Finance-HalluBench: Benchmarking LLM Hallucination and Calibration Failure in SEC Financial Filings

Anonymous submission

Abstract

Analysts who use large language models (LLMs) to read SEC filings need two things: correct answers when the filing contains the answer, and an explicit "not stated" (rather than a confident fabrication) when it does not. We present Finance-HalluBench, a benchmark of 245 audited questions grounded in 10-K excerpts from 12 large U.S. issuers, spanning numerical extraction, temporal (direction-of-change) reasoning, and causal-driver identification. A distinguishing design choice is an explicit *unanswerable* split: for 133 numerical questions the requested line item is provably absent from the excerpt, so any number produced is unsupported by the evidence (a fabrication with respect to the excerpt). We pair the benchmark with a strict first-answer grader, scale-invariant numeric matching, and calibration metrics (ECE, Brier, AUROC) with bootstrap confidence intervals. To isolate the effect of post-training on calibration without cross-family confounds, we evaluate a single base model through its released post-training stages (OLMo-2-1B Base $\rightarrow$ SFT $\rightarrow$ DPO $\rightarrow$ RLVR), and contrast it with four frontier models using verbalized confidence. We also report an audit of our own earlier evaluation, showing that lenient substring grading can move accuracy by by up to 69 points on a task. On the controlled ladder, SFT teaches abstention (correct abstention on unanswerable items rises from 0% to 74%). DPO undoes part of this gain: fabrications rise from 34 to 59 of 133 items, and the model becomes overconfident again (Δ_{CA}=+0.17). The final RLVR stage restores SFT-level behavior. When abstention is offered, frontier models are near ceiling (accuracy ≥ 0.98, ECE ≤ 0.02). Removing the single sentence that offers `NOT_FOUND` has little effect on three of them, which answer only 5–8% of unanswerable items. GPT-5.5, however, then answers 58% of these items with mean confidence 0.76, and most of its answers reproduce the true figure from memory rather than from the excerpt. Abstention is therefore a first-order design variable for financial QA evaluation.

Introduction

LLMs are increasingly used to read and summarize corporate disclosures, and specialized systems such as BloombergGPT (Wu et al. 2023) and FinGPT (Yang, Liu, and Wang 2023) target exactly this setting. In financial use, the cost of an error is asymmetric: a fabricated revenue figure that is stated fluently and with high confidence propagates into models, memos, and trades, whereas an honest "the excerpt does not state this" merely triggers a lookup. What matters is therefore not only accuracy, but whether a model (i) abstains when the evidence is absent and (ii) attaches confidence that separates its correct from its incorrect answers.

Prior work suggests that post-training can damage this second property. Pre-trained models are often reasonably calibrated (Kadavath et al. 2022), while preference optimization has been reported to inflate confidence (OpenAI 2023; Tian et al. 2023; Leng et al. 2025), and recent analyses argue that binary-graded training and evaluation actively reward guessing over abstention (Kalai et al. 2025). Most existing financial QA benchmarks, however, score accuracy only (Chen et al. 2021; Zhu et al. 2021; Islam et al. 2023; Choi et al. 2025), and published comparisons of "base vs. aligned" models typically compare checkpoints from different families or with undisclosed training recipes, which confounds alignment with data and architecture.

Finance-HalluBench targets this gap. Our contributions are:

- **An audited benchmark with an unanswerable split.** 245 questions over 10-K excerpts (FY2021–2024) from 12 issuers, obtained by auditing a 309-item automatically generated pool. Numerical items are split into answerable and unanswerable ones, turning a data-construction artifact into a controlled probe of fabrication (Sec. Benchmark).
- **A strict, reproducible evaluation protocol**: first-answer parsing, scale-invariant numeric matching, explicit abstention detection, and ECE/Brier/AUROC with bootstrap CIs.
- **A confound-free post-training study.** One base model evaluated through its public SFT, DPO, and RLVR checkpoints (Team OLMo et al. 2024; Lambert et al. 2024), plus four frontier models.
- **An evaluation audit.** Re-scoring the generations from an earlier version of this benchmark shows how lenient substring matching distorts accuracy and conclusions; we release both graders.

This is an anonymized submission for review purposes only. Distribution, citation, or public sharing of this manuscript is strictly prohibited. Copyright and publication details will appear in the final version if accepted.

Related Work

Hallucination, abstention, and calibration. Hallucination in generation is surveyed by Ji et al. (2023). Calibration of neural classifiers is typically measured with ECE (Naeini, Cooper, and Hauskrecht 2015; Guo et al. 2017) and the Brier score (Brier 1950); selective prediction asks whether confidence can rank errors (Geifman and El-Yaniv 2017). For LLMs, token-probability confidence (Kadavath et al. 2022) and verbalized confidence (Tian et al. 2023; Xiong et al. 2024) behave differently, and RLHF-style training has been linked to overconfidence (OpenAI 2023; Leng et al. 2025). Unanswerable questions as a test of "knowing what you don't know" date to SQuAD 2.0 (Rajpurkar, Jia, and Liang 2018); AbstentionBench (Kirichenko et al. 2025) recently found that reasoning-tuned models abstain poorly. Kalai et al. (2025) argue that accuracy-only grading incentivizes guessing.

Financial NLP benchmarks. FinQA (Chen et al. 2021) and TAT-QA (Zhu et al. 2021) test numerical reasoning over reports; FinanceBench (Islam et al. 2023) tests open-book QA over SEC filings; FinDER (Choi et al. 2025) evaluates retrieval-augmented QA over 10-Ks; FinBias (Mehrotra et al. 2025) studies bias in financial LLMs. These benchmarks report accuracy (or fairness) but not calibration, and none contains a controlled unanswerable split. Finance-HalluBench is complementary: it is small and focused, and it measures whether confidence and abstention are trustworthy rather than raw capability.

Benchmark

Source and tasks. Instances are drawn from the MD&A and financial-statement sections of 10-K filings retrieved from SEC EDGAR for 12 large-cap issuers (AAPL, ABBV, AMD, AMZN, BRK, CRM, GOOGL, HD, MSFT, NKE, NVDA, TSLA). Each instance pairs a 167–600 character excerpt with a question of one of three types: *numerical extraction* (report the value of a named line item), *temporal reasoning* (did a named metric increase or decrease year-over-year), and *causal reasoning* (state the driver of a change in a named metric, as given in the excerpt).

Audit and cleaning (v1.0 → v1.1). The initial pool of 309 automatically generated items (243 numerical, 41 temporal, 25 causal) had systematic defects that we found only after manual inspection; all fixes are scripted and released. (1) *Answerability.* In 159 numerical items the excerpt window does not contain the gold value at any scale (e.g., "$23.6 billion" counts as containing 23,601 million). Instead of discarding them, we keep those whose excerpt states *no* value for the requested line item in any year and relabel their gold answer as `NOT_FOUND`. This gives the **Num-U** split (n=133; 4 items whose excerpt reports the line item for another year were removed). (2) *Invalid labels.* 46 items had non-financial labels produced by the parser ("Item 6.", "2022", "(Dollars in millions)"); 2 had an artifact gold value ("365365" from "Office 365"); 8 had a gold value that did not correspond to the line item (a percentage change, a common-size percentage, or a mis-rounded per-share value). All were removed. (3) *Fiscal-year alignment.* The year field did not consistently match the fiscal year of the gold value (for some issuers the gold value was the most recent column, for others the prior year). Each question now states the fiscal year explicitly. (4) *Units.* All values had been labelled "millions", including per-share and percentage items. We therefore grade numeric answers scale-invariantly instead of using the unit field. (5) *Underspecified questions.* Temporal and causal questions asked about "the metric described in the passage" even though excerpts often mention several metrics. Each question now names its target metric. Steps (2), (3) and (5) were assisted by an LLM annotator (GPT-6.1-Sol) given the gold label, which it could not change. We added automatic leakage checks (no direction words in temporal questions; no overlap with the gold span in causal questions) and manually reviewed every flagged or rewritten item (24 manual corrections). We removed 2 temporal items with no metric matching the gold label and 2 causal items with malformed gold spans. Because the same model family is also evaluated, we never use annotator outputs as gold answers. Table 1 summarizes v1.1.

Split	n	Gold answer
Num-A (answerable)	50	value in excerpt
Num-U (unanswerable)	133	`NOT_FOUND`
Temporal	39	increased 29 / decreased 10
Causal	23	span from excerpt
Total	245	12 issuers, 10-K excerpts

Table 1: Finance-HalluBench v1.1 composition.

Evaluation Protocol

Prompting. All models receive the same instruction: answer using only the excerpt; for numerical items, output the number only, or `NOT_FOUND` if the excerpt does not state the value; for temporal items, one word; for causal items, the short phrase from the excerpt that states the driver (full prompts in the Appendix). Open-weight models are decoded greedily; API models use provider-default sampling with a single sample per item.

Strict grading. Only the first line of a generation is graded, which prevents a model from receiving credit for continuations that mention both options (a failure mode of substring matching; see Sec. Results). Numerical answers are parsed with optional scale words and accepted if they match the gold value within 0.5% relative tolerance at any power-of-10^3 scale. An answer counts as an abstention if it contains `NOT_FOUND` or an explicit "not stated/reported" phrase. On Num-U, abstention is correct and any non-abstaining answer is a *fabrication* with respect to the excerpt; we further classify fabricated numbers as *memorized* (equal to the true filing value), *misattributed* (equal to another number in the excerpt), or *invented*. Temporal answers are graded on the first direction word; causal answers by token-level F1 ≥ 0.5 against the gold span.

Confidence. For open-weight models we use *token-likelihood confidence*: the mean, over the tokens of the graded answer, of the probability of the greedily chosen token. This is a property of the model's output distribution, not a self-report. For API models, whose token probabilities are not

uniformly available, we elicit *verbalized confidence* (Tian et al. 2023; Xiong et al. 2024) by requesting a JSON object with the answer and the probability that it is correct. The two estimators are not directly comparable, so we compare within each group.

Metrics. We report accuracy, mean confidence, the confidence–accuracy gap $\Delta_{CA} = \overline{\text{conf}} - \text{acc}$ (positive = overconfident), 10-bin ECE, Brier score, and AUROC of confidence for discriminating correct from incorrect answers (0.5 = no discrimination). For Num-U we report abstention and fabrication rates and the mean confidence attached to fabrications. All intervals are 95% percentile bootstrap CIs over items (1,000 resamples).

Models. *Controlled post-training ladder:* OLMo-2-0425-1B Base, SFT, DPO, and Instruct (final RLVR stage) (Team OLMo et al. 2024), which share one pre-trained base and follow the public Tülu 3 recipe (Lambert et al. 2024). The base model is prompted with the same text in a plain completion format; post-trained stages use their chat template. *Frontier models* (verbalized confidence): GPT-6-Luna (no reasoning), GPT-5.5 (low reasoning effort), GPT-6.1-Sol (low), and Gemini-3.8-Flash (low), all with JSON output. *v1.0 generations:* we additionally re-score the released generations of Qwen2.5-7B, Qwen2.5-7B-Instruct, Qwen2.5-1.5B-Instruct (Qwen Team 2024), and Mistral-7B-Instruct-v0.3 from the earlier version.

Results

Finding 1: the v1.0 headline numbers were an evaluation artifact. The earlier version of this benchmark reported accuracy of 0.20–0.23 for all four open models and interpreted the differences between base and aligned models. Table 3 re-scores the same generations, restricted to the 245 items retained in v1.1 and using the v1.0 gold labels, with the strict first-answer grader. Between 24 and 56 correctness labels flip per model. The largest distortion is on temporal items, where Qwen2.5-1.5B-Instruct drops from 0.97 to 0.28 (69 points). Its generations typically continue past the answer and contain both "increased" and "decreased", which a substring grader counts as correct. Conversely, strict scale-invariant matching *raises* several models' numerical accuracy, because the lenient grader failed on formatting ("$23.6 billion" vs. 23,601). The low overall accuracy had a simpler cause: for most numerical items the excerpt did not contain the gold value (Sec. Benchmark), and the v1.0 prompt offered no way to abstain. Under those conditions the models produced a number on 30–100% of the now-unanswerable items. We therefore withdraw the v1.0 cross-model conclusions. The remaining results use v1.1 and the strict grader only.

Finding 2: in a controlled ladder, SFT teaches abstention and DPO erodes it. Figure 1 and the upper block of Table 2 follow one 1B base model through its released post-training stages. (i) The *base* model never abstains: it produces a number for all 133 Num-U items. Its token-likelihood confidence is nevertheless informative. Fabrications receive mean confidence 0.46, compared with 0.74 for correct answers (AUROC 0.88 [0.83, 0.92]), but the model is strongly overconfident overall (Δ_{CA}=+0.27). (ii) *SFT* teaches abstention. Correct abstention on Num-U rises from 0.00 to 0.74, Num-A accuracy rises from 0.58 to 0.74, and overall accuracy rises from 0.32 to 0.73. However, AUROC collapses to chance (0.51). This happens because `NOT_FOUND` strings, although usually correct, are emitted with low token likelihood (mean 0.59). Restricted to non-abstained answers, AUROC is 0.78. (iii) *DPO* reverses part of these gains. Fabrication on Num-U increases from 34 to 59 items, Num-A accuracy falls to 0.44, and the model becomes overconfident again (Δ_{CA}=+0.17 [0.10, 0.23] vs. −0.02 [−0.09, 0.03] after SFT). Fabrications are now made with higher confidence (0.71 vs. 0.60). None of the OLMo fabrications reproduces the true filing value. Base-model fabrications are mostly invented numbers (106 of 133), whereas after post-training most fabrications are misattributed numbers copied from the excerpt (25 of 34 after SFT and 40 of 59 after DPO). The drop in overall accuracy from SFT to DPO (0.73 to 0.58) has non-overlapping 95% CIs. (iv) The final *RLVR* stage returns to SFT-level abstention (0.74) and calibration (Δ_{CA}=+0.02, best ECE 0.21), with AUROC on answered items of 0.82. Thus, in this controlled ladder, the stage that most degrades both abstention and calibration is preference optimization, consistent with reports that RLHF-style objectives inflate confidence (Leng et al. 2025). Verifiable-reward RL does not make it worse. One caveat applies to every stage: all four predict "increased" for every temporal item, so their temporal accuracy (0.74) equals the majority-class rate. Temporal reasoning at 1B scale is therefore not measured meaningfully, and we do not interpret it.

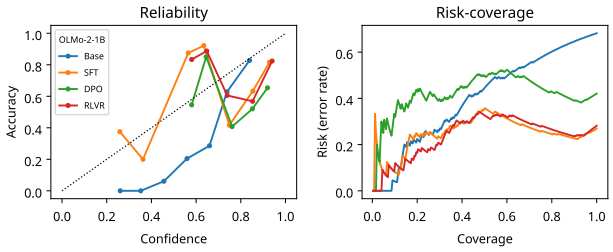

Figure 1: OLMo-2-1B post-training ladder. Left: reliability diagram (bins with ≥5 items). Right: risk–coverage curve when answers are ranked by token-likelihood confidence. Frontier models lie at the top-right corner of the reliability plot (not shown).

Finding 3: frontier models saturate the grounded setting when abstention is offered. With the `NOT_FOUND` option, all four frontier models reach 0.98–0.99 accuracy with ECE ≤ 0.02 (Table 2). They fabricate on at most 2 of 133 Num-U items, and their temporal and causal accuracy is ≥ 0.91. Because overall accuracy is near ceiling, AUROC rests on only 3–6 errors per model and its intervals are wide. We read this as a negative result about the benchmark's difficulty, not as evidence that frontier models are reliable on filings in general: short, single-passage excerpts with an explicit abstention channel are not where current frontier models fail.

Finding 4: without an explicit abstention option, models answer from memory or from the wrong line. Ta-

Model	Acc. [95% CI]	Conf.	Δ_{CA}	ECE [95% CI]	Brier	AUROC [95% CI]	$\text{AUROC}_{\text{ans}}$
OLMo-2-1B Base	0.32 [0.26,0.38]	0.58	+0.27	0.27 [0.23,0.32]	0.22	0.88 [0.83,0.92]	0.88
+SFT	0.73 [0.68,0.79]	0.71	−0.02	0.24 [0.20,0.29]	0.22	0.51 [0.42,0.59]	0.78
+DPO	0.58 [0.52,0.64]	0.75	+0.17	0.28 [0.23,0.34]	0.29	0.46 [0.38,0.54]	0.72
+RLVR (Instruct)	0.72 [0.66,0.77]	0.74	+0.02	0.21 [0.16,0.26]	0.21	0.53 [0.45,0.61]	0.82
GPT-6-Luna	0.98 [0.96,1.00]	1.00	+0.02	0.02 [0.00,0.04]	0.02	0.88 [0.83,0.95]	0.86
GPT-5.5	0.99 [0.97,1.00]	0.99	+0.00	0.01 [0.00,0.02]	0.01	0.99 [0.98,1.00]	1.00
GPT-6.1-Sol	0.98 [0.96,1.00]	1.00	+0.01	0.01 [0.00,0.03]	0.01	0.98 [0.94,1.00]	0.97
Gemini-3.8-Flash	0.98 [0.96,1.00]	0.99	+0.01	0.01 [0.00,0.03]	0.02	0.95 [0.91,0.97]	0.87

Table 2: Overall results on Finance-HalluBench v1.1 (n=245). OLMo-2 rows share one base model and differ only in post-training stage; confidence is the mean token probability of the graded answer. Frontier rows use verbalized confidence. Brackets: 95% bootstrap CIs (1,000 resamples). Δ_{CA} = mean confidence − accuracy. $\text{AUROC}_{\text{ans}}$: AUROC restricted to non-abstained answers.

v1.0 model	Acc_{orig}	$\text{Acc}_{\text{strict}}$	Temp (o/s)	Fab.	Flips
Mistral-7B-Inst.	0.23	0.30	0.82 / 0.82	0.30	24
Qwen2.5-1.5B-Inst.	0.22	0.21	0.97 / 0.28	1.00	56
Qwen2.5-7B	0.20	0.28	0.69 / 0.69	1.00	27
Qwen2.5-7B-Inst.	0.21	0.29	0.77 / 0.72	0.53	31

Table 3: Audit of the v1.0 evaluation: re-scoring the released v1.0 generations with the strict first-answer grader. Fab. = fraction of Num-U items on which a number was produced. Flips = items whose correctness label changes.

	Num-A	Num-U (n=133)			Temp	Caus
Model	Acc	Abst.	Fab.	Conf_{fab}	Acc	Acc
OLMo-2-1B Base	0.58	0.00	1.00	0.46	0.74	0.87
+SFT	0.74	0.74	0.26	0.60	0.74	0.61
+DPO	0.44	0.56	0.44	0.71	0.74	0.74
+RLVR (Instruct)	0.60	0.74	0.26	0.63	0.74	0.78
GPT-6-Luna	0.98	0.99	0.01	0.99	0.97	0.91
GPT-5.5	1.00	0.99	0.01	0.78	0.97	0.96
GPT-6.1-Sol	0.98	0.98	0.02	0.98	0.97	1.00
Gemini-3.8-Flash	0.96	1.00	0.00	–	1.00	0.91

Table 4: Results by split. Num-A: answerable numerical (n=50); Num-U: value absent from the excerpt; Abst. = correct abstention rate, Fab. = fabrication rate (any non-abstaining answer), Conf_{fab} = mean confidence on fabricated answers. Temp n=39, Caus n=23.

	Num-A Acc	Num-U fabrication			
Model	with / w/o	with	w/o [95% CI]	Conf_{fab}	M/A/I
OLMo-2-1B Base	0.58 / 0.62	1.00	1.00 [1.00,1.00]	0.50	0/24/109
+SFT	0.74 / 0.82	0.26	1.00 [1.00,1.00]	0.68	2/66/65
+DPO	0.44 / 0.48	0.44	1.00 [1.00,1.00]	0.72	3/65/65
+RLVR (Instruct)	0.60 / 0.76	0.26	0.98 [0.96,1.00]	0.50	0/48/83
GPT-6-Luna	0.98 / 1.00	0.01	0.05 [0.02,0.09]	0.99	0/7/0
GPT-5.5	1.00 / 1.00	0.01	0.58 [0.50,0.66]	0.76	53/4/20
GPT-6.1-Sol	0.98 / 1.00	0.02	0.08 [0.04,0.14]	0.76	0/11/0
Gemini-3.8-Flash	0.96 / 1.00	0.00	0.05 [0.02,0.09]	0.87	0/7/0

Table 5: Abstention-instruction ablation on the 183 numerical items. "with": the standard prompt offering `NOT_FOUND`; "w/o": the same prompt without that option. Conf_{fab} = mean confidence on fabricated answers in the w/o condition. M/A/I: w/o-condition fabrications that equal the true filing value (memorized), equal another number in the excerpt (misattributed), or neither (invented).

ble 5 deletes the single instruction sentence that offers `NOT_FOUND` and keeps everything else fixed. On answerable items, accuracy rises slightly to 1.00 for every frontier model. The abstention option therefore costs 1–2 false abstentions per model. On Num-U, three frontier models still decline in most cases and answer on only 5–8% of items. Their non-abstaining answers are all *misattributions*: a number that does appear in the excerpt, but for a different line item or year. GPT-5.5 behaves very differently. It answers 58% [49, 66] of the unanswerable items with mean verbalized confidence 0.76. Of these 77 answers, 53 reproduce the *true* figure from the company's 10-K even though it is absent from the excerpt, and 20 are numbers that appear nowhere in it. Such answers are correct about the world but ungrounded in the evidence. A grader that checks only against the filing would reward them, whereas a grounded-reading benchmark, and an auditor who needs a traceable source, should not. The OLMo ladder shows the same dependence in its most extreme form. Without the option, no stage abstains spontaneously: fabrication on Num-U is 0.98–1.00 at every stage, and almost all of these answers are misattributed or invented numbers rather than recalled ones. The 74% abstention reached after SFT and RLVR is therefore produced entirely by the instruction. Removing the option also raises Num-A accuracy for every stage, from 0.74 to 0.82 after SFT and from 0.60 to 0.76 after RLVR. The small post-trained models therefore buy their abstention partly with false abstentions on answerable items. These results have two implications. First, whether a financial QA benchmark offers an abstention option changes what it measures. Without one, part of the score rewards guessing or parametric recall (Kalai et al. 2025). Second, models of similar capability differ substantially in whether they stay within the provided evidence when they are not explicitly allowed to abstain, so this behavior should be tested for each model before deployment.

Limitations

The benchmark is small (245 items, 12 issuers, short excerpts), so per-split estimates for Temporal (n=39) and Causal (n=23) have wide intervals; we report CIs throughout and avoid claims that hinge on differences within them. Excerpts are short windows rather than full filings, so the benchmark tests grounded reading, not retrieval. Num-U is constructed from windows that miss the target value; a model may still know the true figure from pre-training, so we report how many "fabrications" coincide with the true value. The temporal split is imbalanced (29 increased vs. 10 decreased). Question disambiguation used an LLM, mitigated by gold-preserving prompts, leakage checks, and full manual review. The controlled ladder uses a 1B model for compute reasons; whether the stage-wise pattern holds at larger scale is open (OLMo-2 releases 7B–32B ladders). API models were sampled once per item with provider defaults, so their results carry sampling variance that our item-level CIs do not capture. Token-likelihood and verbalized confidence are different estimators, so open-weight and API results should not be compared directly.

Ethics Statement

All source text comes from public SEC filings. The benchmark is intended to evaluate, not to endorse, the use of LLMs for financial analysis; our results indicate that model outputs on filings should not be used without verification.

References

Brier, G. W. 1950. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review*, 78(1): 1–3.

Chen, Z.; Chen, W.; Smiley, C.; Shah, S.; Borova, I.; Langdon, D.; Moussa, R.; Beane, M.; Huang, T.-H.; Routledge, B.; and Wang, W. Y. 2021. FinQA: A Dataset of Numerical Reasoning over Financial Data. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 3697–3711.

Choi, C.; Kwon, J.; Ha, J.; Choi, H.; Kim, C.; Lee, Y.; et al. 2025. FinDER: Financial Dataset for Question Answering and Evaluating Retrieval-Augmented Generation. In *Proceedings of the 6th ACM International Conference on AI in Finance (ICAIF)*. ArXiv:2504.15800.

Geifman, Y.; and El-Yaniv, R. 2017. Selective Classification for Deep Neural Networks. In *Advances in Neural Information Processing Systems*, volume 30.

Guo, C.; Pleiss, G.; Sun, Y.; and Weinberger, K. Q. 2017. On Calibration of Modern Neural Networks. In *Proceedings of the 34th International Conference on Machine Learning (ICML)*, 1321–1330.

Islam, P.; Kannappan, A.; Kiela, D.; Qian, R.; Scherrer, N.; and Vidgen, B. 2023. FinanceBench: A New Benchmark for Financial Question Answering. *arXiv preprint arXiv:2311.11944*.

Ji, Z.; Lee, N.; Frieske, R.; Yu, T.; Su, D.; Xu, Y.; Ishii, E.; Bang, Y. J.; Madotto, A.; and Fung, P. 2023. Survey of Hallucination in Natural Language Generation. *ACM Computing Surveys*, 55(12): 1–38.

Kadavath, S.; Conerly, T.; Askell, A.; Henighan, T.; Drain, D.; Perez, E.; Schiefer, N.; Hatfield-Dodds, Z.; DasSarma, N.; Tran-Johnson, E.; et al. 2022. Language Models (Mostly) Know What They Know. *arXiv preprint arXiv:2207.05221*.

Kalai, A. T.; Nachum, O.; Vempala, S. S.; and Zhang, E. 2025. Why Language Models Hallucinate. *arXiv preprint arXiv:2509.04664*.

Kirichenko, P.; Ibrahim, M.; Chaudhuri, K.; and Bell, S. J. 2025. AbstentionBench: Reasoning LLMs Fail on Unanswerable Questions. *arXiv preprint arXiv:2506.09038*.

Lambert, N.; Morrison, J.; Pyatkin, V.; Huang, S.; Ivison, H.; Brahman, F.; et al. 2024. Tülu 3: Pushing Frontiers in Open Language Model Post-Training. *arXiv preprint arXiv:2411.15124*.

Leng, J.; Huang, C.; Zhu, B.; and Huang, J. 2025. Taming Overconfidence in LLMs: Reward Calibration in RLHF. In *The Thirteenth International Conference on Learning Representations (ICLR)*.

Mehrotra, S.; P, R.; Prajesh, B.; Kambale, H.; and Majumdar, P. 2025. Unmasking Bias in Financial AI: A Robust Framework for Evaluating and Mitigating Hidden Biases in LLMs. In *Proceedings of the 6th ACM International Conference on AI in Finance (ICAIF)*, 951–959.

Naeini, M. P.; Cooper, G.; and Hauskrecht, M. 2015. Obtaining Well Calibrated Probabilities Using Bayesian Binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29.

OpenAI. 2023. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774*.

Qwen Team. 2024. Qwen2.5 Technical Report. *arXiv preprint arXiv:2412.15115*.

Rajpurkar, P.; Jia, R.; and Liang, P. 2018. Know What You Don't Know: Unanswerable Questions for SQuAD. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (ACL)*, 784–789.

Team OLMo; Walsh, P.; Soldaini, L.; Groeneveld, D.; Lo, K.; Arora, S.; et al. 2024. 2 OLMo 2 Furious. *arXiv preprint arXiv:2501.00656*.

Tian, K.; Mitchell, E.; Zhou, A.; Sharma, A.; Rafailov, R.; Yao, H.; Finn, C.; and Manning, C. D. 2023. Just Ask for Calibration: Strategies for Eliciting Calibrated Confidence Scores from Language Models Fine-Tuned with Human Feedback. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 5433–5442.

Wu, S.; Irsoy, O.; Lu, S.; Dabravolski, V.; Dredze, M.; Gehrmann, S.; Kambadur, P.; Rosenberg, D.; and Mann, G. 2023. BloombergGPT: A Large Language Model for Finance. *arXiv preprint arXiv:2303.17564*.

Xiong, M.; Hu, Z.; Lu, X.; Li, Y.; Fu, J.; He, J.; and Hooi, B. 2024. Can LLMs Express Their Uncertainty? An Empirical Evaluation of Confidence Elicitation in LLMs. In *The Twelfth International Conference on Learning Representations (ICLR)*.

Yang, H.; Liu, X.-Y.; and Wang, C. D. 2023. FinGPT: Open-Source Financial Large Language Models. *arXiv preprint arXiv:2306.06031*.

Zhu, F.; Lei, W.; Huang, Y.; Wang, C.; Zhang, S.; Lv, J.; Feng, F.; and Chua, T.-S. 2021. TAT-QA: A Question Answering Benchmark on a Hybrid of Tabular and Textual Content in Finance. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics (ACL)*, 3277–3287.

Appendix A: Prompts

System prompt (all models). "You are a careful financial analyst."

User prompt template.

```
Read the following excerpt from a 10-K
filing and answer the question using only
the excerpt.

Excerpt: "<EXCERPT>"

Question: <QUESTION>
<TASK INSTRUCTION>
```

Task instructions.

- *Numerical:* "Answer with the number exactly as reported in the excerpt (digits only, no explanation). If the excerpt does not state this value, answer NOT_FOUND."
- *Temporal:* "Answer with exactly one word: increased or decreased."
- *Causal:* "Answer with the short phrase from the excerpt that states the driver (no explanation)."

Abstention ablation. In the "w/o" condition the second sentence of the numerical instruction ("If the excerpt does not state this value, answer NOT_FOUND.") is deleted; nothing else changes.

Formatting by model type. OLMo-2 post-trained stages receive the system and user prompt through their chat template. The base model receives the system prompt, a blank line, the user prompt, and a final line "Answer:". Decoding is greedy with at most 24 new tokens. API models receive the user prompt followed by: "Return JSON only: {"answer": <your answer string>, "confidence": <probability 0-1 that your answer is correct>}".

Example questions (v1.1).

- Num-A: "In this excerpt from AAPL's 10-K, what value is reported for 'Rest of Asia Pacific' for fiscal year 2023?"
- Num-U: "In this excerpt from AAPL's 10-K, what value is reported for 'Americas' for fiscal year 2023?" (gold: NOT_FOUND)
- Temporal: "In this excerpt from AAPL's 10-K, did Europe net sales increase or decrease in the most recent fiscal year compared to the prior year?"
- Causal: "In this excerpt from MSFT's 10-K, what was the primary driver of the change in Office Commercial products and cloud services revenue?"

Appendix B: Grader Details

Only the first non-empty line of a generation (or the `answer` field for API models) is graded. Numeric parsing removes commas and currency symbols and reads the first number together with an optional scale word (thousand, million, billion, ...). A parsed value $\hat{v}$ is correct if $|\hat{v} - 10^{3k}v| \leq 0.005\,|10^{3k}v|$ for some $k \in \{-2, \ldots, 2\}$, where v is the gold value. An answer is an abstention if it contains `NOT_FOUND`; a phrase such as "not stated/found/reported/provided/disclosed" or "does not state/report/include"; "N/A"; "unknown"; or "cannot be determined"; or if it is empty or consists only of dashes. Temporal answers are graded on the first occurrence of "increase" or "decrease". Causal answers are correct if the token-level F1 against the gold span (lower-cased alphanumeric tokens) is at least 0.5. The complete grader, all prompts, all raw generations, and the per-item audit decisions are included in the supplementary material.