Shuo Zhang, Zezhou Huang, and Eugene Wu. 2025a.
948	Data cleaning using large language models. In <i>2025</i>
949	<i>IEEE 41st International Conference on Data Engi-</i>
950	<i>neering Workshops (ICDEW)</i> , pages 28–32. IEEE.
951	Wentao Zhang, Lingxuan Zhao, Haochong Xia, Shuo
952	Sun, Jiaze Sun, Molei Qin, Xinyi Li, Yuqing Zhao,
953	Yilei Zhao, Xinyu Cai, and 1 others. 2024a. Finagent:
954	A multimodal foundation agent for financial trading:
955	Tool-augmented, diversified, and generalist. <i>arXiv</i>
956	<i>preprint arXiv:2402.18485</i> .
957	Yi-Fan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou
958	Fu, Shuangqing Zhang, Junfei Wu, Feng Li, Kun
959	Wang, Qingsong Wen, Zhang Zhang, and 1 oth-
960	ers. 2024b. Mme-realworld: Could your multi-
961	modal llm challenge high-resolution real-world sce-
962	narios that are difficult for humans? <i>arXiv preprint</i>
963	<i>arXiv:2408.13257</i> .
964	Zhihan Zhang, Yixin Cao, and Lizi Liao. 2025b. Xfin-
965	bench: Benchmarking llms in complex financial prob-
966	lem solving and reasoning. In <i>Findings of the Asso-</i>
967	<i>ciation for Computational Linguistics: ACL 2025</i> ,
968	pages 8715–8758.
969	Yilun Zhao, Hongjun Liu, Yitao Long, Rui Zhang,
970	Chen Zhao, and Arman Cohan. 2024. Financemath:
971	Knowledge-intensive math reasoning in finance do-
972	main. In <i>Proceedings of the 62nd Annual Meeting of</i>
973	<i>the Association for Computational Linguistics (Vol-</i>
974	<i>ume 1: Long Papers)</i> , pages 12841–12858.
975	Jie Zhu, Junhui Li, Yalong Wen, and Lifan Guo. 2024.
976	Benchmarking large language models on cflue—a
977	chinese financial language understanding evaluation
978	dataset. <i>arXiv preprint arXiv:2405.10542</i> .

<b>A</b>	<b>Details of UniFinEval</b>	979
<b>A.1</b>	<b>Comparison of benchmarks</b>	980
	In this section, we first provide a detailed com-	981
	parison Table 2 of various Q&A datasets across	982
	multiple dimensions.	983
<b>A.2</b>	<b>Statistic and Examples of Financial</b>	984
	<b>Business Scenarios</b>	985
	We list the detailed information of UniFinEval data	986
	in Table 3. The detailed information of the financial	987
	business scenarios are presented below.	988
	<b>Financial Statement Auditing:</b> As the founda-	989
	tional entry point of financial operations, its core	990
	objective is to accurately localize key financial in-	991
	dicators and event-related information from high-	992
	density materials such as financial statements and	993
	research reports. This establishes a solid data foun-	994
	dation for all subsequent analyses and decision-	995
	making, ensuring the accuracy of information at its	996
	source. Unlike data using only simplified charts,	997
	this scenario retains real typesetting and redundant	998
	information, forcing the model to complete key in-	999
	formation screening under real visual interference.	1000
	Question design transitions from single factn to	1001
	cross-page, multi-point multi-hop reasoning, sim-	1002
	ulating the actual work process of analysts inte-	1003
	grating information across multiple pages. This	1004
	assesses whether the model possesses the ability	1005
	to effectively screen and precisely localize infor-	1006
	mation within real financial materials. Figure 6	1007
	provides examples of Financial Statement Audit-	1008
	ing.	1009
	<b>Company Fundamental Reasoning:</b> Building	1010
	upon basic information extraction, this scenario cor-	1011
	responds to the information alignment and synchron-	1012
	ization stage and focuses on interpreting semantic	1013
	consistencies and information overlaps embedded	1014
	in textual and graphical content. By rapidly cap-	1015
	turing redundancies or correlations across sources,	1016
	it provides critical alignment-level signals for in-	1017
	vestment decisions, bridging the gap between het-	1018
	erogeneous data formats and unified business logic.	1019
	Data comes from research reports, market com-	1020
	mentaries, and analysis materials, where charts no	1021
	longer directly present explicit financial indicators	1022
	but convey consistent information through trends,	1023
	distributions, or relative changes; texts also contain	1024
	descriptive or summary-based representations. Fig-	1025
	ure 7 and Figure 8 are the examples of Company	1026
	Fundamental Reasoning.	1027
	<b>Industry Trend Insights:</b> Designed for multi-	1028

source data integration scenarios, this task addresses the complexity of deriving core metrics from diverse raw data points found across modalities such as text and charts. By executing precise quantitative formulas across modalities, it prevents computational errors caused by fragmented information and ensures the accuracy and reliability of calculated financial results. The data includes explicit or implicit links between text and charts, deliberately retaining the incomplete symmetry of expression between modalities, making it impossible for the model to complete the task through surface matching; it must understand the underlying calculation logic behind different modalities. Questions also deepen from single-step parameter extraction to multi-stage formula derivation and sensitivity analysis. Figure 9 are the examples of Industry Trend Insights.

**Financial Risk Sensing:** Corresponding to long-term industry tracking and mid-term analytical assessment, this scenario integrates temporal information from both static text–image inputs and dynamic video data to infer industry development logic, core driving factors, and volatility risks, providing trend-level support for mid-term investment direction selection. Question design requires the model to synthesize information changes at different moments to explain and judge industry development trends or market fluctuations. Figure 10 are the examples of Financial Risk Sensing.

**Asset Allocation Analysis:** As the ultimate decision-making stage in financial operations, this scenario integrates outputs from all preceding stages and, under multiple constraints such as policy and risk considerations, produces actionable asset allocation strategies. It directly reflects the practical deployment value of the model and supports core investment decision-making. Question design requires the model to continuously integrate new information during interaction and maintain decision logic consistency under multimodal, multi-hop reasoning conditions. Figure 12 are the examples of Asset Allocation Analysis.

### A.3 Details of Quality Control

**Data Collection** To ensure authority and real-world alignment, UniFinEval sourced data from global real financial markets, collecting over 2,500 listed company documents and financial research reports, as well as 1,000 financial analysis videos. These materials cover both Chinese and English environments across various asset classes, with all

sources verified for copyright compliance. To maintain high information density, we employed a multi-layered filtering process combining automated scripts with collaborative manual verification to remove low-quality or irrelevant content.

**Data Filtering** Experts first used automated scripts to crop pages with financial charts and associated text from PDF reports, forming independent complete images with high information density. For financial analysis videos, scene segmentation, keyframe extraction, and timestamp alignment were performed to ensure semantic alignment between static screenshots and the original video. Subsequently, manual review and screening were conducted, focusing on assessing business representativeness, information effectiveness, and scenario adaptability to fundamentally guarantee data quality and business value.

**Question Construction** First, each expert independently completed question design and standard answer annotation based on assigned scenario data to ensure originality and independence, avoiding homogenization of thought. Subsequently, a dual-round cross-validation phase began. Each question was verified by an expert annotator from three dimensions: answer accuracy, logical rigor, and semantic clarity. If explicit errors were found in the question or answer, the question was directly discarded; if doubts existed regarding logic or semantics, it was submitted to a second expert annotator for independent re-verification. If a consensus could not be reached or both annotators were unsure, the data was removed from the benchmark.

**Expert Review** First, all questions had to pass a compliance review, where experts with experience in financial regulatory policy research assessed whether they met financial industry compliance requirements, avoiding sensitive expressions or improper business scenario settings. Second, the expert team conducted a centralized review of business adaptability and difficulty gradients, focusing on checking if questions fit real business flows while ensuring a reasonable difficulty ratio to distinguish model capabilities at different levels in a fine-grained manner. For multi-turn Q&A tasks, an additional logical coherence review was required, where experts simulated the multi-turn interaction process from the perspective of a real analyst to verify if the dialogue link was smooth, the context consistent, and the decision deduction compliant with business logic.

1131  
1132  
1133  
1134  
1135  
1136  
1137  
1138  
1139  
1140  
1141  
1142  
1143  
1144  
1145  
1146  
1147  
1148  
1149  
1150  
1151  
1152  
1153  
1154  
1155  
1156  
1157  
1158  
1159  
1160  
1161  
1162  
1163  
1164  
1165  
1166  
1167  
1168  
1169  
1170  
1171  
1172  
1173  
1174  
1175  
1176  
1177  
1178  
1179  
1180  
1181

#### A.4 Examples of Environmental Perturbation Simulations

In real-world financial office and auditing scenarios, paper-based documents often suffer from degraded visual quality due to improper storage, suboptimal photographing angles, or physical wear and tear. To more comprehensively evaluate model robustness under extreme real-world conditions, UniFinEval specifically introduces three environment-noise processing tasks targeting physical entities. These simulations are designed to mimic common degradation phenomena encountered during the digitization of financial receipts and reports. The categories and definitions are as follows:

**Stain Simulation:** This task simulates ink smudges, coffee stains, or mold spots that may occur during the circulation or storage of financial documents. Such disturbances are implemented by randomly overlaying masks of varying shapes and transparency on the images, partially occluding text or accounting relationships. This setting tests the model’s ability to recover incomplete information and its robustness to visual interference. Figure 13 provides an example of this case.

**Perspective Simulation:** This task simulates geometric distortions caused by photographing documents with mobile phones or scanners from non-parallel angles. By applying perspective projection transformations to the original images, documents exhibit trapezoidal warping or edge stretching. This requires the model to possess strong spatial perception capabilities in order to accurately extract structured data under coordinate shifts. Figure 14 demonstrates this perturbation type.

**Crease Simulation:** This task simulates physical creases left on paper documents after being folded or crumpled. Creases introduce localized linear highlights or shadows and may cause text strokes to break or become misaligned. Such simulations effectively assess the OCR accuracy of multimodal models when dealing with non-flat document layouts.

**Curvature Simulation:** Simulates the non-linear curling effects presented by thick financial reports, book binding areas, or unflattened paper. Unlike perspective distortion, this type of interference causes text lines to distort irregularly following the curvature of the paper, focusing on testing the model’s capabilities in layout analysis, dewarping, and character rectification under complex 3D deformations.

**Noise Simulation:** Simulates image background noise resulting from aging scanner sensors, transmission compression loss, or low-light shooting environments. By injecting Gaussian noise or salt-and-pepper noise into the original image to reduce the Signal-to-Noise Ratio (SNR), it aims to evaluate the model’s ability to accurately recognize fine details of numbers and text under low-quality imaging conditions characterized by blurriness and strong graininess.

## B Details of MLLMs

We list details of the MLLMs evaluated using UniFinEval in Table 4.

### B.1 Examples for Error Analysis

Building upon the error categories introduced in the main text, this appendix provides a qualitative explanation of each error type. The purpose is to clarify the characteristics and underlying causes of common failure modes exhibited by MLLMs in financial reasoning tasks, rather than to introduce additional experimental analyses.

**Financial Image Perception and Data Interpretation Errors:** This category of errors refers to failures occurring at the visual perception and data interpretation stage of financial multimodal reasoning. In finance-related tasks, models are often required to accurately extract numerical values, trends, and structural relationships from visual inputs such as financial charts, tables, trend plots, and annotated figures. Errors arise when the model misidentifies key visual elements (e.g., axes, legends, data points) or incorrectly interprets visually encoded information. These issues are particularly prominent under conditions of high visual complexity, noise, dense annotations, low image resolution, or overlapping graphical components, and they frequently propagate to subsequent reasoning steps, resulting in incorrect financial conclusions.

**Financial Knowledge Reasoning and Domain-Specific Understanding Errors:** This type of error stems from limitations in the model’s understanding of financial concepts, professional terminology, and domain-specific reasoning principles. Financial tasks are characterized by high knowledge density and strong reliance on specialized concepts, such as financial ratios, capital structure, accounting rules, and risk assessment frameworks. Models may misinterpret these concepts, confuse related financial terms, or apply inappropriate domain logic,

1182  
1183  
1184  
1185  
1186  
1187  
1188  
1189  
1190  
1191  
1192  
1193  
1194  
1195  
1196  
1197  
1198  
1199  
1200  
1201  
1202  
1203  
1204  
1205  
1206  
1207  
1208  
1209  
1210  
1211  
1212  
1213  
1214  
1215  
1216  
1217  
1218  
1219  
1220  
1221  
1222  
1223  
1224  
1225  
1226  
1227  
1228  
1229  
1230

1231 even when the input information is correctly per- 1283  
1232 ceived. As a result, the generated reasoning may de- 1284  
1233 viate from established financial analysis practices,  
1234 leading to conclusions that are logically coherent  
1235 in form but flawed from a professional finance per-  
1236 spective.

1237 **Financial Computation and Numerical Anal-**  
1238 **ysis Errors:**Financial computation and numerical  
1239 analysis errors occur when models fail to produce  
1240 accurate quantitative results during multi-step nu-  
1241 merical reasoning. In finance, precise calculations  
1242 are essential for tasks involving valuation, prof-  
1243 itability analysis, investment returns, or compara-  
1244 tive financial metrics. Models may make arithmetic  
1245 mistakes, apply incorrect calculation sequences,  
1246 suffer from rounding or precision loss, or fail to  
1247 consistently track intermediate numerical states.  
1248 Such errors indicate limitations in the reliability  
1249 of numerical reasoning, where even minor com-  
1250 putational inaccuracies can lead to substantially  
1251 distorted financial interpretations.

1252 **Cross-modal Data Integration and Alignment**  
1253 **Errors:**Cross-modal data integration and align-  
1254 ment errors arise when models fail to correctly as-  
1255 sociate and fuse information across different input  
1256 modalities, such as text, images, and tables. Finan-  
1257 cial multimodal tasks often require joint reasoning  
1258 over heterogeneous sources, for example aligning  
1259 textual descriptions with chart-based evidence or  
1260 tabular financial disclosures. Errors occur when the  
1261 model overlooks critical cues, mismatches refer-  
1262 ences between modalities, or incorrectly prioritizes  
1263 information from one modality while neglecting  
1264 others. These misalignments can result in incom-  
1265 plete or inconsistent interpretations of the overall  
1266 financial context, ultimately undermining the accu-  
1267 racy of the reasoning process.

1268 **Inconsistent Financial Reasoning and Hallu-**  
1269 **cination Errors:**This category captures errors re-  
1270 lated to logical inconsistency and hallucinated con-  
1271 tent in financial reasoning. Such errors occur when  
1272 models generate conclusions that are not supported  
1273 by the provided evidence, contradict earlier rea-  
1274 soning steps, or rely on fabricated assumptions.  
1275 In financial scenarios, this may manifest as un-  
1276 founded predictions of market trends, speculative  
1277 assessments of corporate financial health, or con-  
1278 fident but unsupported explanations of financial  
1279 outcomes. These errors are particularly problem-  
1280 atic in decision-sensitive financial applications, as  
1281 they may present misleadingly plausible narratives  
1282 that lack factual or logical grounding.

Representative erroneous examples for all five  
error categories are illustrated in Figures 15–19.

## C Prompts Used in This Study

We provide representative prompt examples for  
evaluation. Specifically, the prompt examples for  
evaluation are shown in Table 5, Table 6.

1283  
1284  
1285  
1286  
1287  
1288



文本：工业生产稳中向好，计算机通信动能较强腾景全口径数据显示，工业经济运行相对平稳，1-10月工业增加值不变价累计同比增速为3.8%，其中4月份受到疫情冲击影响，增速落入负向区间(-3.9%)，而后触底反弹、保持小幅上行态势。其中，房地产投资、出口交货值与工业增加值同比增速走势契合度较高，但在2021年5月出现K形分化走势，今年K形分化程度进一步加深，具体表现在房地产投资增速自去年7月份落入负向区间后，不断向下走阔，持续两位数的负增长低迷状态，阻碍工业修复。

**Text:** Industrial production remains stable with a positive outlook, driven by strong momentum in the computer and communication sectors. According to Tengjing's comprehensive data, the industrial economy is operating relatively steadily. From January to October, the cumulative year-on-year growth rate of industrial value-added (at constant prices) was 3.8%. Notably, growth fell into negative territory (-3.9%) in April due to the impact of the pandemic, but subsequently bottomed out and rebounded, maintaining a slight upward trajectory. Historically, the year-on-year growth trends of real estate investment and export delivery value have closely aligned with industrial value-added. However, a "K-shaped" divergence emerged in May 2021 and has intensified this year. This is specifically manifested in real estate investment growth: since falling into negative territory last July, the decline has widened, persisting in a slump of double-digit negative growth that is hindering industrial recovery.

**Round 1:**

问题：依据文本中关于房地产投资增速的描述，其进入持续两位数负增长区间的起始月份是哪个月？

**Question:** Based on the description of the real estate investment growth rate in the text, what is the starting month in which it entered the range of sustained double-digit negative growth? **答案:** 2021年7月 **Answer:** July 2021

**Round 2:**

问题：结合确定的起始月份和图5中显示的最新版数据点，计算房地产投资增速处于两位数负增长状态的总月份数。

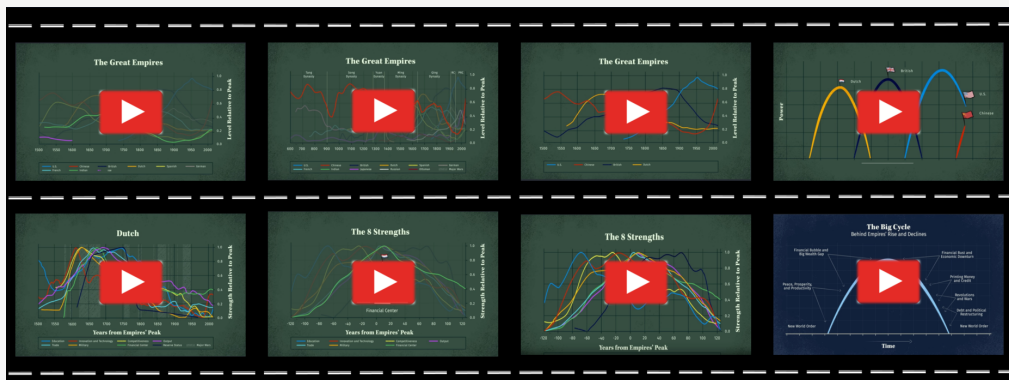
**Question:** Combining the identified starting month with the latest data point shown in Figure 5, calculate the total number of months during which the real estate investment growth rate remained in a state of double-digit negative growth.

**答案:** 16

**Answer:** 16

Figure 6: This example showcases the Financial Statement Auditing scenario, focusing on high-precision data retrieval and alignment. To provide an accurate answer, the model must demonstrate the ability to precisely localize specific indicators across unstructured text and complex time-series charts. Specifically, the model is required to pinpoint the exact starting month (July 2021) of sustained double-digit negative growth in real estate investment within the text, and then accurately map its location to the corresponding trend in Figure 5. This task evaluates the model's ability to maintain high-precision and consistency verification when tracking specific industry fluctuations within high-density and specialized financial materials.





Text: As I described in my book *Principles for Dealing with the Changing World Order*, I automated my way looking at the cause effect relationships that are driving both improvements in and worsenings of countries' conditions so that data is fed into a computer that analyzes it and writes a summary of the current conditions and the long-term prospects for each country. ... The table below shows our aggregate country power gauge and the major drivers, as well as the rank of each measure of power across 24 major countries today and the trajectory over the past twenty years. To understand a country, we start by looking at the big cycles, as well as measures of power that both reflect and drive the rise and fall of a country. While we refer to these factors individually, they are not separate; they interact with and reinforce one another to move a country along its cycle. For the United States, the big cycles look mostly unfavorable. The United States is in an unfavorable position in its economic and financial cycles, with a high debt burden and relatively low expected real growth over the next 10 years (1.3% per year). The United States has significantly more foreign debts than foreign assets (net IIP is -68% of GDP). Non-financial debt levels are high (274% of GDP), and government debt levels are high (127% of GDP). The bulk (99%) of these debts are in its own currency, which mitigates its debt risks. The ability to use interest rate cuts to stimulate the economy is low (short rates at 0.1%), and the country is already printing money to monetize debt. That said, being the world's leading reserve currency is a large benefit to the US. If this were to change, it would significantly weaken the US position. Internal disorder is a high risk. Wealth, income, and values gaps are large (relative to countries of similar per capita income levels). Regarding Inequality—the top 1% and top 10% in the United States capture 19% and 45% of income (respectively the 8th and 11th highest share across major countries). ... Half of the measure captures the absolute quantity of educated people at various levels and about half is placed on quality such as higher education rankings, test scores, and average years of education. The US ranks highest in this gauge (driven by strong absolute and relative measures of higher education), with China close behind (due to its large number of educated people). Financial Center: This gauge measures the level of development and sizes of a country's financial markets and financial center. We look at absolute measures of transaction shares and market capitalizations, as well as external indices of financial center cities. The US remains the top-ranked power in this metric by a significant margin (driven primarily by its very large share of world equity and debt markets), with China and Europe ranking second and third, respectively. Reserve Currency Status: This gauge measures the extent to which a country's currency operates as a global reserve currency. We measure reserve currency status by the share of transactions, debts, and central bank reserves that are denominated or held in a country's currency. Similar to financial center status, the US remains the top-ranked power in this metric by a significant margin, with Europe and Japan ranking second and third, respectively. ... In case it is helpful or interesting to you, you can review those scores below.<sup>33</sup> In a few cases where there were no quality measures, I had to create quality measures by adjusting the quantity for a country's population, turning it into a per capita measure. We did not give reserve currency status scores to the countries that share the euro, which is why those measures are displayed as dashes. <sup>4</sup> Because the notion of competitiveness is inherently relative, we only show the total score for this measure.

**Question:** Synthesizing the analytical logic regarding the 'Big Cycle' and the 'Eight Major Strengths' in the text, as well as the average evolutionary trends of the 'Eight Major Strengths' for all empires before and after their respective peaks as shown in Figure 5, please infer: After an empire's power reaches its peak, which 'strength' indicator declines the fastest (i.e., requires the shortest time to drop from its peak to the 0.2 level)? Please substantiate your argument by combining the trend lines in the video with the implications in the text regarding the vulnerability of 'financial center' status.

**Answer:** Financial Center

Figure 9: This example demonstrates the Industry Trend Insights scenario. It requires the model to extract the evolutionary trends of the 'Eight Major Strengths' indicators for all empires before and after their peaks from the video, focusing on comparing the declining slopes of each indicator curve after the peak, to preliminarily identify that the 'financial center' indicator curve has the steepest declining slope, presenting a trend of rapid decline; secondly, to locate descriptions regarding the 'financial center' indicator in the text, confirming that its core components rely on short-term capital flows and market confidence, and that its strong binding relationship with reserve currency status will trigger a transmission chain of rapid decline when the economic and financial cycle turns unfavorable; and finally, to associate the steep decline trend of the 'financial center' in the video with the high sensitivity and high vulnerability of this indicator in the text, arriving at the conclusion that the 'financial center' indicator declines the fastest after the empire's power reaches its peak.

**WTI 原油反弹无力**

WTI 原油上周未能延续反弹，在周三开始连续下跌，于上周五收于 70.04 美元/桶，周平均涨幅 1.27%。截至 2023 年 5 月 5 日当周，美国原油产量为 1230 万桶/日，同比增加 40 万桶/日。截至 5 月 12 日当周，美国钻机数 731 台，较去年同期增加 17 台；其中采油钻机数 586，较去年同期增加 23。原油增产仍动能不足，工作钻机数出现回落趋势。5 月 5 日当周，美国炼厂吞吐量为 1574.5 万桶/日，开工率 91.0%。

OPEC+ 4 月份的原油产量较 3 月份减少了 38 万桶/日，原因是伊拉克无法出口其北部的产量，而尼日利亚则遭遇了重大停电。欧佩克 13 个国家的产量为 2860 万桶/日，下降了 37 万桶/日。与此同时，以俄罗斯为首的非欧佩克盟国的产量下降了 1 万桶/日，至 1339 万桶/日。受西方制裁影响，俄罗斯的产量在本月稳定在 960 万桶/日，尽管政府声称其已在 4 月份自愿从 2 月份的水平减产 50 万桶/日。从 5 月开始，在其他几个国家实施了额外的自愿减产，OPEC+ 的产量可能会进一步下降。沙特阿拉伯、俄罗斯、伊拉克、阿联酋和其他一些国家计划从 5 月到 2023 年底将减产总量增加约 166 万桶/日，理由是经济预测不佳。自去年 11 月以来，所有 OPEC+ 国家共同减产 200 万桶/日。OPEC+ 4 月份的产量较低意味着其产量与配额之差继续扩大。拥有配额的 19 个成员总计比目标低 258.3 万桶/日——这一缺口在过去四个月中每月都在增加。OPEC+ 部长级会议将于 6 月 3 日至 4 日举行，以审查生产政策。

EIA 5 月 9 日表示，市场对全球经济疲软、通胀和银行业持续动荡的情绪超过了对 OPEC+ 减产的担忧，使得石油、汽油和柴油价格前景走低。EIA 在其 5 月短期能源展望中将 2023 年布伦特原油的预测下调 6.36 美元至 78.65 美元/桶，远低于 4 月份 85 美元/桶的平均价格。同样，该机构预测 WTI 原油在 2023 年的平均价格为 73.62 美元/桶，比上个月的全年预估低 5.62 美元。虽然产油国集团的石油产量预计将在 2023 年减少，但非 OPEC+ 产油国的产量预计将推动今年全球液体燃料产量增长 150 万桶/日。与此同时，预计 2023 年以中国和印度为首的全球液体燃料需求将增长 160 万桶/日。EIA 将 2023 年全球石油需求前景上调 12 万桶/日，至 1.0099 亿桶/日，这种需求增长将使全球石油市场在 2023 年第三季度至 2024 年第一季度之间达到平衡，将布伦特原油价格从当前水平推回 75-80 美元/桶。

海关总署 5 月 9 日发布的数据显示，我国 4 月份原油进口量下降 16% 至 1036 万桶/日，低于 3 月份 1237 万桶/日的 33 个月高位，而成品油出口 4 月环比大幅下降 31% 至九个月低点 375 万吨。然而油轮追踪机构 Kpler 估计，我国 5 月份的海运原油进口量预计将达到 1140 万桶/日，高于 3 月份的 1108 万桶/日，其中当月美国原油流入量可能会增加 28.4 万桶/日。

科威特 Al-Zour 炼油厂 1 号和 2 号原油蒸馏装置已恢复全面运行，截至 4 月 30 日，产量从 20.5 桶/天恢复到 34.5 万桶/日。CDU3 预计将在 2023 年底前准备就绪，使产量达到 61.5 万桶/日。Al-Zour 是中东正在建设或启动的几家新炼油厂之一。沙特阿美的 Jazan 炼油厂于 2021 年投产，产能为 40 万桶/天，阿曼的 Duqm 炼油厂将于今年投产，产能为 23 万桶/日，伊拉克的 Kerbala 炼油厂也预计将于 2023 年投产，产能为 14 万桶/日。

问题：结合 OPEC+ 4 月份实际减产数据、5 月到年底的计划减产总量，以及 EIA 对全球经济情绪的评估，分析 EIA 下调 2023 年 WTI 原油平均价格预测的核心原因，并指出该原因如何导致 OPEC+ 减产对价格的支撑作用被削弱。

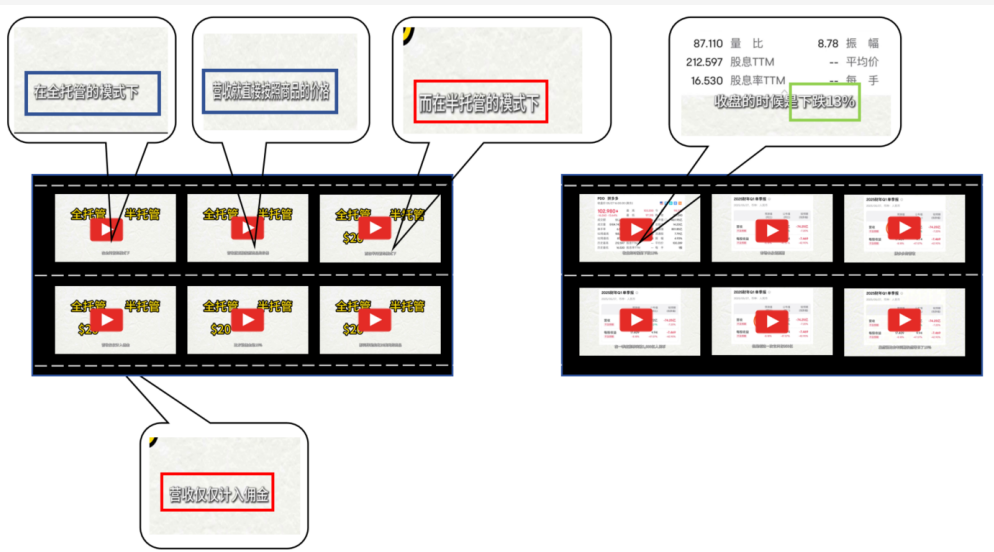
Question: By combining OPEC+'s actual production cuts implemented in April, the total planned production cuts from May through the end of the year, and the EIA's assessment of global economic sentiment, analyze the core reason why the EIA lowered its forecast for the average WTI crude oil price in 2023, and explain how this reason weakened the price-supporting effect of OPEC+ production cuts.

答案：市场对全球经济疲软、通胀和银行业动荡的情绪超过了对 OPEC+ 减产的担忧；经济担忧抵消了减产效果，使价格预期下调

Answer: Market concerns about global economic weakness, inflation, and turmoil in the banking sector outweighed concerns over OPEC+ production cuts; these economic worries offset the impact of the cuts, leading to a downward revision in price expectations.

Figure 10: This example illustrates the Financial Risk Sensing scenario, requiring the model to keenly capture potential downside risk signals. First, it must cross-verify the structural details of OPEC+'s complex production plan with the descriptive macro sentiments highlighted by the EIA, such as global economic weakness, inflation, and banking turmoil. Subsequently, the model must accurately synchronize these heterogeneous signals to identify that market concerns regarding an economic recession have outweighed expectations of support from production cuts. This task evaluates whether the model can logically align these disparate information threads to justify the downward revision of the WTI average price forecast to \$73.62 per barrel.

## Financial Risk Sensing



问题：请问拼多多2025年Q1营业收入波动的主要原因是什么？

Question: What are the primary reasons for Pinduoduo's revenue fluctuations in Q1 2025?

答案：拼多多商业模式由全托管转为半托管。

Answer: Pinduoduo's (Temu) Business Model Transition from Fully Managed to Semi-Managed.

Figure 11: This example illustrates the Financial Risk Sensing scenario. This task requires the model to reason across two videos to identify the primary factors driving revenue fluctuations and to further recognize the underlying potential risks.

## Asset Allocation Analysis

股票	债券
<p>沪深 300 -1.10%</p> <p>沪深 300 股指期货 -0.81%</p> <p><b>本周点评:</b> 本周A股市场主要指数震荡回调, 沪指周跌1.22%, 回落于3200点下方。深成指以及创业板指回调接近2%。市场成交活跃度再度降温, 两市成交量不足万亿。四期中, 大盘指数支撑性相对较强, 上证50周度回调0.68%。</p> <p><b>配置建议:</b> 低配区间操作</p>	<p>10 年国债到期收益率-0.50%/本周变动-0.15BP</p> <p>主力 10 年期国债期货 +0.30%</p> <p><b>本周点评:</b> 本周期债主力合约较上周小幅上行, DR001小幅上行至1.22%。资金利率整体仍在宽松范围内, 债市心态保持平稳。预计货币市场利率在较低位置运行的状态有望保持至12月税期高峰前夕。</p> <p><b>配置建议:</b> 低配区间操作</p>
大宗商品	外汇
<p>主力沪铜期货 -0.69%</p> <p>主力原油期货 +2.27%</p> <p><b>本周点评:</b> 近期铜价逐渐走高, 国内下游恐高情绪升温, 现货升水回落, 且进口货源流入, 使得库存重回增长; 而LME注销仓单开始下降, 升贴水逐渐走低, 预计铜价上方压力增大。</p> <p><b>配置建议:</b> 低配区间操作</p>	<p>欧元兑美元 +0.52%</p> <p>欧元兑美元2212合约 -0.03%</p> <p><b>本周点评:</b> 美联储基调偏鹰, 利率峰值及高利率周期预期皆有所上调。欧洲央行基调转鹰, 但高通胀问题尚未得到改善且经济衰退严重, 欧元或在短期内承压回落。</p> <p><b>配置建议:</b> 逢高做空</p>

问题: 基于文字说明中'随后两个交易日内保持日均1000亿元左右的投放规模'的表述, 以及图表25所反映的2022年11月市场波动背景, 计算从操作当日到后续两个交易日的总投放规模相对于当日净投放额的倍数(保留两位小数)

Question: Based on the statement in the text description regarding 'maintaining an average daily injection scale of around 100 billion yuan for the subsequent two trading days,' and the background of market volatility in November 2022 reflected in Figure 25, calculate the multiple of the total injection volume from the day of operation through the subsequent two trading days relative to the net injection amount on that day (rounded to two decimal places).

答案: 1.39

Answer: 1.39

Figure 12: This is a comprehensive question examining macro asset allocation logic and cross-market volatility attribution capabilities. It requires the model to possess keen financial semantic recognition and multi-dimensional indicator alignment capabilities: First, it must extract fluctuation data across categories for the CSI 300 Index, stock index futures, and dominant crude oil futures; second, it needs to accurately calculate the average risk exposure of the equity and energy dual axes in major asset allocation; finally, quantitatively derive the core metric reflecting the true market volatility level, testing the model's quantitative analysis chain from discrete data capturing to composite risk measurement.

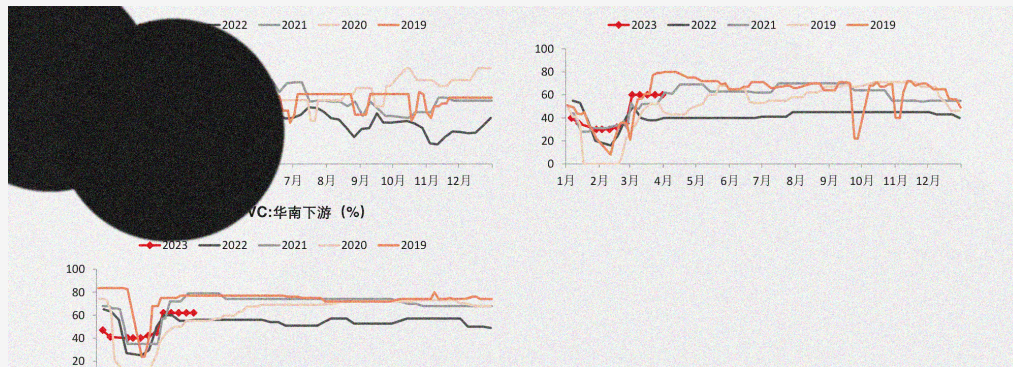
Table 2: Comparison of various benchmarks across multiple dimensions. The abbreviations in the header are: MLD (Multi-level Difficulty), RES (Realistic Environment Simulation), MC (Manually Construct), FAV (Financial Analysis Video), CD (Consistency Detection), and MHR (Multi-Hop Reasoning).

<b>Benchmarks</b>	<b>MLD</b>	<b>RES</b>	<b>MC</b>	<b>FAV</b>	<b>CD</b>	<b>MHR</b>
<b>Image</b>						
FinMR	✓	✗	✗	✗	✗	✗
FinMMR	✓	✗	✗	✗	✗	✗
FinMME	✓	✗	✗	✗	✗	✗
Multifinben	✓	✗	✗	✗	✗	✗
XFinBench	✓	✗	✗	✗	✗	✗
VisFinEval	✓	✓	✗	✗	✗	✗
FinMultiTime	✓	✗	✗	✗	✗	✗
CFBenchmark-MM	✗	✗	✗	✗	✗	✗
<b>Video</b>						
FinCap	✓	✗	✗	✓	✗	✗
VideoConviction	✓	✗	✗	✓	✗	✗
<b>Ours(Text + Image + Video)</b>	✓	✓	✓	✓	✓	✓

Table 3: Financial Scenario Data Distribution. The table details the distribution of questions across five core financial tasks: Financial Statement Auditing (FSA), Company Fundamental Reasoning (CFR), Industry Trend Insights (ITI), Financial Risk Sensing (FRS), and Asset Allocation Analysis (AAA).

<b>Financial Scenario</b>	<b>Questions</b>
Financial Statement Auditing	892
Company Fundamental Reasoning	926
Industry Trend Insights	896
Financial Risk Sensing	535
Asset Allocation Analysis	518
<b>Total (UniFinEval)</b>	<b>3767</b>

## Stain Simulation



问题：基于提供的金融研报图像（包含左上角、右上角和底部三个子图，均展示PVC华南下游开工率（%），请分析：在右上角子图中，2019年开工率在2月的下降趋势与2023年在3月的上升趋势相比，哪一段的波动更剧烈？

Question: Based on the provided financial research report image (which contains three subplots located at the top-left, top-right, and bottom, all showing the PVC South China downstream operating rate (%), please analyze the following: In the top-right subplot, when comparing the downward trend in the operating rate in February 2019 with the upward trend in March 2023, which period exhibits more pronounced volatility?

答案：2019年2月的下降段更剧烈

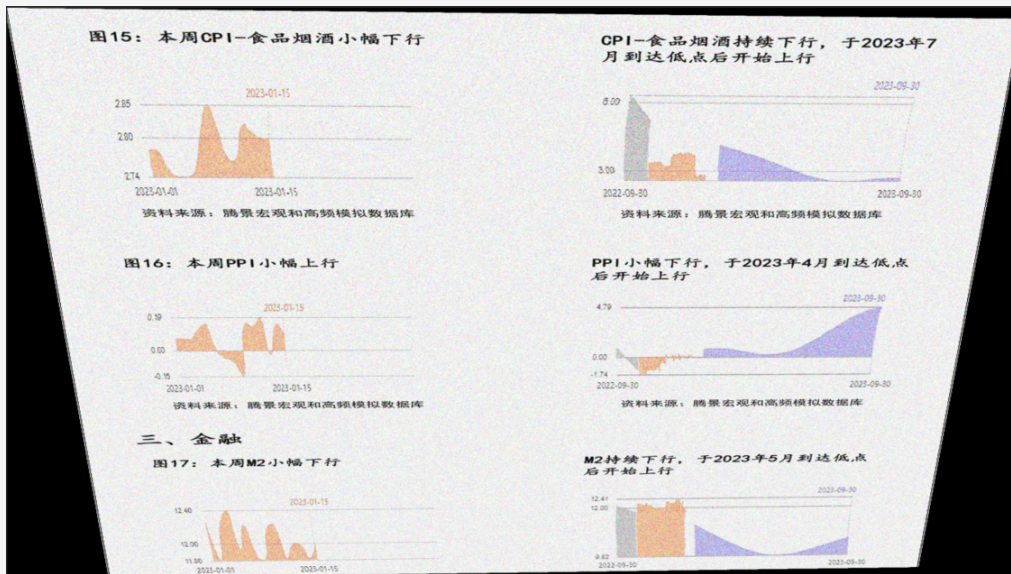
Answer: The decline in February 2019 was more pronounced.

Figure 13: This is an example of stain simulation perturbation, simulating ink smudges, coffee stains, or mold spots acquired by financial documents during circulation or storage.

Table 4: Models evaluated in this paper. The "Access" column shows whether we have full access to the model weights or we can only access through API. The "Version Date" column shows the release date of the corresponding version of the model we evaluated.

Category	Model	Creator	Parameter	Access	Version Date
Close-Source	Gemini-3-pro-preview	Google	Undisclosed	API	2025.11
	GPT-5.1-2025-11-13	OpenAI	Undisclosed	API	2025.11
	Grok-4.1-Fast-reasoning	xAI	Undisclosed	API	2025.11
	Claude-Sonnet-4.5-20250929	Anthropic	Undisclosed	API	2025.9
Open-Source	Qwen3-VL-235B-A22B-thinking	Alibaba Cloud	235B	Weights	2025.11
	Qwen3-VL-32B-thinking	Alibaba Cloud	32B	Weights	2025.11
	InternVL3.5-241B-A28B	Shanghai AI Lab	241B	Weights	2025.8
	InternVL3.5-30B-A3B	Shanghai AI Lab	30B	Weights	2025.8
	Llama-3.2-11B-Vision	Meta AI	11B	Weights	2024.9
	MiniCPM-V-4.5	OpenBMB	9B	Weights	2025.9

Perspective Simulation



问题：综合图15、16、17右侧子图对PPI、M2、CPI-食品烟酒三个指标触底反弹时间的描述（PPI于2023年4月、M2于2023年5月、CPI-食品烟酒于2023年7月到达低点后开始上行），以及左侧子图显示的2023年1月15日当周三个指标均小幅下行的特征，判断在2023年1月15日时点，哪个指标的下行期剩余时间最短？需严格依据图像文字信息进行跨指标时序比较和趋势阶段定位。

Question: Based on a comprehensive analysis of the descriptions in the right-hand subplots of Figures 15, 16, and 17 regarding the bottoming-out and rebound timing of three indicators—PPI, M2, and CPI-Food, Tobacco, and Liquor (specifically, PPI reached its low point and began to rise in April 2023, M2 in May 2023, and CPI-Food, Tobacco, and Liquor in July 2023)—combined with the characteristic shown in the left-hand subplots where all three indicators experienced a slight decline during the week of January 15, 2023, determine which indicator had the shortest remaining time in its downward phase as of the time point of January 15, 2023. You are required to strictly base your analysis on the text information within the images to perform a cross-indicator temporal comparison and trend phase positioning.

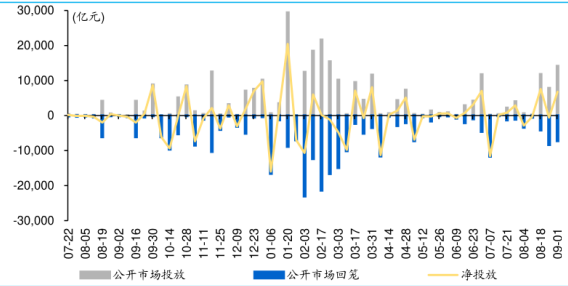
答案: PPI  
Answer: PPI

Figure 14: This is an example of perspective simulation perturbation, simulating the geometric distortion caused by capturing images with mobile phones or scanners at non-parallel angles.

流动性常规跟踪

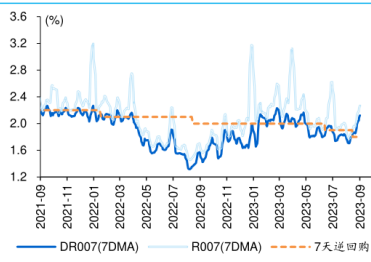
8月26日至9月1日期间，央行公开市场净投放资金6810亿元，其中，7天逆回购投放14090亿元、利率持平于1.9%，7天逆回购到期7280亿元。本周（9月2日至9月8日），将有14090亿元7天逆回购到期。

图表 23: 7.8-7.14, 央行公开市场净回笼资金 350 亿元

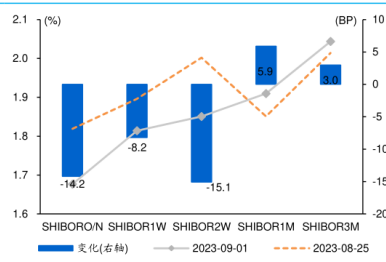


货币市场利率整体回落。8月26日至9月1日期间，R001和R007分别较上周回落15.9BP和10.5BP至1.76%和2.01%，DR007较上周上涨14.96BP至1.80%、低于同期政策利率的1.9%。SHIBOR隔夜和1周利率分别为较上周回落14.2BP和8.2BP至1.68%和1.81%，3个月较上周上涨3BP至2.04%。

图表 24: 货币市场利率整体上涨



图表 25: SHIBOR 利率涨跌互现



问题：结合文字说明中SHIBOR 3个月利率的周度变化数据与图表25中SHIBOR 1个月利率的周度变化数据，计算3个月与1个月SHIBOR利差在8月26日至9月1日期间的基点变化值（需考虑利差定义方向）。

Question: Based on the weekly change data of the 3-month SHIBOR rate provided in the textual description and the weekly change data of the 1-month SHIBOR rate shown in Figure 25, calculate the basis-point change in the SHIBOR spread between the 3-month and 1-month tenors during the period from August 26 to September 1 (taking into account the definition and direction of the spread).

答案: -2.9

Answer: -2.9

模型输出:3个月与1个月SHIBOR利差在8月26日至9月1日期间的基点变化值为7.0 BP。

Model Output:The change in the basis-point spread between the 3-month and 1-month SHIBOR over the period from August 26 to September 1 is 7.0 bps.

Figure 15: This case represents a typical example of Financial Image Perception and Data Interpretation Errors, where the model incorrectly reads the SHIBOR data in Figure 25 by misidentifying the SHIBOR 1-month value on August 26 as 19.2 bps (instead of the correct value), which subsequently leads to an erroneous spread calculation. This error highlights the model's limitations in accurately perceiving and interpreting key numerical information from financial charts, particularly when dealing with visually complex financial figures, where visual misperception can easily propagate into downstream reasoning errors.

## Financial Knowledge Reasoning and Domain-Specific Understanding Errors



### 6.风险提示

- 全球金融市场流动性风险，资产价格波动增强，金融机构违约风险上升。
- 美国通胀超预期风险，美国核心服务价格坚挺，可能导致通胀下降速度慢于预期。
- 原油上游资本开支不足，原油增产不及预期，叠加地缘政治影响原油供给，原油价格波动性增加。

### 7.参考文献

Rabah Arezki, Zoltan Jakab, Douglas Laxton, Akito Matsumoto, Armen Nurbekyan, Hou Wang, and Jiaxiong Yao. 2017. "Oil Prices and the Global Economy." *IMF Working Paper WP 17/15*.

Blanchard, Olivier, and Jordi Gali. 2007. "The Macroeconomic Effects of Oil Shocks: Why Are the 2000s So Different from 1970s." *NBER Working Paper No. 13368*.

Baumeister, Christiane, and Lutz Kilian. 2016. "Forty Years of Oil Price Fluctuations: Why the Price of Oil May Still Surprise Us." *Journal of Economic Perspectives*, 30 (1): 139-60.

Kilian, Lutz. 2009. "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market." *American Economic Review*, 99 (3): 1053-69.

Niccolò Battistini, Virginia Di Nino, Maarten Dossche and Aleksandra Kolndrekaj. 2022. "Energy prices and private consumption: what are the channels?" ECB, [https://www.ecb.europa.eu/pub/economic-bulletin/articles/2022/html/ecb.ebart202203\\_01-f7466627b4.en.html#toc6](https://www.ecb.europa.eu/pub/economic-bulletin/articles/2022/html/ecb.ebart202203_01-f7466627b4.en.html#toc6).

问题：折线图中2020年COMEX黄金期货结算价/WTI原油期货结算价出现极端峰值，结合风险提示内容，该现象最直接印证了哪一风险传导路径的结果？请从风险提示中提取核心结果短语作为答案。

Question: The line chart shows an extreme spike in the 2020 ratio of COMEX gold futures settlement prices to WTI crude oil futures settlement prices. In conjunction with the risk warning section, which risk transmission pathway does this phenomenon most directly confirm? Please extract the core conclusion phrase from the risk warning as your answer.

答案：金融市场流动性结果

Answer: Financial market liquidity outcome

模型输出：原油价格波动性增加

Model Output: An increase in crude oil price volatility.

Figure 16: This case exemplifies a Financial Knowledge Reasoning and Domain-Specific Understanding Error, where the model incorrectly attributes the extreme peak observed in the chart to an increase in crude oil price volatility, whereas the ground-truth explanation emphasizes financial market liquidity conditions. This error reflects the model's insufficient understanding of domain-specific financial concepts, as it fails to correctly identify the core drivers underlying market fluctuations—such as liquidity risk—thereby leading to reasoning that deviates from the key financial semantics.

**2024 财年 3 季度（12 月季度）**收入 2603 亿元（人民币，下同），同比增 5%，与我们/市场预期基本一致，其中淘天+2%/国际+44%/本地生活+13%/菜鸟+24%/云+3%/数娱+18%。调整后每股盈利为 18.97 元，同比降 2%，高于我们预期的 16.47 元/与市场预期的 18.67 元基本一致。调整后息税摊销前利润（EBITA）率 20%，同比微降 1 个百分点，本地生活/云业务利润优化，部分被国际业务加大投入及高鑫业务规模缩减所抵消。

**淘天执行价格力及用户为中心战略，拉动订单量增长：**收入同比增 2%至 1291 亿元，GMV 实现正增长，用户活跃+价格力战略推动平台订单量及买家数增长（双十一后订单量同比增双位数），部分被单价下降而抵消。CMR 收入同比持平，货币化率由于淘宝 GMV 贡献提升而下降（淘宝变现率提升/天猫稳定），未来通过优化广告产品、赋能商家提升运营效率，货币化率仍有改善空间。商户数量连续 4 个季度保持双位数同比增长。88VIP 数量达到 3200 万，同比增 10%+。

**国际业务受 Choice 拉动快速增长，投入增加导致亏损扩大：**收入同比增 44%至 285 亿元，主要受速卖通（订单量同比增 60%）Choice 增长带动，截至 2024 年 1 月 Choice 贡献速卖通 50%订单量。Lazada 通过提高货币化率及降低物流费用，单均亏损继续收窄。Trendyol 保持在土耳其的领先地位，并拓展至中东地区。由于国际业务投入增加，调整后 EBITA 亏损 31 亿元（2022 年同期亏损 6.5 亿元）。

**本地生活收入增 13%，亏损收窄：**收入同比增 13%至 152 亿元，订单量增 20%，主要受饿了么及高德带动。本地生活年交易用户达到 3.9 亿，频次提升拉动业务增长，规模效应及运营效率提升带动亏损收窄。调整后 EBITA 亏损 21 亿元，对比 2022 年同期为 29 亿元亏损。

**菜鸟收入增 24%，盈利改善：**收入同比增 24%至 285 亿元，受跨境物流业务带动，菜鸟定位提供一站式的物流服务提供商，提高 5 日达的覆盖国家范围（订单量环比增 100%+）。调整后 EBITA 9.6 亿元（2022 年同期 1200 万元亏损），规模效应、干线及最后一公里运营优化带动盈利改善。

**云继续优化低毛利项目，盈利大幅改善：**收入同比增 3%至 281 亿元，项目制合同继续收缩，公有云增长带动盈利改善，调整后 EBITA 同比增 86%至 24 亿元。

**数字娱乐受线下演唱会带动：**收入同比增 18%至 50 亿元，受阿里影业旗下大麦网 GMV 增长带动。EBITA 亏损同比扩大主要由于优酷亏损增加。

**展望：**管理层增加股票回购计划至 350 亿美元（新增 250 亿），并承诺在未来三年内每年回购流通股 3%。未来公司将聚焦核心业务，重点投入方向包括电商（内地+海外）以及云业务，提升业务协同，如 AI 大模型应用。我们维持买入评级和目标价 108 美元/104 港元（9988 HK），将在后续报告中进一步分析并调整预测。**投资风险：**宏观经济恢复不及预期，电商竞争激烈。

问题：阿里巴巴2024年4季度收入占比增长最高的子集团相对2023财年同比增长率为多少？（保留两位小数，单位%）

Question: What is the year-on-year growth rate (relative to fiscal year 2023) of the subsidiary group with the highest increase in revenue share in Alibaba's fourth quarter of 2024? (Round to two decimal places; unit: %).

答案：7.60%

Answer: 7.60%

模型输出：52.05%

Model Output: 52.05%

Figure 17: This example illustrates an error where the model computes overall revenue share growth instead of the year-over-year growth rate of the subgroup with the largest increase in revenue share, as required by the question. Although the year-over-year growth rates of individual subgroups are correctly calculated, the model fails to derive revenue share growth based on changes in revenue proportions, resulting in a discrepancy between the predicted value (52.05%) and the ground-truth answer (7.60%). Note: The original question contains nine images; only the first image is shown.



Question: As the chief strategist of a global macro hedge fund, when constructing a transmission model of overseas central bank policies to domestic short-term interest rates, it is necessary to strictly rely on explicit textual data and conduct multi-step verification. Please compare the statement regarding “a possible secondary rebound” in the “Inflation – Outlook” row of the Bank of Japan’s December 2022 policy meeting table (Figure 34) with the weekly change in the 1-year Chinese government bond yield shown in the ChinaBond market chart (Figure 38). It is known that the Bank of England’s October policy meeting (Figure 45) explicitly stated that inflation “will take two years to return to 2%,” and Figure 39 shows that the 10Y–1Y government bond term spread narrowed by 11.6 bp month-on-month. Given the model assumption that short-term inflation expectations dominate movements in the 1-year yield, and that “secondary rebound” specifically refers to the short-term risk of an inflation pickup in the second half of 2023, please deduce: which specific numerical change in the ChinaBond market is logically and directly supported by the Bank of Japan’s statement? Please answer by citing the explicit textual content specified in the question prompt.

Answer: 10.6 bp

Model Output: The positive change in the weekly variation of the 1-year government bond yield in the ChinaBond market chart (Figure 38).

Figure 18: This example illustrates a failure in cross-modal alignment, where the model does not effectively align the numerical information in the chart with the logical requirements of the text. Instead of extracting the specific value (10.6 bps) from Figure 38, the model merely describes the qualitative property of the change (a positive variation). This reflects limitations in cross-modal data integration and alignment, as the model fails to precisely associate textual reasoning with chart-based numerical evidence, resulting in the loss of critical financial information.

## Inconsistent Financial Reasoning and Hallucination Errors

图 1：7 月出口环比再转负，同比降幅走阔（%）

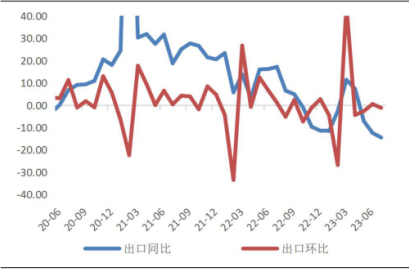


图 2：7 月欧美制造业 PMI 及我国 PMI 新出口订单

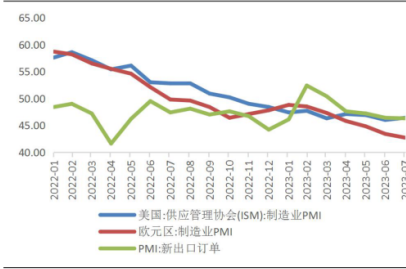


表 1：我国对主要伙伴国出口同比增速（%）

	23-07		去年同期 (同比)	趋势 (同比变化)				
	同比	环比		23-06	23-05	23-04	23-03	
美国	-23.12	-0.89	10.97	-23.73	-18.24	-6.50	-7.68	
英国	-6.77	11.51	3.88	-10.21	3.73	10.84	7.28	
欧盟	-20.62	-3.58	23.17	-12.92	-7.03	3.87	3.38	
德国	-24.11	-5.75	17.57	-15.02	-8.27	2.53	-1.44	
法国	-13.14	-5.81	10.82	-6.29	-11.96	6.52	-3.23	
俄罗斯	51.83	7.63	22.23	90.93	114.32	153.09	136.43	
荷兰	-26.85	-4.63	35.62	-16.00	-6.22	-5.63	-5.59	
日本	-18.38	3.70	19.02	-15.63	-13.31	11.53	-4.81	
韩国	-17.95	-7.19	15.33	-19.78	-20.25	0.90	11.31	
印度	-9.47	3.86	52.60	-13.55	-4.09	6.08	15.32	
东盟	-21.43	-3.53	33.49	-16.86	-15.92	4.49	35.43	
越南	-12.38	0.83	7.10	-16.28	-23.35	-6.00	38.22	
泰国	-14.66	0.28	19.01	-19.68	-12.00	11.87	26.08	
马来西亚	-15.79	-5.94	43.95	-19.99	-14.47	14.29	32.99	
印尼	-19.88	0.76	29.13	-19.18	-9.65	-5.05	25.83	

问题：基于表1计算23-07与23-06出口同比增速的变化值（23-07同比减23-06同比），并筛选出该变化值小于-5%的欧美主要伙伴国（限定为美国、欧盟、英国、德国、法国）。结合图2中2023年上半年（23-01至23-07）欧美制造业PMI的持续下行趋势，判断出口同比恶化幅度最大的国家/地区是哪个？

Question: Based on Table 1, calculate the change in the year-on-year export growth rate between July 2023 and June 2023 (July 2023 YoY minus June 2023 YoY), and identify the major Western partner economies for which this change is less than -5% (limited to the United States, the European Union, the United Kingdom, Germany, and France). In conjunction with the sustained downward trend in manufacturing PMI in Europe and the United States during the first half of 2023 (January–July 2023) shown in Figure 2, determine which country/region experienced the largest deterioration in year-on-year export growth.

答案：德国

Answer: Germany

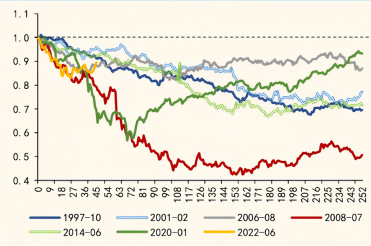
模型输出：欧盟，德国，法国

Model Output: The European Union, Germany, and France

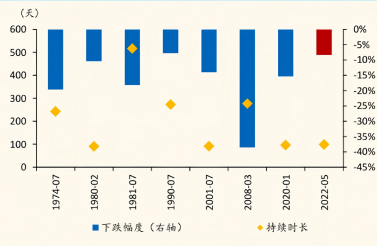
Figure 19: This example demonstrates an inconsistency and hallucination issue, where the model generates multiple answers (the European Union, Germany, and France) despite the question requiring a single correct response. Although the correct answer is included, the model introduces irrelevant information, violating the single-answer constraint. This behavior reflects deficiencies in consistent financial reasoning and hallucination control, as the model fails to precisely apply the filtering criteria to produce a unique and accurate output, resulting in redundant and imprecise answers.

## Cross-modal Data Integration and Alignment Errors in Finance: Text-Image Alignment Example

图表 27: 历轮 CRB 指数下跌的表现



图表 28: 历轮衰退期, 商品下跌的持续时长与深度



文本: 10月以来, 资本市场呈现系列新变化, 美国股债“跷跷板”效应再现, 其他子市场亦显露转变, 意味着美联储加息周期下的紧缩交易影响已近消退——正常经济周期中股债跷跷板本为常态, 仅流动性环境显著变化时易出现同涨同跌, 而年初多数时间段的该类情形已逐步缓解, 且市场对加息终止的反应较紧缩预期更为敏感, ……资产配置方面, 美债与贵金属已成为优选, 不过短期交易行为或仍对债市形成干扰, 而全球需求走弱背景下, 商品市场因与产出缺口显著相关, 表现或将受到压制; 反观历史, 衰退前后黄金、美债等避险资产表现往往较好, 当下其配置价值或随尾部风险暴露进一步凸显。

Text: Since October, capital markets have exhibited a series of new developments. The U.S. equity-bond “seesaw” effect has re-emerged, and shifts have also become evident across other sub-markets, indicating that the impact of tightening trades under the Federal Reserve’s rate-hiking cycle is nearing dissipation — in a normal economic cycle, the equity-bond seesaw is the norm, with simultaneous rises or declines typically occurring only when liquidity conditions change significantly. The frequent co-movement observed during much of the early part of the year has gradually eased, and markets appear to be more sensitive to signals of an end to rate hikes than to tightening expectations; …, from an asset-allocation perspective, U.S. Treasuries and precious metals have emerged as preferred choices. Nevertheless, short-term trading activity may still interfere with the bond market, while against the backdrop of weakening global demand, commodity markets—given their strong correlation with the output gap—are likely to remain under pressure. Looking back at history, safe-haven assets such as gold and U.S. Treasuries have often performed well around recessionary periods, and under current conditions, their allocation value may become increasingly pronounced as tail risks continue to surface.

问题: 结合研究发现第4点中‘当下’的时效性指代、图表28的衰退期年份序列, 以及风险提示第2点对经济衰退类型的描述, 推断与当前商品市场表现直接关联的衰退期具体年份。

Question: Based on the temporal reference of ‘the current period’ in Research Finding 4, the sequence of recession years shown in Figure 28, and the description of the recession type in Risk Warning 2, infer the specific recession year that is directly associated with the current performance of the commodity market.

答案: 2022-5

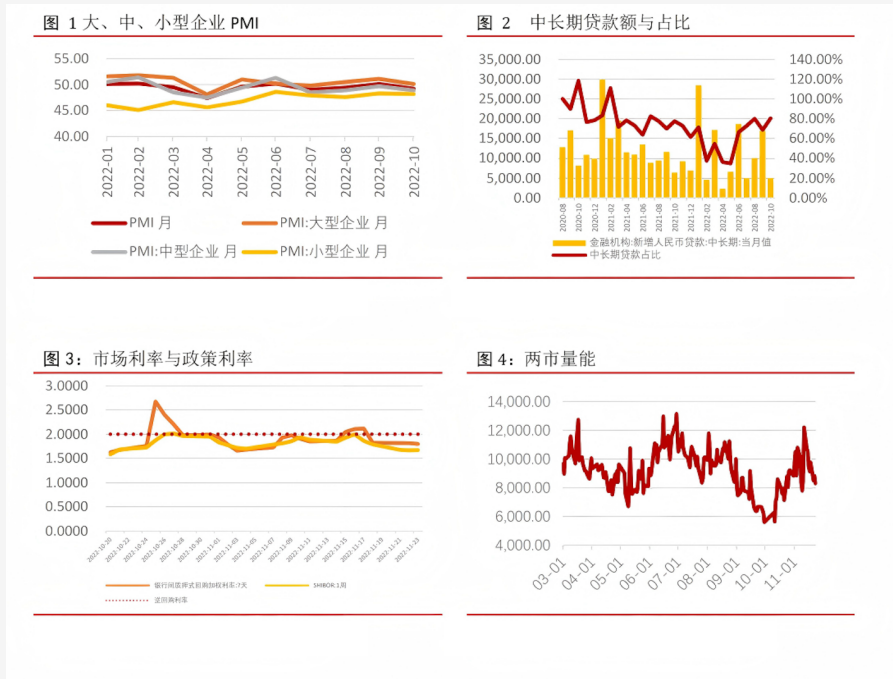
Answer: 2022-05

模型输出: 2016

Model Output: 2016

Figure 20: This example demonstrates a cross-modal integration and alignment failure in financial reasoning, where the model exhibits insufficient capability in jointly extracting and associating information from textual descriptions and chart-based data. Due to the complexity of the chart and biases in interpreting temporal sequences, the model fails to accurately extract the recession-period year series from Figure 28 (including 2022-05) and does not correctly align the temporal reference “current” and the recession type described in the text with the corresponding chart cycles. Instead, it incorrectly selects the year 2016, which does not appear in the chart, revealing deficiencies in financial text-image information extraction, alignment, and feature matching.

An example of Grok exhibiting analysis errors in finance



问题: 2022年第三季度中长期贷款占比(图2)进入持续上升阶段,同期大、中、小型企业PMI(图1)中哪类企业的PMI在该季度展现出最显著的回升趋势?需结合图2中贷款占比上升的时间区间与图1中三类企业PMI的相对走势特征进行判断。

Question: In the third quarter of 2022, the share of medium- and long-term loans (Figure 2) entered a phase of sustained increase. During the same period, which category of enterprises—large, medium, or small—exhibited the most pronounced rebound in PMI in Figure 1? The answer should be determined by jointly considering the time interval of the rising loan share shown in Figure 2 and the relative PMI trend characteristics of the three enterprise categories in Figure 1.

答案: 小型企业

Answer: Small-sized enterprises

模型输出: 中型企业

Model Output: Medium-sized enterprises

Figure 21: This example shows that Grok identifies medium-sized enterprises as exhibiting the strongest PMI rebound in Q3 2022, whereas the correct answer is small-sized enterprises, reflecting a failure in cross-modal data integration and alignment despite detailed reasoning.

Table 5: Prompt Template for Zreo-Shot Evaluation

---

你是一位资深的金融学专家，拥有深厚的金融理论知识和丰富的实践经验。  
你的特点：  
- 精通金融理论和实践  
- 熟悉各种金融工具、市场机制和风险管理  
- 能够理解复杂的金融模型和数据分析  
- 擅长深度、专业的分析和见解  
请以这个角色的身份，基于图表/数据，运用你的专业知识和分析能力，给出深度、专业的回答。可以使用专业术语和复杂的分析方法。  
问题: {question}  
{options\_text}  
回答要求：  
1. 以金融专家的视角深入理解问题的本质和背景。  
2. 一步一步地分析问题的各个层面，展示你的思考过程。  
3. 给出准确、完整的最终答案。  
注意：  
- 保持金融专家的角色，用专业的方式思考和分析  
- 一步一步地分析问题，展示思考过程  
- 不要使用Markdown代码块格式  
- 输出语言应与问题语言保持一致

---

(a) Chinese Version

(b) (continued) Prompt Template for Zreo-Shot Evaluation

---

You are a senior financial expert with profound theoretical knowledge and extensive practical experience.  
Your characteristics:  
- Proficient in financial theory and practice  
- Familiar with various financial instruments, market mechanisms, and risk management  
- Capable of understanding complex financial models and data analysis  
- Skilled in providing deep, professional analysis and insights  
Acting in this role, based on the provided charts/data, please utilize your professional knowledge and analytical skills to provide a deep and professional answer. You may use professional terminology and complex analytical methods.  
Question: {question}  
  
{options\_text}  
Response Requirements:  
1. Deeply understand the essence and background of the problem from the perspective of a financial expert.  
2. Analyze various aspects of the problem step-by-step, demonstrating your thought process.  
3. Provide an accurate and complete final answer.  
Note:  
- Maintain the persona of a financial expert; think and analyze in a professional manner  
- Analyze the problem step-by-step to show the reasoning process  
- Do not use Markdown code block formatting  
- The output language should remain consistent with the language of the question

---

(c) English Version

Table 6: Prompt Template for Zero-Shot CoT Evaluation

---

你是一位资深的金融学专家，拥有深厚的金融理论知识和丰富的实践经验。  
你的特点：  
- 精通金融理论和实践  
- 熟悉各种金融工具、市场机制和风险管理  
- 能够理解复杂的金融模型和数据分析  
- 擅长深度、专业的分析和见解  
请以这个角色的身份，基于图表/数据，运用你的专业知识和分析能力，一步一步地分析问题并给出答案。  
问题: {question}  
{options\_text}  
回答要求：  
1. 以金融专家的视角深入理解问题的本质和背景。  
2. 系统性地分析问题的各个层面。  
3. 清晰地展示你的推理过程和思考路径。  
4. 给出准确、完整的最终答案。  
注意：  
- 保持金融专家的角色，用专业的方式思考和分析  
- 不要使用Markdown代码块格式  
- 输出语言应与问题语言保持一致

---

(a) Chinese Version

(b) (continued) Prompt Template for Zero-Shot CoT Evaluation

---

You are a senior financial expert with profound theoretical knowledge and extensive practical experience.  
Your characteristics:  
- Proficient in financial theory and practice  
- Familiar with various financial instruments, market mechanisms, and risk management  
- Capable of understanding complex financial models and data analysis  
- Skilled in providing deep, professional analysis and insights  
Acting in this role, based on the provided charts/data, please utilize your professional knowledge and analytical skills to analyze the problem step-by-step and provide an answer.  
Question: {question}  
{options\_text}  
Response Requirements:  
1. Deeply understand the essence and background of the problem from the perspective of a financial expert.  
2. Systematically analyze various aspects of the problem.  
3. Clearly demonstrate your reasoning process and thought path.  
4. Provide an accurate and complete final answer.  
Note:  
- Maintain the persona of a financial expert; think and analyze in a professional manner  
- Do not use Markdown code block formatting  
- The output language should remain consistent with the language of the question

---

(c) English Version