

# Read Your Own Mind: Reasoning Helps Surface Self-Confidence Signals in LLMs

Jakub Podolak

University of Amsterdam

jakub.podolak.241 [at] gmail.com

Rajeev Verma

University of Amsterdam

## Abstract

We study the source of uncertainty in DeepSeek R1-32B by analyzing its self-reported verbal confidence on question answering (QA) tasks. In the default answer-then-confidence setting, the model is regularly over-confident, whereas semantic entropy—obtained by sampling many responses—remains reliable. We hypothesize that this is because of semantic entropy’s larger test-time compute, which lets us explore the model’s predictive distribution. We show that granting DeepSeek the budget to explore its distribution by forcing a long chain-of-thought before the final answer greatly improves its verbal score effectiveness, even on simple fact-retrieval questions that normally require no reasoning. Our analysis concludes that reliable uncertainty estimation requires explicit exploration of the generative space, and self-reported confidence is trustworthy only after such exploration.

## 1 Introduction

Generative language models (GLMs) like GPT, LLaMA, or Deepseek families have achieved great performance on diverse tasks (Dubey et al., 2024; DeepSeek-AI et al., 2025), yet they are prone to failure modes such as “hallucinations” (Huang et al., 2023). These inaccuracies can undermine trust and lead to poor decisions in LLM-assisted systems (Huang et al., 2024). To mitigate this issue, quantification and the communication of model’s uncertainty in its outputs is seen as a potential to entrust these models with reliability.

Numerous uncertainty-quantification (UQ) approaches have been proposed in this direction: from Monte-Carlo sampling based, such as Semantic Entropy (SE) (Farquhar et al., 2024) to simpler Verbalized Confidence estimation (VC), which just asks the model directly to state its confidence (Xiong et al., 2024).

While Verbalized Confidence estimation is easy to use, there is no scientific consensus on what it

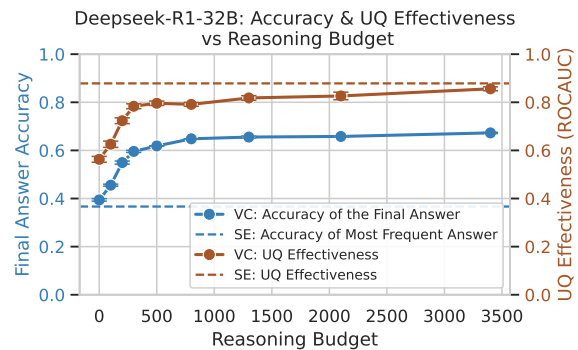


Figure 1: **DeepSeek R1-32B’s Verbalized Confidence (VC) improves and matches Semantic Entropy’s (SE) effectiveness, when longer reasoning is forced.** Our work suggests that it is the test-time exploration of the model’s predictive space, not the particular uncertainty heuristic, that makes confidence estimates reliable.

represents or what its source is, potentially making it unreliable to use in critical scenarios. Furthermore, prior work has shown that Verbalized Confidence is often pathological - for the same question a model might first return “Answer A (100 % confidence)” and, in a second sample, “Answer B (95 % confidence),” even though these probabilities cannot coexist (Xiong et al., 2024). On a more practical side, verbalized scores may be over-confident (Yang et al., 2024; Pawitan and Holmes, 2024), whereas Semantic Entropy remains comparatively well calibrated (Farquhar et al., 2024).

Semantic Entropy’s effectiveness can be attributed to its test-time compute—allocating extra tokens at inference (Snell et al., 2024) to explore the predictive distribution. Test-time compute can also come in the form of an extended reasoning chain that precedes the final answer (Wei et al., 2023; DeepSeek-AI et al., 2025), and most recent works show that such reasoning can improve verbalized or token-level calibration as well (Zeng et al., 2025; Jurayj et al., 2025). These findings prompted us to pose a hypothesis: Is the model

able to directly quantify and express verbally its uncertainty, or is the test-time token exploration necessary for the model to reliably summarize its confidence?

In this paper, we try to answer this question and better understand the source of VC by performing a set of experiments with DeepSeek R1-32B (DeepSeek-AI et al., 2025) as a representative model. Our results show that without any chain-of-thought, DeepSeek’s verbalized scores carry little information about correctness. As we grant the model progressively larger reasoning budgets, its calibration improves and approaches the reliability of Semantic Entropy, even on simple fact-retrieval items. This trend suggests that meaningful uncertainty estimates emerge only after the model’s predictive space has been explored, and that the final confidence percentage largely summarizes the diversity exposed in this process. We further enforce this hypothesis by using a separate reader model that, by just analyzing DeepSeek’s reasoning trace, matches the reliability of DeepSeek’s own Verbalized Confidence.

## 2 Background and Related Work

Generative language models frequently generate fluent but incorrect answers that can cause downstream harm (Band et al., 2024; Huang et al., 2024). When no external verifier is available, a model’s self-reported confidence is the only proxy for correctness, making reliable uncertainty estimates essential.

**Calibration of LLM Confidence Scores.** A confidence score is *calibrated* if, for example, predictions tagged “80 % confident” are correct roughly 80 % of the time. Common approaches to obtain the confidence scores include token-level probabilities treated as a classification score (Dhuliawala et al., 2022), semantic-level measures that evaluate agreement across multiple sampled completions (Farquhar et al., 2024), and explicitly verbalized percentages in a model’s output (Xiong et al., 2024; Tian et al., 2023).

### Semantic Entropy vs. Verbalized Confidence

A generative model, given a question  $Q$ , defines a distribution over semantically distinct answers  $P(A | Q)$ . The uncertainty of this distribution is naturally quantified by its Shannon entropy, and while computing it exactly is infeasible, we can approximate it by Monte-Carlo sampling and clus-

tering semantically equivalent answers. This is exactly how the Semantic Entropy (SE) method (Farquhar et al., 2024) works, leading to well-calibrated scores. The big downside of this method is that it requires sampling data on test-time (larger test-time-compute budget).

Another way to obtain a confidence score is simply to ask the model for one, for instance, “I am 85 % sure.” This Verbalized Confidence (VC) is easy to collect and works with any black-box API (Xiong et al., 2024; Yang et al., 2024; Ni et al., 2024). Yet, opposed to SE, it is unclear what the number represents: is the model sampling its own distribution, recalling similar training examples, or just guessing? To our best knowledge, no study has answered these questions, leaving the method too uncertain for safety-critical use.

The most recent works find that reasoning-tuned models that generate more tokens at the test time give better calibrated verbalized score (Hammoud et al., 2025; Wei et al., 2024; Xiong et al., 2024; Zhao et al., 2024). This hints that exploring test-time compute budgets’ impact on Verbalized Confidence calibration might be crucial to understand its source, yet we are not familiar with any research work that tries to answer our questions specifically.

In this work, we systematically compare Verbalized Confidence and Semantic Entropy under matched test-time compute budgets, examine several task domains, and analyze the reasoning trace to see where the verbalized score comes from and why it lags behind Semantic Entropy.

## 3 Methodology

Our objective is to uncover where a model’s Verbalized Confidence comes from. We identify two competing views:

**Intrinsic latent variable:** the model can read out a hidden latent belief state and use it to express its uncertainty, and **Self-sampling:** model does not have access to any reliable latent source of confidence, and reliable confidence emerges only after the model explicitly explores its own predictive space, as Semantic Entropy does by sampling many answers.

We test these views through a set of experiments that measure the behavior of VC when the model is forced to reason before answering, compare the effectiveness and accuracy to the SE, and analyze the uncertainty exposed in the reasoning traces. We describe our experimental setup in Appendix A.

## 4 Results

Without any reasoning tokens, the score is barely better than random, and with enough exploration budget, VC can approach SE’s effectiveness (subsection 4.1). Furthermore, an external reader can recover essentially the same uncertainty signal by inspecting the chain of thought alone (subsection 4.2), suggesting that the Self-sampling hypothesis might be true.

### 4.1 Extended Reasoning is Necessary for VC to Reach SE-Level Effectiveness

Figure 3 shows final-answer accuracy, UQ effectiveness, and average stated confidence for correct and incorrect answers as a function of the reasoning budget, with Semantic Entropy shown for comparison. We see that granting just 100–500 reasoning tokens raises accuracy 41% → 63% and boosts verbalized-confidence ROC-AUC 0.56 → 0.80.

For fact-retrieval questions (Fig. 3b), answer accuracy does not improve with longer reasoning budgets, yet UQ effectiveness continues to improve with additional tokens. We can reach very long reasoning traces for fact-retrieval questions thanks to the employed forced reasoning technique (Muenighoff et al., 2025) presented in Figure 4.

Verbalized Confidence is initially weaker than superior Semantic Entropy but reaches near parity at 200 tokens for fact retrieval and 3,500 tokens for mathematical items, while maintaining higher answer accuracy due to the reasoning process. For comparison, in our experiments, SE used 218 tokens per sample on average, meaning the two methods are very similar both in computational efficiency and UQ effectiveness for fact retrieval questions.

These results confirm that allocating test-time compute to reasoning is essential for reliable uncertainty estimates, and extended CoT effectively mitigates DeepSeek’s over-confidence without sacrificing performance. The sheer scale of the improvement in effectiveness: from near-random 0.56 ROCAUC to 0.88 suggests that there is no latent uncertainty information available for the model, and self-sampling is necessary to obtain a good uncertainty estimate.

### 4.2 External Reader Model Recovers VC Calibration from Reasoning Trace Alone

If there is no hidden latent variable from which Verbalized Confidence is drawn, then the reasoning

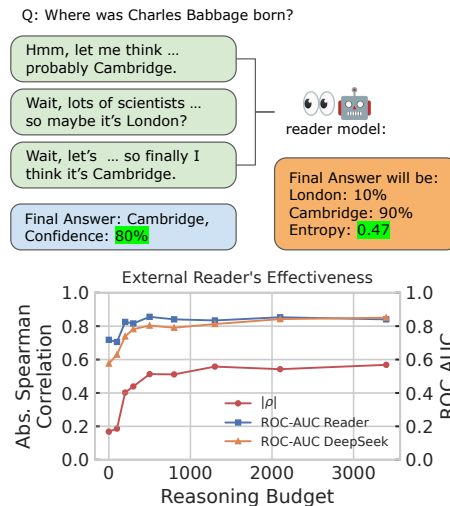


Figure 2: **Separate reader matches the reliability of DeepSeek’s own Verbalized Confidence by just looking at the reasoning trace.** With more reasoning tokens, the agreement between them (measured as absolute Spearman correlation) increases, and the effectiveness of both scores changes similarly.

trace has to contain all the uncertainty information needed to explain Deepseek’s final score. We can verify it using an external reader model that, given only DeepSeek’s chain of thought, tries to predict its final answer and confidence.

Figure 2 illustrates our experimental setup and results. As a reader model, we used OpenAI’s GPT-4o-mini (OpenAI, 2024), we provide more information about the setup in Appendix C. We display (i) the absolute Spearman correlation  $|\rho|$  between DeepSeek’s self-reported confidence and the reader entropy  $H_{\text{reader}}$ , and (ii) the ROC-AUC of each score in detecting incorrect answers.

With no reasoning tokens exposed, the correlation between Reader’s and Deepseek’s scores is low, however, with more reasoning tokens, the effectiveness of the reader goes up in tandem with Deepseek’s effectiveness, and the correlation between the two goes up. At 3.4 k tokens, DeepSeek reaches ROC-AUC = 0.851 and the reader 0.841 with  $|\rho| = 0.57$ , indicating that almost the entire confidence signal is now accessible in the trace.

These results support our claim that there is no directly accessible notion of uncertainty, and uncertainty information must be surfaced through test-time token sampling. When the model provides a Verbalized Score after the reasoning process, it most likely just reads its reasoning trace and summarizes the alternatives and uncertainty exposed in it.

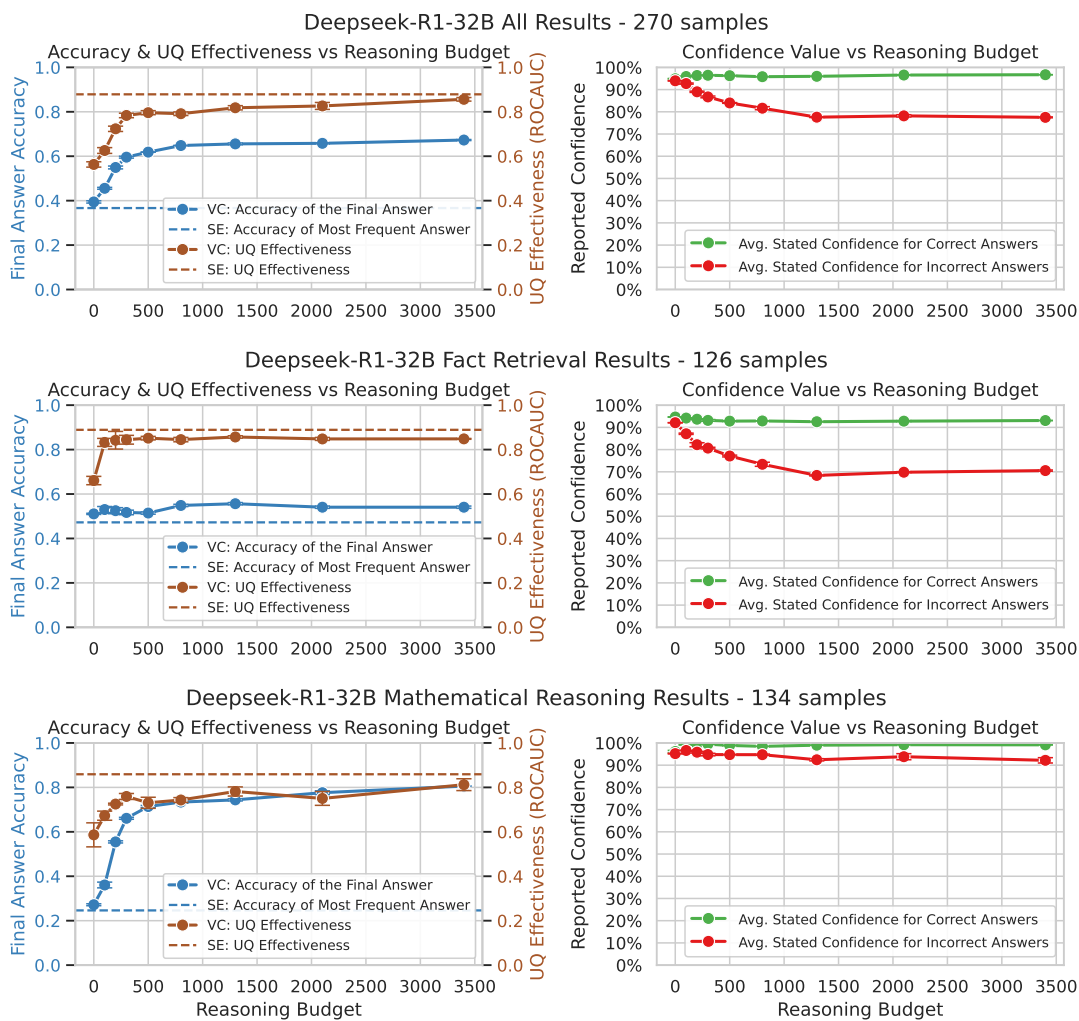


Figure 3: **Effectiveness and Accuracy of Verbalized Confidence with Forced Reasoning vs Semantic Entropy.** (a) Full overview. (b) Fact retrieval results. (c) Mathematical reasoning results. Note: The remaining 10 samples not falling into the Fact Retrieval or Mathematical Reasoning categories are included in the Full overview but not presented as separate plots.

## 5 Discussion and Future Work

We aimed to determine whether large language models can directly verbalize well-calibrated uncertainty or whether reliable confidence estimates only emerge after explicit exploration of their predictive space, via additional test-time compute such as parallel sampling (Semantic Entropy) or extended reasoning.

Our experiments suggest that *test-time compute, not the particular uncertainty heuristic, is the decisive factor for obtaining reliable confidence estimates in DeepSeek-R1-32B*. Left to produce only a short answer, the model remains over-confident because its belief state cannot be accessed directly. Granting the model additional tokens, either by sampling independent continuations (Semantic Entropy) or by forcing a longer chain of thought, al-

lows it to externalize alternative hypotheses. This exposes a big issue with Verbalized Confidence - its appeal lies in the simplicity and how fast it is, yet it works well only after a significant reasoning computation is done.

While these results are encouraging, they may not be generalizable since we’ve tested only one model and used a very compact QA dataset. Furthermore, assuming it is true that test-time compute is the decisive factor for reliable UQ, it still might be the case that some methods of test-time compute may be more efficient in eliciting uncertainty than others. Future work could focus on making models reason more efficiently or explore their uncertainty in a more structured way. That could help Verbalized Confidence inherit the Semantic-Entropy-level of calibration with less computation needed.



## References

- AIME. 2024. Aime 2024 dataset. [https://huggingface.co/datasets/Maxwell-Jia/AIME\\_2024](https://huggingface.co/datasets/Maxwell-Jia/AIME_2024). Accessed: 2025-05-07.
- Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. 2024. [Linguistic calibration of long-form generations](#). *Preprint*, arXiv:2404.00474.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, and 181 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *Preprint*, arXiv:2501.12948.
- Shehzaad Dhuliawala, Leonard Adolphs, Rajarshi Das, and Mrinmaya Sachan. 2022. [Calibration of machine reading systems at scale](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 1682–1693, Dublin, Ireland. Association for Computational Linguistics.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, and 1 others. 2024. [The llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- S. Farquhar, J. Kossen, L. Kuhn, and Y. Gal. 2024. [Detecting hallucinations in large language models using semantic entropy](#). *Nature*, 630(8017):625–630.
- Hasan Abed Al Kader Hammoud, Hani Itani, and Bernard Ghanem. 2025. [Beyond the last answer: Your reasoning trace uncovers more than you think](#). *Preprint*, arXiv:2504.20708.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). *Preprint*, arXiv:2009.03300.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2023. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *Preprint*, arXiv:2311.05232.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, and Ting Liu. 2024. [A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions](#). *ACM Transactions on Information Systems*.
- Mandar Joshi, Eunsol Choi, Daniel Weld, and Luke Zettlemoyer. 2017. [triviaqa: A Large Scale Distantly Supervised Challenge Dataset for Reading Comprehension](#). *arXiv e-prints*, arXiv:1705.03551.
- William Jurayj, Jeffrey Cheng, and Benjamin Van Durme. 2025. [Is that your final answer? test-time scaling improves selective question answering](#). *Preprint*, arXiv:2502.13962.
- Niklas Muennighoff, Zitong Yang, Weijia Shi, Xiang Lisa Li, Li Fei-Fei, Hannaneh Hajishirzi, Luke Zettlemoyer, Percy Liang, Emmanuel Candès, and Tatsunori Hashimoto. 2025. [s1: Simple test-time scaling](#). *Preprint*, arXiv:2501.19393.
- Shiyu Ni, Keping Bi, Lulu Yu, and Jiafeng Guo. 2024. [Are large language models more honest in their probabilistic or verbalized confidence?](#) *Preprint*, arXiv:2408.09773.
- OpenAI. 2024. [Hello, gpt-4o](#). Accessed: 2025-01-05.
- Yudi Pawitan and Chris Holmes. 2024. [Confidence in the reasoning of large language models](#). *Preprint*, arXiv:2412.15296.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. 2024. [Scaling llm test-time compute optimally can be more effective than scaling model parameters](#). *Preprint*, arXiv:2408.03314.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher D. Manning. 2023. [Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback](#). *Preprint*, arXiv:2305.14975.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. 2023. [Chain-of-thought prompting elicits reasoning in large language models](#). *Preprint*, arXiv:2201.11903.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2024. [Can llms express their uncertainty? an empirical evaluation of confidence elicitation in llms](#). *Preprint*, arXiv:2306.13063.
- Daniel Yang, Yao-Hung Hubert Tsai, and Makoto Yamada. 2024. [On verbalized confidence scores for llms](#). *Preprint*, arXiv:2412.14737.
- Qingcheng Zeng, Weihao Xuan, Leyang Cui, and Rob Voigt. 2025. [Do reasoning models show better verbalized calibration?](#) *Preprint*, arXiv:2504.06564.

Xinran Zhao, Hongming Zhang, Xiaoman Pan, Wenlin Yao, Dong Yu, Tongshuang Wu, and Jianshu Chen. 2024. [Fact-and-reflection \(FaR\) improves confidence calibration of large language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 8702–8718, Bangkok, Thailand. Association for Computational Linguistics.

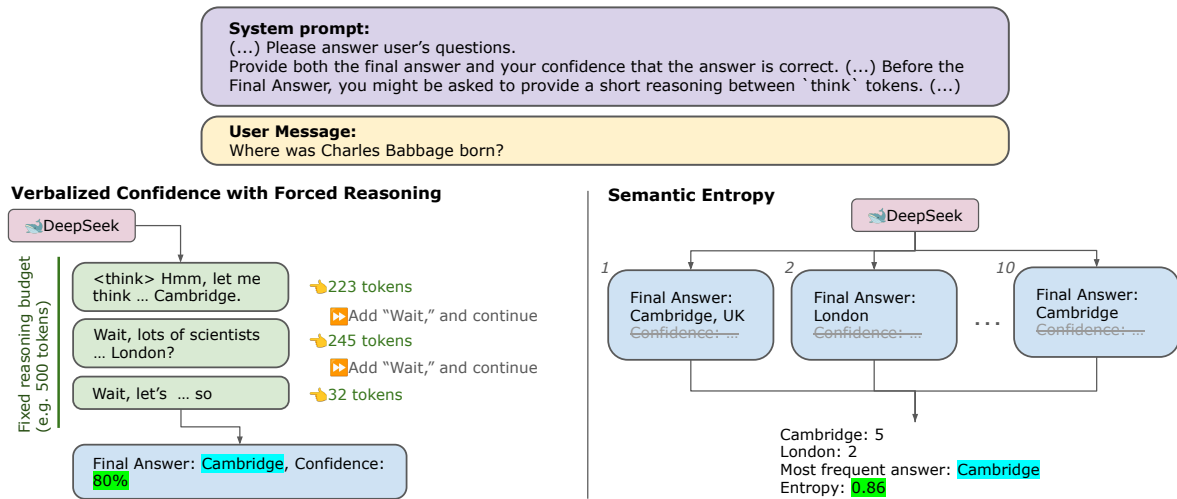


Figure 4: **Two tested methods of obtaining Final Answer and Confidence** - Verbalized Confidence with Forced Reasoning (VC) works by prompting the model to reason for longer-until the fixed budget is exhausted - before stating the answer and confidence. Semantic Entropy (SE) obtains 10 independent answers that are later clustered semantically to identify the most frequent one, and to calculate the entropy in the answer distribution.

## A Experimental Setup

**Data Sources.** Because long-trace experiments are computationally expensive, we built a small (270 samples) but diverse benchmark instead of using full datasets. We sampled questions from five popular, open-source sources: TriviaQA, MMLU, and SimpleQA for fact retrieval (Joshi et al., 2017; Hendrycks et al., 2021; Wei et al., 2024), plus GSM8K and AIME-2024 for mathematical reasoning (Cobbe et al., 2021; AIME, 2024). Our goal is open-ended QA in natural language, so we stripped away multiple-choice options in MMLU and any figure references in AIME-2024, manually discarding questions that could not stand alone after this edit, such as “Which of the following is true?”. Every surviving example was then hand-labeled with its knowledge domain and the skills needed to answer it, such as “Fact Retrieval” or “Mathematical Reasoning”. Full sampling details and the final label distribution appear in Appendix B.

**Model and Prompts.** We chose Deepseek-R1-32B<sup>1</sup> (DeepSeek-AI et al., 2025) following Jurayj et al. (2025) for its strong reasoning capabilities at a manageable model size. Furthermore, it is one of the most popular open-sourced reasoning-tuned models. All experiments were run on two NVIDIA A100 GPUs.

We provide an elaborate discussion on prompting and inference we adopted in Figure 4. Across all setups, we used a single system prompt that directs the model to (1) think step by step, and then (2) provide a final answer along with a Verbalized Confidence score. The full prompt text, as well as an interaction example, is available in Appendix C. To regulate the length of the reasoning chain, we applied the budget-based truncation method of Muennighoff et al. (2025): when the reasoning budget is exhausted (or set to zero), the chain terminates immediately. If the budget remains, the system appends “Wait, ” tokens, and asks to generate more tokens. For experiments with Verbalized Confidence, we lowered the decoding temperature to 0.1 to prevent the model from going off-topic in long reasoning. For parallel sampling in Semantic Entropy experiments, we set it to 1.0 to obtain more diverse responses and approximate the predictive distribution more efficiently.

**UQ Methods.** Next, we describe how we obtain the estimates of verbalized score and semantic entropy: For Verbalized Confidence, we ask the model to provide the final answer and its confidence between 0% and 100% after (optional) forced reasoning (refer to Figure 4 for visualization).

For Semantic Entropy, we follow Farquhar et al. (2024), and generate  $n = 10$  answers for each question with no reasoning chain. Afterwards, we use OpenAI’s gpt-4o-mini (OpenAI, 2024) to cluster semantically equivalent generations. We select the majority cluster – the cluster with the most members (i.e., the answer

<sup>1</sup>deepseek-ai/DeepSeek-R1-Distill-Qwen-32B

that appears most frequently once semantically equivalent responses are grouped) – as the predicted answer and compute the Shannon entropy of the cluster-size distribution as the uncertainty score.

**Evaluation and Metrics.** We report two main metrics: the accuracy of the final answer and the effectiveness of the UQ method measured as the area under ROC (ROC AUC) in the task of classifying the model’s final answer correctness (hallucination classification). To calculate if the model’s final answer is correct, we query OpenAI’s GPT-4o-mini (OpenAI, 2024) model with the question if this proposed answer is equivalent to the ground truth answer in the dataset given a question.

We repeat experiments with verbalized score confidence across varying reasoning budgets 3 times and show mean and 95% confidence intervals. For the rest of the experiments, we repeat them once unless stated otherwise.

## B Data Creation and Composition

Because our longest-trace runs are expensive, we limited the benchmark to **270 open-ended questions** drawn from five well-known, permissively licensed QA datasets. We first sampled 310 items uniformly at random (seed 42) to balance fact-retrieval and mathematical-reasoning content while keeping the total below the  $\approx 300$ -sample budget we could process. Items whose solutions required figures (AIME-2024), multiple-choice candidates (MMLU), or extra context passages (TriviaQA) were discarded after manual inspection, leaving the 270 used in all experiments (Table ??). By doing so, we ensured that all the incorrect answers were caused by the model’s mistakes, instead of missing context in the data.

Each example received two human labels - Knowledge Domain, and Skill Required. A large language model (OpenAI o3) proposed initial tags for 100 random questions. The first author then reviewed every instance, correcting tags where needed, and used these tags to manually label all 270 samples. You can find the specific tags and number of datapoints in Figure 5. In the main paper, we break out results for the full dataset and for the two most common *Skill Required* tags only; *Knowledge Domain* splits are omitted because several categories are too small. Per-dataset results can be found in Appendix D. Five representative questions and their tags are shown in Table 1.

Example Question (truncated)	Dataset	Skill	Domain
In what year did Augustus De Morgan publish the article "Trochoidal Curve" in the Penny Cyclopaedia?	SimpleQA	Fact Retrieval	History and Past Events
There exist real numbers $x$ and $y$ , both greater than 1, such that $\log_x(y^x) = \log_y(x^{4y}) = 10$ . Find $xy$ .	AIME2024	Mathematical Reasoning	Mathematics
James runs 12 miles a day for 5 days a week. If he runs 10 miles an hour how many hours does he run a week?	GSM8K	Mathematical Reasoning	Mathematics
In Python 3, which of the following function removes all leading and trailing whitespace in string?	MMLU	Fact Retrieval	IT and Engineering
Anaphylaxis is what sort of life-threatening illness?	TriviaQA	Fact Retrieval	Science, Nature and Medicine

Table 1: Five representative items from the 270-question benchmark.

## C Prompts and Inference

**Main VC vs SE runs.** You can find the full system prompt used in the main VC vs SE experiments, as well as a real interaction example with Verbalized Confidence and budget forcing in Figure 6.

**Reader-model experiments.** Our goal is to let an external model read DeepSeek’s reasoning trace and predict a probability distribution over possible DeepSeek’s answers. After obtaining the distribution,



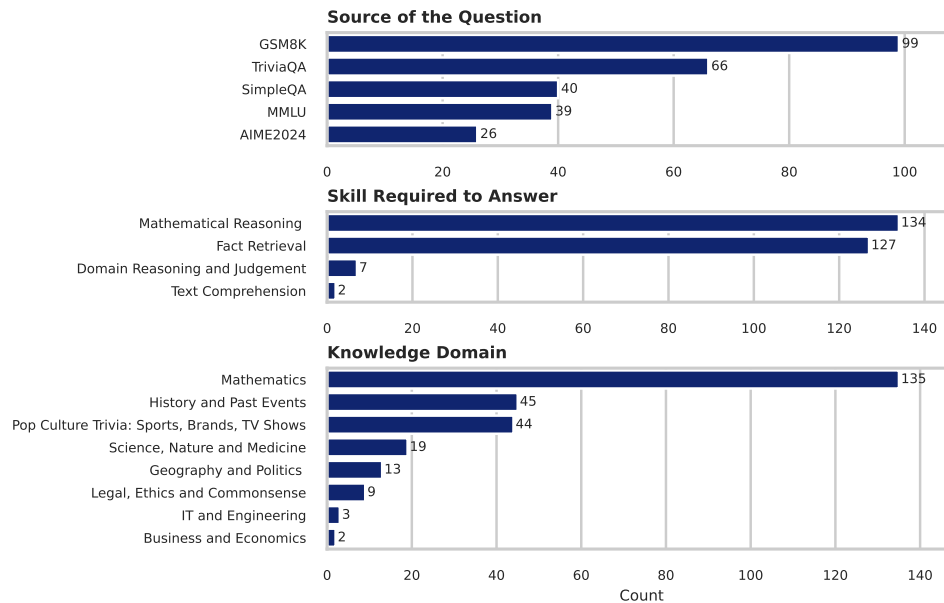


Figure 5: Internal composition of our used data sample.

**System prompt:**  
 You are a trivia expert, please answer user questions. Provide both the final answer and your confidence that the answer is correct. Be honest and precise in your confidence estimation, as the users will rely on this information to make decisions. Try not to overestimate your confidence. Answer in the format: `Final Answer: <final\_answer>`. Confidence: <confidence>` where <confidence> is a percentage between 0% and 100%, and <final\_answer> is just the final answer in one brief sentence.

Before the Final Answer, you might be asked to provide a short reasoning between `think` tokens. In your reasoning process, try to constantly reevaluate your confidence, by considering the information you have and the information you are missing or may have calculated wrong. Reason not only about the final answer, but also about your confidence.

**User Message:**  
 The Fitzwilliam Museum is in which British town or city?

**Model response:**  
 <think> Okay, so I need to figure out where the Fitzwilliam Museum is located. I remember hearing about it before, but I'm not exactly sure. Let me think. I know there are several famous museums in the UK, like the British Museum in London and the Victoria and Albert Museum, also in London. But the Fitzwilliam Museum doesn't ring a bell as being in London. Maybe it's in another city.\n\nI think Cambridge is known for its universities, especially the University of Cambridge. There are a lot of museums there associated with the university. I recall that the Fitzwilliam Museum is one of them. It's a major art museum, right? They have a lot of paintings, sculptures, and other artifacts. I'm pretty sure it's in Cambridge because I've seen pictures of it before, and it's often mentioned alongside other Cambridge landmarks.\n\nWait, could it be in Oxford? No, Oxford has the Ashmolean Museum, which is also a big one. So, no, the Fitzwilliam is definitely in Cambridge. I'm pretty confident about that. I don't think it's in any other city like Manchester or Birmingham. It's specifically linked to the University of Cambridge, so that must be the location.\n\nI don't have any conflicting information, so I think I'm safe to say it's in Cambridge. My confidence is high because I've heard it mentioned in the context of Cambridge and its university.\n\nWait, just to double-check, I</think>  
 Final Answer: Cambridge. Confidence: 95%

Figure 6: Used system prompt and real interaction example. The question comes from the TriviaQA dataset. Moments where the model was forced to reason for longer are highlighted.

Dataset	Sampled	After Filtering	Removed
GSM8K	100	99	1
TriviaQA	70	66	4
SimpleQA	40	40	0
MMLU	70	39	31
AIME2024	30	26	4
<b>Total</b>	<b>310</b>	<b>270</b>	<b>40</b>

Table 2: **Number of samples per dataset before and after manual filtering.**

we calculate Shannon entropy, which is used as a notion of uncertainty. The procedure of obtaining a distribution over possible DeepSeek’s answers has four steps:

1. **Candidate extraction.** For each question, we feed the entire 3.4k-token reasoning chain to gpt-4o-mini, prompting it to list all candidate answers mentioned in the trace.
2. **Multiple-choice reformulation.** We label the distinct candidates with letters A, B, C..., we also add an option "None Unknown".
3. **One-token completion.** We prompt gpt-4o-mini again, asking which of the candidates is most likely to be returned by DeepSeek as a final answer. We ask the API for exactly one letter A, B, C... as a response. We keep the logits of this single-token prediction, apply softmax, and obtain a categorical distribution  $p(A), p(B), \dots$ .
4. **Reader entropy.** The Shannon entropy of that distribution,  $H_{\text{reader}} = -\sum_i p(i) \log p(i)$ , is the reader’s uncertainty score used in our analyses.

You can find prompts that we used to extract the list of all candidates, as well as the final prediction, in figure [Figure 7](#).

## D Results Across Datasets

We provide additional comparison of VC and SE with varying reasoning budgets across different sources of data in [Figure 8](#). While the trends are much noisier because of fewer samples, we see that for all the sources the difference in confidence between correct and incorrect answers increases with more reasoning tokens.

### Possible-candidates prompt

“You are a helpful assistant. We asked a person to answer an open-ended question. The person wrote a reasoning trace and then gave a final answer. I want to know what all the possible options were that the person considered before giving the final answer. This is for research on the answers people consider. I want to use these possible answers to create a multiple-choice question. Please give me all the possible answers the person considered, without duplicates, keeping the answers distinct and suitable for a multiple-choice question. Respond as a JSON list of strings, like this: FINAL LIST: ["answer1", "answer2", ...], which I can later use in a quiz. Add the correct answer, the final answer, and Other / Unknown to the list, making sure there are no duplicates. Reason about possible answers and how to remove duplicates, before giving the final list. Include only answers mentioned in the reasoning trace plus the correct answer. Do not add answers that are missing from the trace or the correct / unknown items.”

### Final-prediction prompt

“You are a helpful assistant. We asked a person to answer a question; the person started reasoning about the possible answers. Your role is to serve as an autocompletion model that predicts the most likely final answer the person will give, based on the current reasoning trace. You are not allowed to give your own answer—only the person’s most likely final answer. Choose from the list of possible answers (A, B, C, ...) and output exactly one letter, with no whitespace or comments.”

+ A) {Option A}  
+ B) {Option B}  
+ C) {Option C}  
...

Figure 7: **Reader-model prompts.** First prompt is used to extract candidate answers from the reasoning trace, the second one shows the follow-up prompt that asks the reader to predict DeepSeek’s final choice. The last block of the second prompt is filled dynamically with the candidate list produced by the first one.

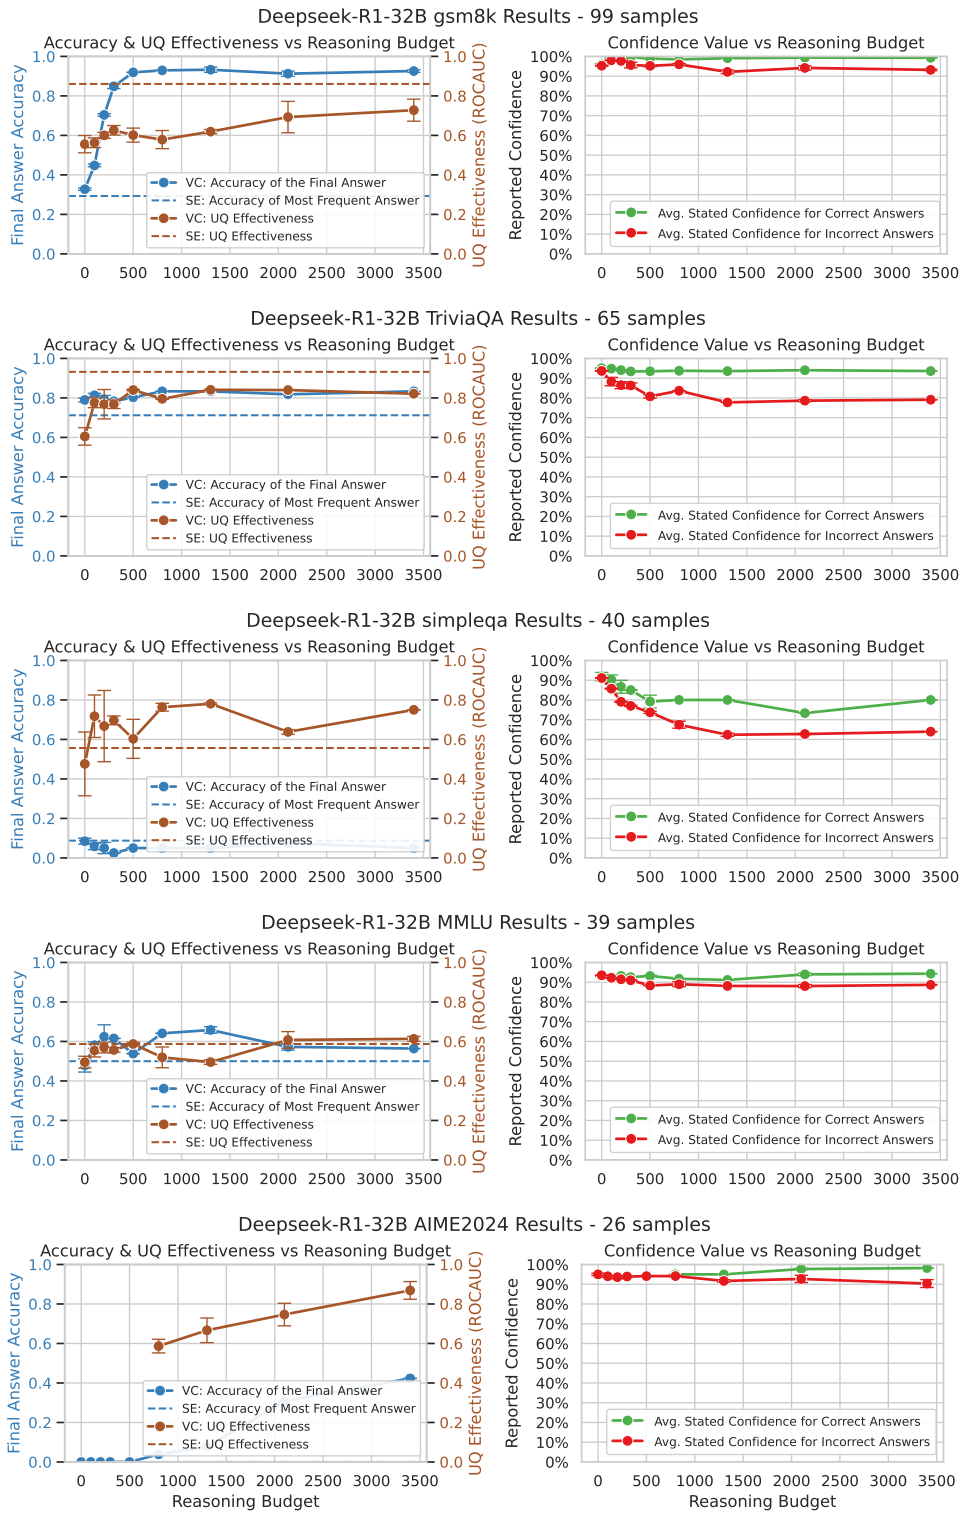


Figure 8: **Effectiveness and Accuracy of Verbalized Confidence with Forced Reasoning vs Semantic Entropy.** Despite noise from limited samples, the right-hand plots show a consistent and increasingly pronounced divergence in reported confidence between correct and incorrect answers as the reasoning budget increases.