

# Understanding the Impact of Confidence in Retrieval Augmented Generation: A Case Study in the Medical Domain

Shintaro Ozaki<sup>♣</sup> Yuta Kato<sup>♣</sup> Siyuan Feng<sup>♣</sup> Masayo Tomita<sup>♣</sup> Kazuki Hayashi<sup>♣</sup>  
Wataru Hashimoto<sup>♣</sup> Ryoma Obara<sup>◇</sup> Masafumi Oyamada<sup>◇</sup>  
Katsuhiko Hayashi<sup>♣</sup> Hidetaka Kamigaito<sup>♣</sup> Taro Watanabe<sup>♣</sup>

<sup>♣</sup>Nara Institute of Science and Technology (NAIST)

<sup>♣</sup>The University of Tokyo <sup>◇</sup>NEC Corporation

{ozaki.shintaro.ou6, kamigaito.h, taro.watanabe}@naist.ac.jp

## Abstract

Retrieval Augmented Generation (RAG) complements the knowledge of Large Language Models (LLMs) by leveraging external information to enhance response accuracy for queries. This approach is widely applied in several fields by taking its advantage of injecting the most up-to-date information, and researchers are focusing on understanding and improving this aspect to unlock the full potential of RAG in such high-stakes applications. However, despite the potential of RAG to address these needs, the mechanisms behind the confidence levels of its outputs remain underexplored. Our study focuses on the impact of RAG, specifically examining whether RAG improves the confidence of LLM outputs in the medical domain. We conduct this analysis across various configurations and models. We evaluate confidence by treating the model's predicted probability as its output and calculating several evaluation metrics which include calibration error method, entropy, the best probability, and accuracy. Experimental results across multiple datasets confirmed that certain models possess the capability to judge for themselves whether an inserted document relates to the correct answer. These results suggest that evaluating models based on their output probabilities determine whether they function as generators in the RAG framework. Our approach allows us to evaluate whether the models handle retrieved documents.<sup>1</sup>

## 1 Introduction

Retrieval Augmented Generation (RAG) (Lewis et al., 2020) serves as a method to not only mitigate hallucinations but also supplement the knowledge of Large Language Models (LLMs) (Achiam et al., 2023; Dubey et al., 2024; Aizawa et al., 2024). By leveraging external information, RAG enhances response accuracy and alignment with queries, mak-

<sup>1</sup>The code is available at [https://github.com/naist-nlp/CC\\_RAG](https://github.com/naist-nlp/CC_RAG).

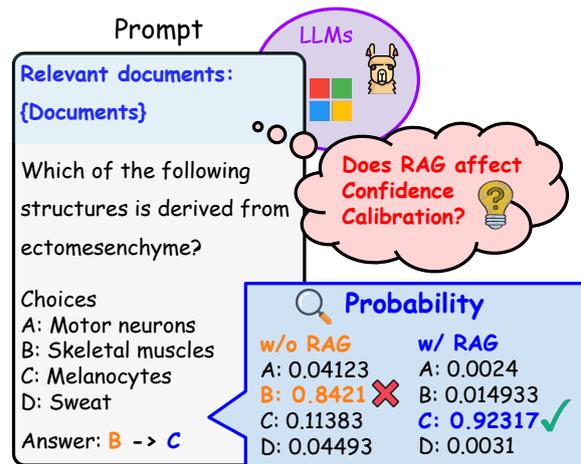


Figure 1: The focus of our research is to analyze whether RAG improves the confidence of the model response.

ing it widely applicable in industries. Notable domains include finance (Yepes et al., 2024; Setty et al., 2024) and healthcare (Xiong et al., 2024), where the reliability of information is critical. This study focuses on the medical domain, which has relatively more text data than other fields and involves complex factors directly related to the human body. (Sohn et al., 2024; Jeong et al., 2024)

While researchers explore performance improvements for LLMs using RAG, as illustrated in Figure 1, analyses focusing on prediction confidence remain limited. Although RAG enhances answer accuracy, it may lead to overconfidence, where models exhibit unwarranted self-assurance (Chen et al., 2024). We hypothesize that retrieving documents to support the correct answer through RAG improve the model's confidence, potentially leading to errors in confidence calibration. Based on this, we pose a research question: Do LLMs improve the confidence for outputs with RAG?

To address this question, we conduct a systematic analysis of multiple tasks and models in the medical domain, exploring diverse scenarios using

Dataset	Size	Option
PubMedQA	1,000	3
MedMCQA (Extract)	2,206	4

Table 1: The dataset used in our study. We select datasets that not only contain QA pairs but also include explanatory passages that justify the answers.

PubMedQA (Jin et al., 2019) and MedMCQA (Pal et al., 2022). In particular, we create pseudo-RAG to manipulate document content – such as adding irrelevant documents deliberately or including only those directly related to the answer – to simulate the range of situations RAG encounter.

Our result shows that inserting documents deliberately containing answer-supporting information improve confidence in many models, aligning with expectations, although some models exhibited behavior contrary to this prediction. Additionally, inserting documents unrelated to the correct answer rarely improve the confidence, suggesting that LLMs can discriminate whether an inserted document relates to the answer. These results indicate that evaluating models based on output probabilities can lead to reveal the suitable generator model.

## 2 Related Work

### 2.1 Confidence of LLMs

Research on confidence has been prevalent since before the era of LLMs (Jiang et al., 2021) and continues to be extensively explored (Geng et al., 2024). Becker and Soatto (2024) proposed a framework that measures confidence by leveraging explanation-generating text produced by LLMs. Zhao et al. (2021) identified the issue that few-shot prompting significantly impacts model confidence and alters its inherent performance, and they proposed methods to address this problem. Confidence estimation is used as a technique to suppress hallucinations, where models generate false information (Zhang et al., 2023). Cole et al. (2023) demonstrated that by utilizing model confidence, it is possible to suppress outputs for ambiguous questions. Our study contributes to this body of research by specifically analyzing how RAG influences confidence calibration in LLM outputs. Unlike prior works that primarily optimize retrieval mechanisms, we directly investigate confidence calibration dynamics.

### 2.2 Boosting RAG with Confidence

Recent advances in RAG have leveraged model confidence (e.g., output probability) to optimize retrieval and generation processes. For instance, Jiang et al. (2023) introduced FLARE, which dynamically decides whether to retrieve additional information based on token-level confidence during generation, ensuring efficient retrieval by minimizing unnecessary searches. Similarly, query rewriting techniques using reinforcement learning (Ma et al., 2023) and strategies such as Recitation-Augmented Generation (Sun et al., 2023), which searches for text resembling hypothetical answers, have shown promise in enhancing retrieval accuracy. Moreover, recent studies like Self-RAG (Asai et al., 2024) integrate retrieval into the generation process itself. In many of these approaches, confidence plays a crucial role either in deciding when to retrieve or in re-ranking retrieved documents based on their relevance. However, these studies focus on improving RAG performance without analyzing how confidence itself is influenced by the RAG. Specifically, while confidence thresholds and re-ranking mechanisms are employed to optimize retrieval and generation, the underlying dynamics of confidence calibration within the RAG pipeline remain underexplored. Our study analyzes confidence calibration with and without RAG to address this gap, verify the implicit assumptions of prior works, and contribute to a deeper understanding of confidence-based mechanisms in RAG.

## 3 Methods

Our study analyzes whether the confidence improves through RAG by calculating the model’s confidence or entropy from the predicted probability by the model. Each input is formatted by concatenating a system prompt, a question prompt, and its answer options (e.g., a four-choice question), following the design of Medical Information Retrieval Augmented Generation Evaluation (MI-RAGE) (Xiong et al., 2024). We also analyze the optimal position for inserting documents retrieved via pseudo-RAG, i.e., inserting a document directly relevant to the answer or irrelevant deliberately into the model input prompt. Specifically, we evaluate three insertion patterns: before the question (Pre-Question, denoted as Pre-Q), between the question and the answer choices (After-Question, denoted as Aft-Q), and after the answer choices (After-Choice, denoted as Aft-C). This setup allows us to exam-

ine the Lost-in-the-Middle phenomenon (Liu et al., 2024), where models tend to overlook intermediate content when processing long-context inputs. Moreover, in order to focus on the impact of retrieved document positions, we use documents that contain the correct answer to the question. We validate our research question under three scenarios: (1): inserting only the explanation related to the answer (denoted as Ans1). (2): combining the correct explanation with two irrelevant documents (denoted as Ans1-Oth2). (3): inserting three irrelevant documents (denoted as Oth3). The irrelevant documents are selected from unrelated questions, ensuring that they do not contain the correct answer or semantically similar content.

Directly generating the choice answer by the model complicates evaluation, because differences in reported metrics arise even under identical conditions across studies (Xiong et al., 2024; Chen et al., 2023; Wu et al., 2024). In some studies, researchers select the final candidate using regular expressions, while in others, they treat the output of a specific word (such as Yes or No) as the correct answer. Thus, evaluation methods are not uniquely defined if the sentence generated. In our study, we predict the most plausible option from the given choices as follows:

$$v_i = \log P(x_i | \text{prompt})$$

$$P(x_i) = \frac{\exp(v_i)}{\sum_{j=1}^J \exp(v_j)}$$

where  $v_i$  represents the log probability corresponding to each choice  $x_i$  and the prompt refers to the provided question or context.  $P(x_i)$  denotes the probability that the choice  $x_i$  is the correct answer, normalized by dividing the exponential of  $v_i$  by the sum of exponentials of all  $v_j$  values, while  $J$  is the number of options, which is 3 or 4.

## 4 Experimental Setup

### 4.1 Datasets

We focus on the application of RAG in the medical domain. For the dataset, we select PubMedQA (Jin et al., 2019) and MedMCQA (Pal et al., 2022), both of which include multiple-choices QA data along with explanatory passages that justify the answers. These datasets follow the experimental setup of MIRAGE (Xiong et al., 2024), as shown in Table 1. For MedMCQA, we extract only the questions that include supporting evidence for the answer, resulting in a total of 2,206 instances.

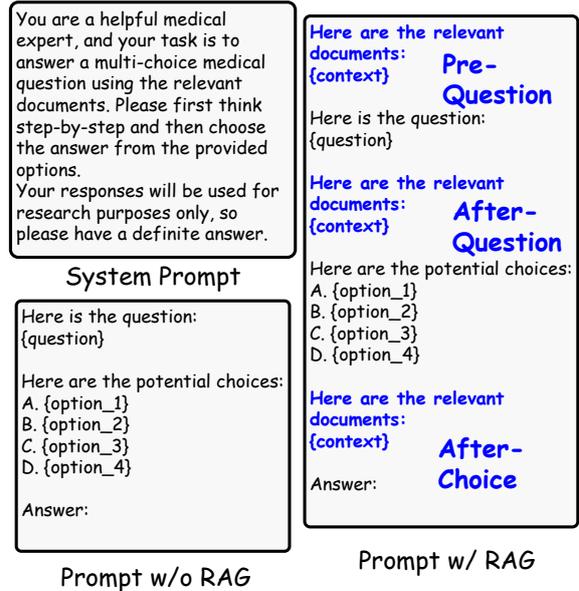


Figure 2: Prompts used in our research. Each prompt begins with a concatenated of the system prompt. Following MIRAGE (Xiong et al., 2024), we design the templates to enable the calculation of probabilities.

### 4.2 Inference Models

Following prior research (Xiong et al., 2024), we select the following models for evaluation: Phi-3.5 (3.8B) (Abdin et al., 2024), PMC-Llama (13B) (Wu et al., 2024), Llama2 (70B) (Touvron et al., 2023b), LLaMA3.1 (8B / 70B) (Dubey et al., 2024), and Meditron (70B) (Chen et al., 2023). To ensure fair evaluation across models with different architectures and parameter sizes, we also include Gemma2 (2B) (Team et al., 2024) and Qwen2.5 (14B / 70B) (Yang et al., 2024), bringing the total to nine models. PMC-Llama is fine-tuned on medical domain documents based on Llama (Touvron et al., 2023a), while Meditron undergoes continual pretraining on Llama2 (Touvron et al., 2023b). For 70B models, we apply 4-bit quantization, and for PMC-Llama, we use half-precision quantization to compute probabilities. Detailed model configurations are provided in Appendix A.1.

### 4.3 Templates

Our study modifies the approach based on the MIRAGE paper (Xiong et al., 2024) by excluding Chain of Thought (CoT) (Wei et al., 2022), allowing direct probability computation. (In other words, when using CoT, the model must generate responses, which, as discussed in Section 3, prevents a valid evaluation.) Figure 2 presents the prompts used in our study. Each prompt incorpo-

rates system prompts from prior research (Xiong et al., 2024) at the beginning of the input prompt. To investigate whether the Lost-in-the-Middle phenomenon (Liu et al., 2024), also occurs in RAG, our study inserts retrieved documents at three positions: before the question (Pre-Q), after the question (Aft-Q), and after the answer choices (Aft-C).

#### 4.4 Evaluation Metrics

We evaluate if RAG boosts LLM confidence using entropy, best probability, accuracy, and Adaptive Calibration Error. In our multiple-choice QA task, each question has one correct answer, and output probabilities classify responses as correct or not.

**Entropy.** We examine how entropy changes for candidate answer choices under the influence of RAG, calculating an entropy. Ideally, inserting an answer-containing document should decrease entropy (indicating a more confident selection of the correct choice), while inserting entirely unrelated documents should improve entropy. The entropy is computed as:

$$H(P) = - \sum_{i=1}^J P(x_i) \log P(x_i)$$

$$P(x_i) = \frac{\exp(v_i)}{\sum_{j=1}^J \exp(v_j)}$$

Here,  $x_i$  represents a candidate answer among  $J$  total options, and  $v_i$  denotes the logit score (i.e., the unnormalized log-probability) assigned to  $x_i$ . The softmax function transforms these logits into a probability distribution  $P(x_i)$ , from which the entropy  $H(P)$  is calculated. Lower entropy indicates higher model confidence in a particular choice, while higher entropy implies uncertainty.

**Best Probability.** We define ‘‘Best Probability’’ as the highest output probability among the candidate choices given to the model. In our study, we evaluate this metric as confidence. A high output probability shows strong confidence for correct answers, while a low output probability is preferred for incorrect answers (we want irrelevant documents to lower the model’s confidence).

The notation of best probability is as follows:

$$x^* = \arg \max_{x_i \in \mathcal{X}} (\log P(x_i | \text{prompt}))$$

$$P(x_i) = \frac{\exp(v_i)}{\sum_{j=1}^J \exp(v_j)}$$

Settings	Options
QA	PubMedQA, MedMCQA
Model	Gemma2, Phi3.5, Llama2, Llama3.1 Qwen2.5, PMC-Llama, Meditron
Template	w/o RAG, Pre-Q, Aft-Q, Aft-C
Evaluation	Entropy, Best Prob, Accuracy, ACE

Table 2: Experimental settings used in our research.

Here,  $\mathcal{X}$  is the set of all candidate answer choices, and  $x^*$  denotes the choice with the highest log-probability. Each  $v_i$  represents the model’s logit for the candidate  $x_i$ , and the softmax function converts these logits into a probability distribution over all choices. The selected  $x^*$  corresponds to the most confident prediction the model makes under the given prompt. This Best Probability reflects how strongly the model favors its top prediction, and it serves as an interpretable confidence score in our evaluations.

**Adaptive Calibration Error (ACE).** Adaptive Calibration Error (ACE) (Nixon et al., 2019) is a metric proposed to address the shortcomings of Expected Calibration Error (ECE) (Naeini et al., 2015; Guo et al., 2017), specifically aiming to reduce the risk of bins with a small number of samples. Proskurina et al. (2024) and Ulmer et al. (2022) have pointed out that ACE is a more suitable calibration error metric for multi-class classification problems. Based on these findings, we adopt ACE in our evaluation. Table 2 provides a complete list of all combinations and Appendix A.6 the details of evaluation metrics.

## 5 Results

Table 3 presents the experimental results using MedMCQA, while Table 4 shows the results for PubMedQA. When distinguishing between correctly answered and incorrectly answered questions, Phi and Qwen exhibited ideal behavior from an entropy perspective. Specifically, inserting supporting documents for the correct answers led to higher entropy, whereas inserting only unrelated documents resulted in lower entropy. In contrast, other models, e.g., Llama2, Llama3.1, and Gemma2, produced unexpected results, suggesting that Llama and Gemma may struggle to process inserted documents effectively. Furthermore, Qwen and Phi demonstrated the ability to determine whether an inserted document was relevant

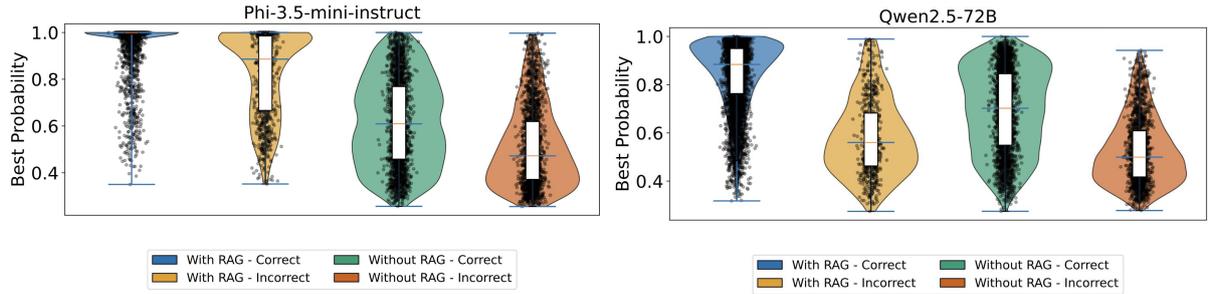


Figure 3: The transition of experimental results using MedMCQA. The figure classifies correctly answered and incorrectly answered questions, illustrating how their distributions shift. This visualization corresponds to the Ans1 setting, with plots for all three conditions: Pre-Q, Aft-Q, and Aft-C.

MedMCQA (Entropy and Best Probability)																	
Model	Pattern	Entropy (Correct) ↓				Best Prob (Correct) ↑				Entropy (Incorrect) ↓				Best Prob (Incorrect) ↓			
		None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3
Llama-2-70b-chat-hf	w/o RAG	1.24 $\pm$ 1.11	-	-	-	0.42 $\pm$ 0.55	-	-	-	1.28 $\pm$ 1.27	-	-	-	0.38 $\pm$ 0.38	-	-	-
	Pre-Q	-	1.12 $\pm$ 0.16	1.12 $\pm$ 0.16	1.27 $\pm$ 0.10	-	0.54 $\pm$ 0.12	0.54 $\pm$ 0.12	0.41 $\pm$ 0.09	-	1.27 $\pm$ 0.08	1.28 $\pm$ 0.08	1.31 $\pm$ 0.06	-	0.41 $\pm$ 0.07	0.40 $\pm$ 0.08	<b>0.38</b> $\pm$ 0.07
	Aft-Q	-	<b>1.11</b> $\pm$ 0.16	1.15 $\pm$ 0.16	1.29 $\pm$ 0.10	-	<b>0.55</b> $\pm$ 0.12	0.52 $\pm$ 0.13	0.40 $\pm$ 0.09	-	<b>1.27</b> $\pm$ 0.09	1.29 $\pm$ 0.08	1.31 $\pm$ 0.06	-	0.41 $\pm$ 0.08	0.40 $\pm$ 0.08	0.38 $\pm$ 0.07
	Aft-C	-	1.15 $\pm$ 0.16	1.23 $\pm$ 0.12	1.30 $\pm$ 0.10	-	0.51 $\pm$ 0.13	<b>0.46</b> $\pm$ 0.11	<b>0.39</b> $\pm$ 0.09	-	1.28 $\pm$ 0.09	1.31 $\pm$ 0.07	1.31 $\pm$ 0.07	-	0.40 $\pm$ 0.08	0.38 $\pm$ 0.08	0.38 $\pm$ 0.08
Llama-3.1-70B	w/o RAG	1.24 $\pm$ 1.24	-	-	-	0.42 $\pm$ 0.42	-	-	-	1.31 $\pm$ 1.31	-	-	-	0.35 $\pm$ 0.32	-	-	-
	Pre-Q	-	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	-	0.32 $\pm$ 0.03	0.33 $\pm$ 0.03	0.33 $\pm$ 0.03	-	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	-	0.32 $\pm$ 0.01	0.32 $\pm$ 0.01	0.33 $\pm$ 0.01
	Aft-Q	-	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	-	0.33 $\pm$ 0.03	0.33 $\pm$ 0.03	0.33 $\pm$ 0.03	-	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	-	0.32 $\pm$ 0.01	0.33 $\pm$ 0.01	0.33 $\pm$ 0.01
	Aft-C	-	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	1.35 $\pm$ 0.07	-	0.33 $\pm$ 0.03	0.33 $\pm$ 0.03	0.33 $\pm$ 0.03	-	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	1.35 $\pm$ 0.02	-	0.32 $\pm$ 0.01	0.33 $\pm$ 0.01	0.33 $\pm$ 0.01
Llama-3.1-8B	w/o RAG	1.38 $\pm$ 1.38	-	-	-	0.28 $\pm$ 0.28	-	-	-	1.38 $\pm$ 1.38	-	-	-	0.27 $\pm$ 0.27	-	-	-
	Pre-Q	-	1.38 $\pm$ 0.07	<b>1.38</b> $\pm$ 0.07	1.38 $\pm$ 0.07	-	0.28 $\pm$ 0.03	0.28 $\pm$ 0.03	0.28 $\pm$ 0.03	-	1.38 $\pm$ 0.02	1.38 $\pm$ 0.02	1.38 $\pm$ 0.02	-	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01
	Aft-Q	-	1.38 $\pm$ 0.07	1.38 $\pm$ 0.07	1.38 $\pm$ 0.06	-	0.28 $\pm$ 0.03	0.28 $\pm$ 0.03	0.28 $\pm$ 0.03	-	1.38 $\pm$ 0.02	<b>1.38</b> $\pm$ 0.02	1.38 $\pm$ 0.02	-	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01
	Aft-C	-	1.38 $\pm$ 0.07	1.38 $\pm$ 0.07	1.38 $\pm$ 0.07	-	0.28 $\pm$ 0.03	<b>0.28</b> $\pm$ 0.03	0.28 $\pm$ 0.03	-	1.38 $\pm$ 0.02	1.38 $\pm$ 0.02	1.38 $\pm$ 0.02	-	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01	0.27 $\pm$ 0.01
meditron-70b	w/o RAG	1.23 $\pm$ 1.05	-	-	-	0.42 $\pm$ 0.57	-	-	-	1.25 $\pm$ 1.19	-	-	-	0.40 $\pm$ 0.40	-	-	-
	Pre-Q	-	1.11 $\pm$ 0.17	<b>1.05</b> $\pm$ 0.19	1.23 $\pm$ 0.11	-	0.54 $\pm$ 0.13	<b>0.57</b> $\pm$ 0.14	0.47 $\pm$ 0.09	-	1.24 $\pm$ 0.10	<b>1.19</b> $\pm$ 0.11	1.27 $\pm$ 0.08	-	0.43 $\pm$ 0.08	0.47 $\pm$ 0.08	0.43 $\pm$ 0.08
	Aft-Q	-	1.10 $\pm$ 0.17	1.09 $\pm$ 0.17	1.21 $\pm$ 0.11	-	0.54 $\pm$ 0.13	0.55 $\pm$ 0.13	0.48 $\pm$ 0.09	-	1.23 $\pm$ 0.10	1.20 $\pm$ 0.10	1.23 $\pm$ 0.10	-	0.44 $\pm$ 0.08	0.47 $\pm$ 0.08	0.47 $\pm$ 0.09
	Aft-C	-	1.15 $\pm$ 0.17	1.28 $\pm$ 0.10	1.30 $\pm$ 0.08	-	0.51 $\pm$ 0.13	0.43 $\pm$ 0.09	0.41 $\pm$ 0.07	-	1.27 $\pm$ 0.08	1.31 $\pm$ 0.06	1.31 $\pm$ 0.06	-	0.41 $\pm$ 0.07	0.40 $\pm$ 0.07	0.41 $\pm$ 0.07
PMC-LLaMA-13B	w/o RAG	1.00 $\pm$ 1.00	-	-	-	0.56 $\pm$ 0.56	-	-	-	1.05 $\pm$ 1.05	-	-	-	0.53 $\pm$ 0.53	-	-	-
	Pre-Q	-	1.36 $\pm$ 0.06	1.36 $\pm$ 0.06	1.36 $\pm$ 0.06	-	0.33 $\pm$ 0.05	0.32 $\pm$ 0.05	0.32 $\pm$ 0.05	-	1.37 $\pm$ 0.03	1.37 $\pm$ 0.02	1.37 $\pm$ 0.02	-	0.31 $\pm$ 0.04	<b>0.31</b> $\pm$ 0.03	0.31 $\pm$ 0.03
	Aft-Q	-	1.36 $\pm$ 0.06	1.36 $\pm$ 0.06	1.36 $\pm$ 0.06	-	0.33 $\pm$ 0.05	0.32 $\pm$ 0.05	0.31 $\pm$ 0.05	-	1.36 $\pm$ 0.03	1.37 $\pm$ 0.03	1.37 $\pm$ 0.03	-	0.32 $\pm$ 0.04	0.31 $\pm$ 0.03	0.31 $\pm$ 0.04
	Aft-C	-	1.36 $\pm$ 0.06	1.35 $\pm$ 0.07	1.36 $\pm$ 0.07	-	0.33 $\pm$ 0.05	0.33 $\pm$ 0.05	0.33 $\pm$ 0.05	-	1.36 $\pm$ 0.03	1.36 $\pm$ 0.03	1.36 $\pm$ 0.03	-	0.32 $\pm$ 0.04	0.32 $\pm$ 0.04	0.32 $\pm$ 0.04
Gemma-2-2b	w/o RAG	1.17 $\pm$ 1.11	-	-	-	0.52 $\pm$ 0.56	-	-	-	1.18 $\pm$ 1.13	-	-	-	0.52 $\pm$ 0.51	-	-	-
	Pre-Q	-	1.12 $\pm$ 0.08	1.13 $\pm$ 0.07	1.15 $\pm$ 0.06	-	0.55 $\pm$ 0.05	0.54 $\pm$ 0.04	0.52 $\pm$ 0.04	-	1.17 $\pm$ 0.05	1.16 $\pm$ 0.04	1.16 $\pm$ 0.04	-	<b>0.51</b> $\pm$ 0.04	0.51 $\pm$ 0.04	0.52 $\pm$ 0.03
	Aft-Q	-	1.13 $\pm$ 0.07	1.14 $\pm$ 0.06	1.15 $\pm$ 0.06	-	0.55 $\pm$ 0.05	0.53 $\pm$ 0.04	0.52 $\pm$ 0.04	-	1.17 $\pm$ 0.05	1.16 $\pm$ 0.04	1.15 $\pm$ 0.04	-	0.51 $\pm$ 0.04	0.51 $\pm$ 0.04	0.52 $\pm$ 0.03
	Aft-C	-	<b>1.11</b> $\pm$ 0.08	1.12 $\pm$ 0.07	1.13 $\pm$ 0.06	-	<b>0.56</b> $\pm$ 0.05	0.55 $\pm$ 0.05	0.54 $\pm$ 0.04	-	1.16 $\pm$ 0.05	1.14 $\pm$ 0.05	<b>1.13</b> $\pm$ 0.05	-	0.53 $\pm$ 0.04	0.54 $\pm$ 0.04	0.54 $\pm$ 0.04
Phi-3.5	w/o RAG	0.93 $\pm$ 0.05	-	-	-	0.62 $\pm$ 0.98	-	-	-	1.09 $\pm$ 0.39	-	-	-	0.51 $\pm$ 0.51	-	-	-
	Pre-Q	-	0.06 $\pm$ 0.17	0.07 $\pm$ 0.18	0.24 $\pm$ 0.32	-	0.98 $\pm$ 0.08	0.98 $\pm$ 0.08	0.90 $\pm$ 0.15	-	<b>0.39</b> $\pm$ 0.34	0.43 $\pm$ 0.35	0.49 $\pm$ 0.38	-	0.84 $\pm$ 0.18	0.82 $\pm$ 0.18	0.80 $\pm$ 0.19
	Aft-Q	-	<b>0.05</b> $\pm$ 0.16	0.07 $\pm$ 0.18	0.34 $\pm$ 0.36	-	<b>0.98</b> $\pm$ 0.07	0.97 $\pm$ 0.08	0.87 $\pm$ 0.17	-	0.45 $\pm$ 0.35	0.46 $\pm$ 0.35	0.50 $\pm$ 0.37	-	0.81 $\pm$ 0.18	0.81 $\pm$ 0.18	0.80 $\pm$ 0.19
	Aft-C	-	0.09 $\pm$ 0.19	0.14 $\pm$ 0.22	0.27 $\pm$ 0.32	-	0.97 $\pm$ 0.09	0.95 $\pm$ 0.10	0.90 $\pm$ 0.15	-	0.45 $\pm$ 0.34	0.44 $\pm$ 0.36	0.42 $\pm$ 0.35	-	0.81 $\pm$ 0.18	0.82 $\pm$ 0.19	0.84 $\pm$ 0.17
Qwen2.5-14B	w/o RAG	0.86 $\pm$ 0.48	-	-	-	0.67 $\pm$ 0.85	-	-	-	1.06 $\pm$ 1.03	-	-	-	0.55 $\pm$ 0.49	-	-	-
	Pre-Q	-	0.52 $\pm$ 0.33	0.53 $\pm$ 0.35	0.89 $\pm$ 0.30	-	0.84 $\pm$ 0.15	0.83 $\pm$ 0.16	0.65 $\pm$ 0.18	-	<b>1.03</b> $\pm$ 0.23	1.05 $\pm$ 0.23	1.07 $\pm$ 0.22	-	0.56 $\pm$ 0.15	0.55 $\pm$ 0.15	0.54 $\pm$ 0.15
	Aft-Q	-	<b>0.48</b> $\pm$ 0.32	0.51 $\pm$ 0.33	0.92 $\pm$ 0.29	-	<b>0.85</b> $\pm$ 0.14	0.84 $\pm$ 0.16	0.63 $\pm$ 0.17	-	1.04 $\pm$ 0.24	1.05 $\pm$ 0.23	1.06 $\pm$ 0.21	-	0.56 $\pm$ 0.15	0.55 $\pm$ 0.15	0.55 $\pm$ 0.14
	Aft-C	-	0.66 $\pm$ 0.35	0.80 $\pm$ 0.32	1.00 $\pm$ 0.24	-	0.77 $\pm$ 0.17	0.71 $\pm$ 0.18	0.59 $\pm$ 0.16	-	1.12 $\pm$ 0.20	1.12 $\pm$ 0.20	1.15 $\pm$ 0.17	-	0.51 $\pm$ 0.13	0.51 $\pm$ 0.14	<b>0.49</b> $\pm$ 0.12
Qwen2.5-72B	w/o RAG	0.82 $\pm$ 0.44	-	-	-	0.69 $\pm$ 0.86	-	-	-	1.09 $\pm$ 0.87	-	-	-	0.53 $\pm$ 0.53	-	-	-
	Pre-Q	-	0.52 $\pm$ 0.31	0.56 $\pm$ 0.32	0.75 $\pm$ 0.32	-	0.83 $\pm$ 0.14	0.80 $\pm$ 0.15	0.71 $\pm$ 0.18	-	0.91 $\pm$ 0.26	<b>0.87</b> $\pm$ 0.24	0.94 $\pm$ 0.24	-	0.61 $\pm$ 0.16	0.63 $\pm$ 0.15	0.62 $\pm$ 0.15
	Aft-Q	-	<b>0.44</b> $\pm$ 0.30	0.54 $\pm$ 0.31	0.83 $\pm$ 0.31	-	<b>0.86</b> $\pm$ 0.13	0.82 $\pm$ 0.15	0.68 $\pm$ 0.17	-	0.99 $\pm$ 0.27	0.97 $\pm$ 0.24	0.99 $\pm$ 0.22	-	0.57 $\pm$ 0.16	0.59 $\pm$ 0.15	0.60 $\pm$ 0.15
	Aft-C	-	0.53 $\pm$ 0.32	0.65 $\pm$ 0.32	0.85 $\pm$ 0.28	-	0.82 $\pm$ 0.15	0.77 $\pm$ 0.16	0.67 $\pm$ 0.17	-	1.00 $\pm$ 0.25	1.02 $\pm$ 0.21	0.99 $\pm$ 0.22	-	0.57 $\pm$ 0.15	0.57 $\pm$ 0.14	0.61 $\pm$ 0.14

Table 3: Experimental result on MedMCQA. **Bold** indicates the best value among the models. Specifically, the lowest entropy and the highest best probability (Correct case) are highlighted. This table has numerical values and their standard deviations.

to the answer, leading to provide strong evidence that they function as suitable generators.

## 6 Analysis & Discussion

### How Does RAG Affect Prediction Probabilities?

Figure 3 presents partial experimental results using MedMCQA, while Figure 4 shows results from PubMedQA. These figures correspond to the Ans1 setting, where all three phases—Pre-Q, Aft-Q, and Aft-C—are plotted. A detailed analysis focuses on Phi and Qwen, which exhibited ideal behavior. When RAG was not applied, i.e., evaluating the models’ intrinsic accuracy, the output probabili-

ties were evenly distributed across both datasets. Furthermore, the results of Phi-3.5 on PubMedQA reveal that the incorrect predictions tend to concentrate at the upper end, i.e., where output probabilities are high. This pattern suggests that the model exhibits overconfidence, making incorrect predictions despite assigning high probabilities. When solving a QA task under a deliberate setting that includes supporting documents for correct answers (similar to pseudo-RAG), all models (Phi and Qwen) showed improved output probabilities. This suggests that the models can assess whether retrieved documents contain useful information.

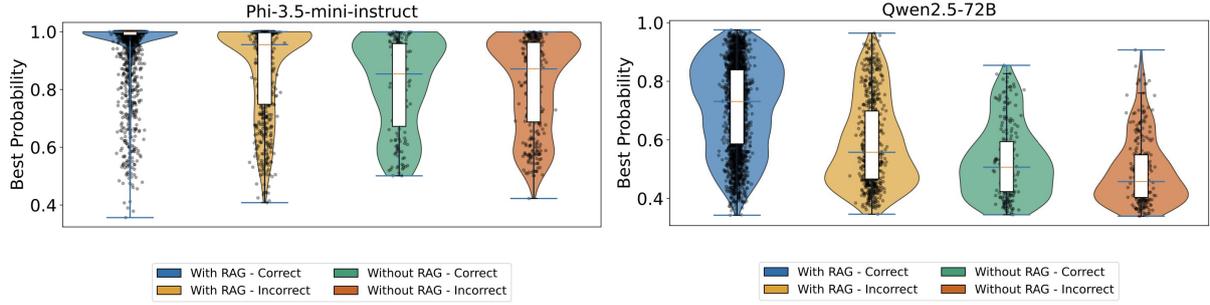


Figure 4: The transition of experimental results using PubMedQA. The figure classifies correctly answered and incorrectly answered questions, illustrating how their distributions shift. This visualization corresponds to the Ans1 setting, with plots for all three conditions: Pre-Q, Aft-Q, and Aft-C.

PubMedQA (Entropy and Best Probability)																	
Model	Pattern	Entropy (Correct) ↓				Best Prob (Correct) ↑				Entropy (Incorrect) ↓				Best Prob (Incorrect) ↓			
		None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3
Llama-2-70b-chat-hf	w/o RAG	0.83±0.83	-	-	-	0.62±0.62	-	-	-	0.93±0.93	-	-	-	0.55±0.38	-	-	-
	Pre-Q	-	1.12±0.16	1.12±0.16	1.27±0.10	-	0.54±0.12	0.54±0.12	0.41±0.09	-	1.27±0.08	1.28±0.08	1.31±0.06	-	0.41±0.07	0.40±0.08	<b>0.38±0.07</b>
	Aft-Q	-	1.11±0.16	1.15±0.16	1.29±0.10	-	0.55±0.12	0.52±0.13	0.40±0.09	-	1.27±0.09	1.29±0.08	1.31±0.06	-	0.41±0.08	0.40±0.08	0.38±0.07
	Aft-C	-	1.15±0.16	1.23±0.12	1.30±0.10	-	0.51±0.13	0.46±0.11	0.39±0.09	-	1.28±0.09	1.31±0.07	1.31±0.07	-	0.40±0.08	0.38±0.08	0.38±0.08
Llama-3.1-70B	w/o RAG	0.86±0.86	-	-	-	0.59±0.59	-	-	-	0.87±0.87	-	-	-	0.57±0.32	-	-	-
	Pre-Q	-	1.35±0.07	1.35±0.07	1.35±0.07	-	0.32±0.03	0.33±0.03	0.33±0.03	-	1.35±0.02	1.35±0.02	1.35±0.02	-	<b>0.32±0.01</b>	0.32±0.01	0.33±0.01
	Aft-Q	-	1.35±0.07	1.35±0.07	1.35±0.07	-	0.33±0.03	0.33±0.03	0.33±0.03	-	1.35±0.02	1.35±0.02	1.35±0.02	-	0.32±0.01	0.33±0.01	0.33±0.01
	Aft-C	-	1.35±0.07	1.35±0.07	1.35±0.07	-	0.33±0.03	0.33±0.03	0.33±0.03	-	1.35±0.02	1.35±0.02	1.35±0.02	-	0.32±0.01	0.33±0.01	0.33±0.01
Llama-3.1-8B	w/o RAG	1.09±1.09	-	-	-	0.36±0.36	-	-	-	1.09±1.09	-	-	-	0.36±0.27	-	-	-
	Pre-Q	-	1.38±0.07	1.38±0.07	1.38±0.07	-	0.28±0.03	0.28±0.03	0.28±0.03	-	1.38±0.02	1.38±0.02	1.38±0.02	-	0.27±0.01	0.27±0.01	<b>0.27±0.01</b>
	Aft-Q	-	1.38±0.07	1.38±0.07	1.38±0.06	-	0.28±0.03	0.28±0.03	0.28±0.03	-	1.38±0.02	1.38±0.02	1.38±0.02	-	0.27±0.01	0.27±0.01	0.27±0.01
	Aft-C	-	1.38±0.07	1.38±0.07	1.38±0.07	-	0.28±0.03	0.28±0.03	0.28±0.03	-	1.38±0.02	1.38±0.02	1.38±0.02	-	0.27±0.01	0.27±0.01	0.27±0.01
meditron-70b	w/o RAG	0.94±0.94	-	-	-	0.54±0.57	-	-	-	0.94±0.94	-	-	-	0.52±0.40	-	-	-
	Pre-Q	-	1.11±0.17	1.05±0.19	1.23±0.11	-	0.54±0.13	<b>0.57±0.14</b>	0.47±0.09	-	1.24±0.10	1.19±0.11	1.27±0.08	-	0.43±0.08	0.47±0.08	0.43±0.08
	Aft-Q	-	1.10±0.17	1.09±0.17	1.21±0.11	-	0.54±0.13	0.55±0.13	0.48±0.09	-	1.23±0.10	1.20±0.10	1.25±0.10	-	0.44±0.08	0.47±0.08	0.47±0.09
	Aft-C	-	1.15±0.17	1.28±0.10	1.30±0.08	-	0.51±0.13	0.43±0.09	0.41±0.07	-	1.27±0.08	1.31±0.06	1.31±0.06	-	0.41±0.07	<b>0.40±0.07</b>	0.41±0.07
PMC-LLaMA-13B	w/o RAG	1.08±1.08	-	-	-	0.40±0.40	-	-	-	1.08±1.08	-	-	-	0.40±0.31	-	-	-
	Pre-Q	-	1.36±0.06	1.36±0.06	1.36±0.06	-	0.33±0.05	0.32±0.05	0.32±0.05	-	1.37±0.03	1.37±0.02	1.37±0.02	-	0.31±0.04	<b>0.31±0.03</b>	0.31±0.03
	Aft-Q	-	1.36±0.06	1.36±0.06	1.36±0.06	-	0.33±0.05	0.32±0.05	0.31±0.05	-	1.36±0.03	1.37±0.03	1.37±0.03	-	0.32±0.04	0.31±0.03	0.31±0.04
	Aft-C	-	1.36±0.06	1.35±0.07	1.36±0.07	-	0.33±0.05	0.33±0.05	0.33±0.05	-	1.36±0.03	1.36±0.03	1.36±0.03	-	0.32±0.04	0.32±0.04	0.32±0.04
Gemma-2-2b	w/o RAG	0.93±0.93	-	-	-	0.61±0.61	-	-	-	0.93±0.93	-	-	-	0.61±0.51	-	-	-
	Pre-Q	-	1.12±0.08	1.13±0.07	1.15±0.06	-	0.55±0.05	0.54±0.04	0.52±0.04	-	1.17±0.05	1.16±0.04	1.16±0.04	-	<b>0.51±0.04</b>	0.51±0.04	0.52±0.03
	Aft-Q	-	1.13±0.07	1.14±0.06	1.15±0.06	-	0.55±0.05	0.53±0.04	0.52±0.04	-	1.17±0.05	1.16±0.04	1.15±0.04	-	0.51±0.04	0.51±0.04	0.52±0.03
	Aft-C	-	1.11±0.08	1.12±0.07	1.13±0.06	-	0.56±0.05	0.55±0.05	0.54±0.04	-	1.16±0.05	1.14±0.05	1.13±0.05	-	0.53±0.04	0.54±0.04	0.54±0.04
Phi-3.5	w/o RAG	0.40±0.05	-	-	-	0.81±0.98	-	-	-	0.41±0.39	-	-	-	0.82±0.80	-	-	-
	Pre-Q	-	0.06±0.17	0.07±0.18	0.24±0.32	-	0.98±0.08	0.98±0.08	0.90±0.15	-	<b>0.39±0.34</b>	0.43±0.35	0.49±0.38	-	0.84±0.18	0.82±0.18	<b>0.80±0.19</b>
	Aft-Q	-	<b>0.05±0.16</b>	0.07±0.18	0.34±0.36	-	<b>0.98±0.07</b>	0.97±0.08	0.87±0.17	-	0.45±0.35	0.46±0.35	0.50±0.37	-	0.81±0.18	0.81±0.18	<b>0.80±0.19</b>
	Aft-C	-	0.09±0.19	0.14±0.22	0.27±0.32	-	0.97±0.09	0.95±0.10	0.90±0.15	-	0.45±0.34	0.44±0.36	0.42±0.35	-	0.81±0.18	0.82±0.19	<b>0.84±0.17</b>
Qwen2.5-14B	w/o RAG	0.90±0.48	-	-	-	0.59±0.85	-	-	-	0.92±0.92	-	-	-	0.58±0.49	-	-	-
	Pre-Q	-	0.52±0.33	0.53±0.35	0.89±0.30	-	0.84±0.15	0.83±0.16	0.65±0.18	-	1.03±0.23	1.05±0.23	1.07±0.22	-	0.56±0.15	0.55±0.15	0.54±0.15
	Aft-Q	-	<b>0.48±0.32</b>	0.51±0.33	0.92±0.29	-	<b>0.85±0.14</b>	0.84±0.16	0.63±0.17	-	1.04±0.24	1.05±0.23	1.06±0.21	-	0.56±0.15	0.55±0.15	0.55±0.14
	Aft-C	-	0.66±0.35	0.80±0.32	1.00±0.24	-	0.77±0.17	0.71±0.18	0.59±0.16	-	1.12±0.20	1.12±0.20	1.15±0.17	-	0.51±0.13	0.51±0.14	<b>0.49±0.12</b>
Qwen2.5-72B	w/o RAG	0.97±0.44	-	-	-	0.53±0.86	-	-	-	1.00±0.87	-	-	-	0.49±0.49	-	-	-
	Pre-Q	-	0.52±0.31	0.56±0.32	0.75±0.32	-	0.83±0.14	0.80±0.15	0.71±0.18	-	0.91±0.26	<b>0.87±0.24</b>	0.94±0.24	-	0.61±0.16	0.63±0.15	0.62±0.15
	Aft-Q	-	<b>0.44±0.30</b>	0.54±0.31	0.83±0.31	-	<b>0.86±0.13</b>	0.82±0.15	0.68±0.17	-	0.99±0.27	0.97±0.24	0.99±0.22	-	0.57±0.16	0.59±0.15	0.60±0.15
	Aft-C	-	0.53±0.32	0.65±0.32	0.85±0.28	-	0.82±0.15	0.77±0.16	0.67±0.17	-	1.00±0.25	1.02±0.21	0.99±0.22	-	0.57±0.15	0.57±0.14	0.61±0.14

Table 4: Experimental results using PubMedQA. **Bold** indicates the best value among the models. Specifically, the lowest entropy and the highest best probability (Correct case) are highlighted. This table has numerical values and their standard deviations.

**Model Behavior When Inserting Answer-Containing Documents.** When explicitly inserting documents that contain the correct answers, Phi and Qwen demonstrated ideal behavior. For instance, from a correct entropy perspective in Table 3, Phi had a value of 0.933 under the w/o RAG setting, which dropped to 0.051 after document insertion. Similarly, for Qwen2.5 (72B), entropy decreased from 0.819 to 0.444. This observation indicates that the models can assess whether an inserted document is relevant to the question. Moreover, when they determine that the document is unnecessary, they attempt to answer using their

own knowledge. Further evidence supporting this conclusion comes from cases where inserting unrelated documents did not improve accuracy. This suggests that the models selectively utilize external information only when it is deemed useful.

**Behavior of Calibration Error.** Table 5 presents the evaluation result of ACE using MedMCQA, while Table 6 shows the results for PubMedQA. A detailed analysis of Llama and Gemma reveals substantial differences in behavior depending on the model. Notably, even when inserting entirely correct documents (Ans1), Llama3.1 (70B) experi-

Model	Pattern	ACE ↓				Accuracy ↑			
		None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3
Llama2 (70B)	w/o RAG	2.208	-	-	-	38.322	-	-	-
	Pre-Q	-	22.359	25.113	7.181	-	72.575	75.340	32.094
	Aft-Q	-	23.912	21.132	10.781	-	75.567	69.628	28.105
	Aft-C	-	19.653	17.514	9.803	-	67.498	60.743	28.876
Llama3.1 (70B)	w/o RAG	19.582	-	-	-	58.977	-	-	-
	Pre-Q	-	11.496	11.580	11.658	-	20.898	20.943	20.943
	Aft-Q	-	11.518	11.671	11.714	-	20.898	20.898	20.943
	Aft-C	-	11.504	11.707	11.795	-	20.943	20.943	20.943
Llama3.1 (8B)	w/o RAG	5.423	-	-	-	22.209	-	-	-
	Pre-Q	-	4.701	4.254	3.644	-	23.345	23.209	23.799
	Aft-Q	-	4.473	4.632	3.892	-	23.028	23.209	24.025
	Aft-C	-	4.476	4.746	4.990	-	23.209	23.255	23.663
Meditron (70B)	w/o RAG	6.412	-	-	-	35.525	-	-	-
	Pre-Q	-	17.684	7.652	8.665	-	67.724	54.034	36.038
	Aft-Q	-	15.894	9.467	15.334	-	66.682	47.144	31.958
	Aft-C	-	15.101	6.946	9.006	-	62.829	34.180	31.913
PMC-Llama (13B)	w/o RAG	15.671	-	-	-	38.107	-	-	-
	Pre-Q	-	4.943	4.367	4.357	-	32.729	31.641	26.972
	Aft-Q	-	4.003	2.550	5.032	-	32.910	30.009	26.972
	Aft-C	-	3.496	3.780	4.397	-	33.454	28.740	28.060
Gemma2 (2B)	w/o RAG	19.568	-	-	-	32.297	-	-	-
	Pre-Q	-	25.511	23.160	20.520	-	31.233	31.278	31.278
	Aft-Q	-	24.160	21.072	20.618	-	31.188	31.278	31.278
	Aft-C	-	24.814	23.118	22.916	-	31.324	31.324	31.324
Phi3.5 (3.8B)	w/o RAG	5.624	-	-	-	51.518	-	-	-
	Pre-Q	-	9.786	10.378	33.709	-	86.083	84.950	51.813
	Aft-Q	-	7.636	9.270	43.415	-	88.486	85.947	39.393
	Aft-C	-	7.682	15.952	42.476	-	87.353	76.111	44.334
Qwen2.5 (14B)	w/o RAG	12.125	-	-	-	49.151	-	-	-
	Pre-Q	-	8.646	8.740	11.892	-	89.483	88.441	47.280
	Aft-Q	-	7.013	7.257	17.592	-	89.121	87.534	40.798
	Aft-C	-	7.746	9.778	8.531	-	79.329	75.884	45.014
Qwen2.5 (72B)	w/o RAG	4.030	-	-	-	60.483	-	-	-
	Pre-Q	-	9.393	7.896	20.412	-	89.982	85.766	45.739
	Aft-Q	-	8.782	9.781	18.652	-	93.246	89.574	44.696
	Aft-C	-	9.270	5.990	23.564	-	88.622	79.284	39.483

Table 5: Evaluation results with MedMCQA. Red highlights areas where performance improved compared to the non-RAG setting, while Blue indicates areas where performance deteriorated.

ences a drop in accuracy, whereas Llama3.1 (8B) shows improved accuracy even when inserting completely unrelated documents (Oth3). This stark contrast indicates that even within the same model family, behavior can vary largely. Moreover, a comparison between Llama2, Meditron, and Llama3.1 shows considerable differences in behavior, ruling out parameter size as the primary cause. These findings suggest that while the Llama series performs well under specific instruction formats, it may negatively impact performance in other cases. On the other hand, Qwen and Phi exhibit a clear pat-

tern: inserting entirely unrelated documents (Oth3) worsens ACE, while inserting answer-containing documents (Ans1 or Ans1-Oth2) leads to improvements. This tendency implies that Phi and Qwen possess the ability to assess whether retrieved documents provide useful information. These results show that analyzing LLM confidence through predicted probabilities effectively reveals the model’s ability to identify meaningful documents.

**Did “Lost in the Middle” Phenomenon Occur?** Our study also examined the “Lost in the Middle” phenomenon (Liu et al., 2024) by evaluating the

Model	Pattern	ACE ↓				Accuracy ↑			
		None	Ans1	Ans1-Oth2	Oth3	None	Ans1	Ans1-Oth2	Oth3
Llama2 (70B)	w/o RAG	12.107	-	-	-	46.400	-	-	-
	Pre-Q	-	29.791	30.422	14.146	-	82.200	79.800	56.500
	Aft-Q	-	30.380	31.220	13.234	-	81.600	74.100	53.000
	Aft-C	-	13.494	13.430	11.322	-	57.200	53.100	50.700
Llama3.1 (70B)	w/o RAG	6.091	-	-	-	58.600	-	-	-
	Pre-Q	-	11.329	11.513	11.521	-	55.200	55.200	55.200
	Aft-Q	-	11.343	11.532	11.539	-	55.200	55.200	55.200
	Aft-C	-	11.370	11.534	11.543	-	55.200	55.200	55.200
Llama3.1 (8B)	w/o RAG	24.939	-	-	-	11.000	-	-	-
	Pre-Q	-	23.683	23.368	23.975	-	12.200	12.600	12.000
	Aft-Q	-	23.085	23.576	23.988	-	12.800	12.400	12.000
	Aft-C	-	23.930	23.854	24.370	-	11.900	12.100	11.600
Meditron (70B)	w/o RAG	18.115	-	-	-	34.800	-	-	-
	Pre-Q	-	11.540	18.483	8.365	-	57.300	69.800	57.200
	Aft-Q	-	9.159	6.645	6.270	-	56.700	55.600	54.800
	Aft-C	-	4.171	5.050	7.915	-	54.700	54.800	55.100
PMC-Llama (13B)	w/o RAG	17.261	-	-	-	22.800	-	-	-
	Pre-Q	-	10.462	4.650	3.387	-	28.800	37.900	36.600
	Aft-Q	-	10.322	4.000	3.985	-	28.900	39.200	40.000
	Aft-C	-	4.169	5.421	7.250	-	41.200	44.500	46.100
Gemma2 (2B)	w/o RAG	6.387	-	-	-	55.200	-	-	-
	Pre-Q	-	5.794	5.409	5.394	-	55.300	55.200	55.200
	Aft-Q	-	6.159	5.098	4.188	-	55.200	55.200	55.200
	Aft-C	-	9.081	6.161	6.376	-	55.200	55.200	55.200
Phi3.5 (3.8B)	w/o RAG	48.176	-	-	-	33.400	-	-	-
	Pre-Q	-	14.640	14.777	57.831	-	81.600	81.200	21.900
	Aft-Q	-	13.677	31.960	52.083	-	82.300	62.700	41.300
	Aft-C	-	16.771	33.297	47.123	-	73.700	52.300	33.800
Qwen2.5 (14B)	w/o RAG	15.874	-	-	-	42.800	-	-	-
	Pre-Q	-	4.746	4.816	18.425	-	83.400	83.200	32.600
	Aft-Q	-	3.460	5.013	26.783	-	82.800	76.100	33.900
	Aft-C	-	7.616	3.088	23.229	-	74.500	63.900	32.100
Qwen2.5 (72B)	w/o RAG	7.205	-	-	-	46.400	-	-	-
	Pre-Q	-	10.477	3.801	25.283	-	74.900	78.100	33.000
	Aft-Q	-	8.024	10.931	17.828	-	80.300	71.200	34.300
	Aft-C	-	8.877	6.995	13.543	-	76.800	71.000	42.500

Table 6: Evaluation results on PubMedQA. Red highlights areas where performance improved compared to the non-RAG setting, while Blue indicates areas where performance deteriorated.

impact of document placement within the template across multiple positions (Pre-Q, Aft-Q, and Aft-C). Focusing on Phi and Qwen, which exhibited expected behavior in terms of entropy and accuracy, an intriguing pattern emerged. From an entropy perspective, inserting the document after the answer choices yielded the best results, while from an accuracy perspective, placing it before the answer choices was optimal. These findings suggest that when prioritizing the reliability of information, entropy should be the primary metric.

**Error Analysis.** Appendix B.1 presents the results of the error analysis, which examines how the model makes mistakes. PubMedQA consists of three answer choices: yes, no, and maybe, allowing us to analyze the types of errors the model makes. For Llama3.1 (8B, 70B), PMC-Llama, and Gemma2, the bin colors remain unchanged, indicating that these models do not incorporate arbitrarily inserted supporting documents (Ans1, Ans1-Oth2). Meditron, w/o RAG, outputs “No” for all incorrect answers. However, when a document is provided, it changes all responses to “Yes,” revealing an ex-

tremely sensitive behavior.

## 7 Conclusion

Our research explored the impact of Retrieval Augmented Generation (RAG) on model confidence in the medical domain where information reliability is crucial. We found that when models retrieve relevant documents, they not only boost accuracy but also show higher confidence scores. In contrast, irrelevant documents have little effect on improving confidence. Several models demonstrate the ability to judge if the retrieved documents connect to the correct answer, indicating a more discerning use of external information than we anticipated. Our evaluation metrics provide a clear framework for spotting the best generator models within RAG systems. The findings reveal that models adjust their output probabilities in response to the quality of the retrieved documents, which opens up new ways to measure and improve model performance. These insights help refine RAG methods, making them more reliable for high-stakes applications.

## 8 Limitations

### 8.1 The Experiments of the Other Domain

Our study prioritizes domains where RAG is applied, focusing specifically on the medical domain to analyze confidence. To advance further, it becomes necessary to validate RAG in domains such as finance and analyze its confidence in contexts requiring highly reliable information.

### 8.2 Further Analyzing New RAG Architecture

Our study focused exclusively on analyzing the basic RAG architecture. While the standard RAG framework directly utilizes retrieved documents within the LLM, newer RAG architectures incorporate various control mechanisms. Moving forward, it is essential to analyze these advanced architectures from the perspective of confidence as well.

### 8.3 Other Metrics for Evaluation

The evaluation metrics used in this study, ACE, may have some drawbacks. (Kull et al., 2019; Kumar et al., 2019; Baan et al., 2022). Since LLMs outputs are not always strictly correct or incorrect, researchers often use Prediction Rejection Ratio (PRR), which measures the correlation between confidence scores and output quality. (Fadeeva et al., 2023; Vashurin et al., 2025; He et al., 2024; Ozaki et al., 2025b). However, our study focuses

on a multiple-QA task, where each question has a uniquely defined correct answer. Additionally, the models were evaluated using force decoding. Given these conditions, ACE serves as appropriate evaluation metrics.

## 8.4 Methods for Generating Model Outputs

This study deliberately avoids generating free-text responses from models. Instead, it retrieves answer candidates using force decoding. This decision stems from an observation in prior research: many studies rely heavily on regular expressions to extract correct answers, leading to substantial accuracy variations even when using the same QA task and model. (<https://github.com/Teddy-XiongGZ/MedRAG>, <https://github.com/epfLLM/meditron>, <https://github.com/chaoyi-wu/PMC-LLaMA>.) To address this issue, we select answer choices based on the model’s inherent output probabilities. This approach avoids introducing dependencies on specific evaluation metrics, which would otherwise occur if the model were required to generate explanations using Chain-of-Thought (CoT) or produce confidence scores.

## 9 Ethical Considerations

### 9.1 The Possibility of Dataset Bias

The datasets and retrieval mechanisms employed in our study may carry inherent biases, which could influence the model’s predictions and potentially affect fairness in decision-making. Recognizing these biases, we advocate for the use of diverse and representative datasets to minimize their impact. Additionally, we uphold transparency by analyzing the interplay between confidence and accuracy, providing users with clearer insights into the system’s limitations and confidence. However, we emphasize the need for human oversight, as no automated system can guarantee infallibility.

### 9.2 AI Assistant Tools

We used ChatGPT<sup>2</sup> and DeepL<sup>3</sup> to translate sentences to English and accelerate our research.

## References

Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat

<sup>2</sup><https://chatgpt.com/>

<sup>3</sup><https://www.deepl.com/ja/translator>

- Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219*.
- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Akiko Aizawa, Eiji Aramaki, Bowen Chen, Fei Cheng, Hiroyuki Deguchi, Rintaro Enomoto, Kazuki Fujii, Kensuke Fukumoto, Takuya Fukushima, Namgi Han, et al. 2024. Llm-jp: A cross-organizational project for the research and development of fully open japanese llms. *arXiv preprint arXiv:2407.03963*.
- Akari Asai, Zeqiu Wu, Yizhong Wang, Avirup Sil, and Hannaneh Hajishirzi. 2024. **Self-RAG: Learning to retrieve, generate, and critique through self-reflection**. In *The Twelfth International Conference on Learning Representations*.
- Joris Baan, Wilker Aziz, Barbara Plank, and Raquel Fernandez. 2022. **Stop measuring calibration when humans disagree**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 1892–1915, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Evan Becker and Stefano Soatto. 2024. **Cycles of thought: Measuring llm confidence through stable explanations**. *Preprint*, arXiv:2406.03441.
- Lu Chen, Ruqing Zhang, Jiafeng Guo, Yixing Fan, and Xueqi Cheng. 2024. **Controlling risk of retrieval-augmented generation: A counterfactual prompting framework**. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 2380–2393, Miami, Florida, USA. Association for Computational Linguistics.
- Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. 2023. **Meditron-70b: Scaling medical pretraining for large language models**. *Preprint*, arXiv:2311.16079.
- Jeremy Cole, Michael Zhang, Daniel Gillick, Julian Eisenschlos, Bhuwan Dhingra, and Jacob Eisenstein. 2023. **Selectively answering ambiguous questions**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 530–543, Singapore. Association for Computational Linguistics.
- Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2022. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. *Advances in Neural Information Processing Systems*, 35:30318–30332.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Ekaterina Fadeeva, Roman Vashurin, Akim Tsvigun, Artem Vazhentsev, Sergey Petrakov, Kirill Fedyanin, Daniil Vasilev, Elizaveta Goncharova, Alexander Panchenko, Maxim Panov, Timothy Baldwin, and Artem Shelmanov. 2023. **LM-polygraph: Uncertainty estimation for language models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 446–461, Singapore. Association for Computational Linguistics.
- Jiahui Geng, Fengyu Cai, Yuxia Wang, Heinz Koepl, Preslav Nakov, and Iryna Gurevych. 2024. **A survey of confidence estimation and calibration in large language models**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6577–6595, Mexico City, Mexico. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017. On calibration of modern neural networks. In *International Conference on Machine Learning*, pages 1321–1330. PMLR.
- Kazuki Hayashi, Yusuke Sakai, Hidetaka Kamigaito, Katsuhiko Hayashi, and Taro Watanabe. 2024. **Towards artwork explanation in large-scale vision language models**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 705–729, Bangkok, Thailand. Association for Computational Linguistics.
- Jianfeng He, Runing Yang, Linlin Yu, Changbin Li, Ruoxi Jia, Feng Chen, Ming Jin, and Chang-Tien Lu. 2024. **Can we trust the performance evaluation of uncertainty estimation methods in text summarization?** In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 16514–16575, Miami, Florida, USA. Association for Computational Linguistics.
- Minbyul Jeong, Jiwoong Sohn, Mujeen Sung, and Jae-woo Kang. 2024. Improving medical reasoning through retrieval and self-reflection with retrieval-augmented large language models. *Bioinformatics*, 40(Supplement\_1):i119–i129.
- Zhengbao Jiang, Jun Araki, Haibo Ding, and Graham Neubig. 2021. **How can we know when language models know? on the calibration of language models for question answering**. *Transactions of the Association for Computational Linguistics*, 9:962–977.
- Zhengbao Jiang, Frank Xu, Luyu Gao, Zhiqing Sun, Qian Liu, Jane Dwivedi-Yu, Yiming Yang, Jamie Callan, and Graham Neubig. 2023. **Active retrieval**

- augmented generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7969–7992, Singapore. Association for Computational Linguistics.
- Qiao Jin, Bhuwan Dhingra, Zhengping Liu, William Cohen, and Xinghua Lu. 2019. **PubMedQA: A dataset for biomedical research question answering**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2567–2577, Hong Kong, China. Association for Computational Linguistics.
- Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. 2019. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *Advances in neural information processing systems*, 32.
- Ananya Kumar, Percy S Liang, and Tengyu Ma. 2019. Verified uncertainty calibration. *Advances in Neural Information Processing Systems*, 32.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2024. **Lost in the middle: How language models use long contexts**. *Transactions of the Association for Computational Linguistics*, 12:157–173.
- Xinbei Ma, Yeyun Gong, Pengcheng He, Hai Zhao, and Nan Duan. 2023. **Query rewriting in retrieval-augmented large language models**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5303–5315, Singapore. Association for Computational Linguistics.
- Mahdi Pakdaman Naeini, Gregory F Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2901–2907.
- Jeremy Nixon, Michael W. Dusenberry, Linchuan Zhang, Ghassen Jerfel, and Dustin Tran. 2019. Measuring calibration in deep learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Shintaro Ozaki, Kazuki Hayashi, Miyu Oba, Yusuke Sakai, Hidetaka Kamigaito, and Taro Watanabe. 2024. Bqa: Body language question answering dataset for video large language models. *arXiv preprint arXiv:2410.13206*.
- Shintaro Ozaki, Kazuki Hayashi, Yusuke Sakai, Hidetaka Kamigaito, Katsuhiko Hayashi, and Taro Watanabe. 2025a. **Towards cross-lingual explanation of artwork in large-scale vision language models**. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 3773–3809, Albuquerque, New Mexico. Association for Computational Linguistics.
- Shintaro Ozaki, Kazuki Hayashi, Yusuke Sakai, Jingu Kwon, Hidetaka Kamigaito, Katsuhiko Hayashi, Manabu Okumura, and Taro Watanabe. 2025b. Text-tiger: Text-based intelligent generation with entity prompt refinement for text-to-image generation. *arXiv preprint arXiv:2504.18269*.
- Ankit Pal, Logesh Kumar Umaphathi, and Malaikannan Sankarasubbu. 2022. **Medmcqa: A large-scale multi-subject multi-choice dataset for medical domain question answering**. In *Proceedings of the Conference on Health, Inference, and Learning*, volume 174 of *Proceedings of Machine Learning Research*, pages 248–260. PMLR.
- Irina Proskurina, Luc Brun, Guillaume Metzler, and Julien Velcin. 2024. **When quantization affects confidence of large language models?** In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 1918–1928, Mexico City, Mexico. Association for Computational Linguistics.
- Yusuke Sakai, Hidetaka Kamigaito, and Taro Watanabe. 2024. **mCSQA: Multilingual commonsense reasoning dataset with unified creation strategy by language models and humans**. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14182–14214, Bangkok, Thailand. Association for Computational Linguistics.
- Spurthi Setty, Harsh Thakkar, Alyssa Lee, Eden Chung, and Natan Vidra. 2024. Improving retrieval for rag based question answering models on financial documents. *arXiv preprint arXiv:2404.07221*.
- Jiwoong Sohn, Yein Park, Chanwoong Yoon, Sihyeon Park, Hyeon Hwang, Mujeen Sung, Hyunjae Kim, and Jaewoo Kang. 2024. Rationale-guided retrieval augmented generation for medical question answering. *arXiv preprint arXiv:2411.00300*.
- Zhiqing Sun, Xuezhi Wang, Yi Tay, Yiming Yang, and Denny Zhou. 2023. **Recitation-augmented language models**. In *International Conference on Learning Representations*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Dennis Ulmer, Jes Frellsen, and Christian Hardmeier. 2022. *Exploring predictive uncertainty and calibration in NLP: A study on the impact of method & data scarcity*. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2707–2735, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Roman Vashurin, Ekaterina Fadeeva, Artem Vazhentsev, Lyudmila Rvanova, Akim Tsvigun, Daniil Vasilev, Rui Xing, Abdelrahman Boda Sadallah, Kirill Grishchenkov, Sergey Petrakov, Alexander Panchenko, Timothy Baldwin, Preslav Nakov, Maxim Panov, and Artem Shelmanov. 2025. *Benchmarking uncertainty quantification methods for large language models with lm-polygraph*. *Preprint*, arXiv:2406.15627.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems*, 35:24824–24837.

Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. *Transformers: State-of-the-art natural language processing*. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.

Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. 2024. Pmc-llama: toward building open-source language models for medicine.

Guangzhi Xiong, Qiao Jin, Zhiyong Lu, and Aidong Zhang. 2024. *Benchmarking retrieval-augmented generation for medicine*. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 6233–6251, Bangkok, Thailand. Association for Computational Linguistics.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.

Antonio Jimeno Yepes, Yao You, Jan Milczek, Sebastian Laverde, and Renyu Li. 2024. Financial report chunking for effective retrieval augmented generation. *arXiv preprint arXiv:2402.05131*.

Jiaxin Zhang, Zhuohang Li, Kamalika Das, Bradley Malin, and Sricharan Kumar. 2023. *SAC<sup>3</sup>: Reliable hallucination detection in black-box language models via semantic-aware cross-check consistency*. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 15445–15458, Singapore. Association for Computational Linguistics.

Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and Sameer Singh. 2021. Calibrate before use: Improving few-shot performance of language models. In *International conference on machine learning*, pages 12697–12706. PMLR.

## A Example Appendix

### A.1 Detailed Model Settings

The PMC-Llama model was quantized to half-precision, while the 70B / 72B models were quantized to 4-bit precision for experimentation. The implementation relied on the Transformers library (Wolf et al., 2020) and bitsandbytes (Dettmers et al., 2022).

Model	Params	HuggingFace Name
Phi-3.5	3.8B	microsoft/Phi-3.5-mini-instruct
PMC-Llama	13B	axiong/PMC_LLAMA_13B
LLama2	70B	meta-llama/LLama-2-70b-chat-hf
Meditron	70B	epfl-llm/meditron-70b
Llama3.1	8B	meta-llama/LLama-3.1-8B
Llama3.1	70B	meta-llama/LLama-3.1-70B
Gemma2	2B	google/gemma-2-2b
Qwen2.5	14B	Qwen/Qwen2.5-14B
Qwen2.5	72B	Qwen/Qwen2.5-72B

Table 7: Detailed name of models.

### A.2 Dataset Selection

The dataset selection is based on prior research by Xiong et al. (2024). From the datasets they used, we select those that include both QA pairs and explanatory passages that justify the answers (MedMCQA and PubMedQA) for this study.

Since the test set for MedMCQA is not publicly available, our study used the dev set as the test set, following the approach adopted in MIRAGE<sup>4</sup>. We used the datasets, especially MedMCQA<sup>5</sup>, PubMedQA<sup>6</sup>.

### A.3 Details of the Input Format

As described in Section 3, we determine the selected choice based on the output probabilities as-

<sup>4</sup><https://huggingface.co/datasets/openlifescienceai/medmcqa>

<sup>5</sup><https://github.com/MedMCQA/MedMCQA>

<sup>6</sup><https://github.com/pubmedqa/pubmedqa>

signed by LLMs to the given candidates. To prevent answer choices from being split into multiple tokens by the tokenizer, we replace them with A, B, C, and D before feeding them into the model. This approach ensures a fair comparison across models, even for answer choices that would otherwise span multiple tokens.

#### A.4 Inference Settings

In this study, as far as inference which needs to use GPUs, all experiments were conducted on a single NVIDIA RTX A6000 and NVIDIA GeForce RTX 3090 GPU.

#### A.5 Why Do We Focus on the Medical Domain?

Among the various domains where information reliability is crucial (e.g., finance, law, autonomous driving, and healthcare), we chose to focus on healthcare for the following reasons:

- **Complexity and Scale of Medical Texts:** Medical documents are inherently complex and vast in scope, making them particularly suitable for RAG-based approaches. Combined with the critical importance of information reliability in this field, focusing on healthcare becomes a highly significant choice.
- **Challenges in Real-World Applications:** Questions involving detailed patient information, such as medical histories and symptoms, often overwhelm retrieval systems, making it difficult to identify crucial diagnostic clues. Furthermore, in practical applications, patient conditions and individual characteristics vary widely. Differences in age, medical history, genetic factors, and lifestyle often lead to variations in treatment for the same disease. Providing inaccurate information in such scenarios can result in severe consequences. (Sohn et al., 2024)
- **Established Significance of BioNLP:** The prominence of the healthcare domain is evident from the long-standing ‘‘BioNLP’’ workshop, which has been held for over two decades.<sup>7</sup>
- **Emerging Trends in Healthcare RAG:** Efforts to improve RAG performance in the medical

domain have led to developments like Self-BioRAG, which leverages confidence scores. Its popularity and significant citation count highlight this field as a trending area of research. (Jeong et al., 2024) These points illustrate the rationale behind our focus on the healthcare domain.

#### A.6 Details of Evaluation Metrics

**Expected Calibration Error (ECE)** Calibration error metrics evaluate whether a model’s predicted probabilities align with actual accuracy in QA tasks. For instance, if a model assigns a 90% probability to an answer, the accuracy of such predictions should also be 90% for optimal calibration. Expected Calibration Error (ECE) (Naeini et al., 2015; Guo et al., 2017) quantifies this discrepancy by segmenting the predicted probability range into multiple bins and computing the difference between the predicted probability and the observed accuracy within each bin as follows:

$$\text{ECE} = \sum_{m=1}^M \frac{|B_m|}{n} |\text{acc}(B_m) - \text{conf}(B_m)| \quad (1)$$

Here,  $M$  denotes the number of bins,  $B_m$  represents the set of samples within bin  $m$ ,  $|B_m|$  is the number of samples in bin  $m$ , and  $n$  is the total number of samples.  $\text{acc}(B_m)$  refers to the accuracy within bin  $B_m$ , while  $\text{conf}(B_m)$  indicates the average confidence of predictions in bin  $m$ . ECE is computed as the weighted average of the absolute differences between the accuracy and confidence across bins, where the weights correspond to the proportion of samples in each bin.

**Adaptive Calibration Error (ACE)** ACE performs binning so that the number of samples in each bin remains constant. This approach ensures a more stable evaluation within each bin:

$$\text{ACE} = \frac{1}{KR} \sum_{k=1}^K \sum_{r=1}^R |\text{acc}(r, k) - \text{conf}(r, k)| \quad (2)$$

Here,  $K$  denotes the number of classes,  $R$  represents the number of bins,  $\text{acc}(r, k)$  indicates the accuracy in bin  $r$  for class  $k$ , and  $\text{conf}(r, k)$  denotes the confidence of predictions in the same bin and class.

<sup>7</sup>[https://aclweb.org/aclwiki/BioNLP\\_Workshop](https://aclweb.org/aclwiki/BioNLP_Workshop)

### A.7 The results using Expected Calibration Error (ECE)

The results using ECE are presented in Table 8 and Table 9. As discussed in Section 4.4, Proskurina et al. (2024) and Ulmer et al. (2022) have pointed out that ACE is a more suitable calibration error metric for multi-class classification problems, while ECE is better suited for binary classification. Nevertheless, we include ECE results for completeness and additional verification.

### A.8 Violin plot

Figures 5a and 5b present the violin plot results for Llama3.1 (8B), while Figures 5c and 5d show the results for Llama3.1 (70B). The Llama models exhibit notably low output probabilities for candidate answer choices when no supporting documents are inserted. Furthermore, even when explicitly inserting documents containing supporting evidence, the output probabilities do not improve significantly. This suggests that these models may strictly adhere to predefined instructions and struggle to incorporate additional contextual information.

### B Future Direction

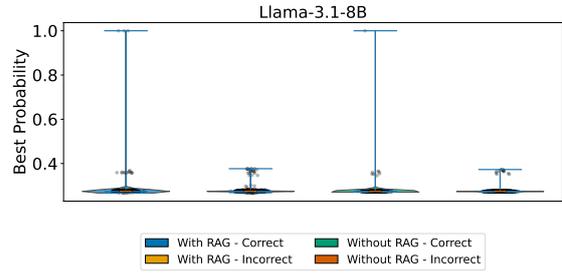
In this study, we used a dataset containing correct answer choices along with supporting rationale passages for QA tasks. In the future, it may be possible to focus on non-medical domains by drawing on previous work that semi-automatically generates questions using LLMs (Ozaki et al., 2024; Sakai et al., 2024). There are also studies on explanation generation (Ozaki et al., 2025a; Hayashi et al., 2024), which could inform the generation of supporting rationales.

Model	Pattern	ECE ↓			
		None	Ans1	Ans1-Oth2	Oth3
Llama2 (70B)	w/o RAG	0.02	–	–	–
	Pre-Q	–	0.22	0.25	0.08
	Aft-Q	–	0.24	0.21	0.12
	Aft-C	–	0.20	0.17	0.10
Llama3.1 (70B)	w/o RAG	0.20	–	–	–
	Pre-Q	–	0.14	0.14	0.14
	Aft-Q	–	0.14	0.14	0.14
	Aft-C	–	0.14	0.14	0.14
Llama3.1 (8B)	w/o RAG	0.03	–	–	–
	Pre-Q	–	0.02	0.02	0.01
	Aft-Q	–	0.02	0.02	0.01
	Aft-C	–	0.02	0.02	0.01
Meditron (70B)	w/o RAG	0.07	–	–	–
	Pre-Q	–	0.18	0.08	0.09
	Aft-Q	–	0.16	0.09	0.15
	Aft-C	–	0.15	0.07	0.09
PMC-Llama (13B)	w/o RAG	0.16	–	–	–
	Pre-Q	–	0.01	0.01	0.04
	Aft-Q	–	0.01	0.02	0.06
	Aft-C	–	0.02	0.05	0.05
Gemma2 (2B)	w/o RAG	0.20	–	–	–
	Pre-Q	–	0.24	0.22	0.21
	Aft-Q	–	0.23	0.21	0.21
	Aft-C	–	0.25	0.23	0.23
Phi-3.5 (3.8B)	w/o RAG	0.05	–	–	–
	Pre-Q	–	0.06	0.07	0.32
	Aft-Q	–	0.04	0.06	0.42
	Aft-C	–	0.04	0.13	0.41
Qwen2.5 (14B)	w/o RAG	0.12	–	–	–
	Pre-Q	–	0.09	0.09	0.12
	Aft-Q	–	0.07	0.07	0.18
	Aft-C	–	0.08	0.10	0.09
Qwen2.5 (72B)	w/o RAG	0.04	–	–	–
	Pre-Q	–	0.10	0.08	0.20
	Aft-Q	–	0.09	0.10	0.19
	Aft-C	–	0.10	0.06	0.24

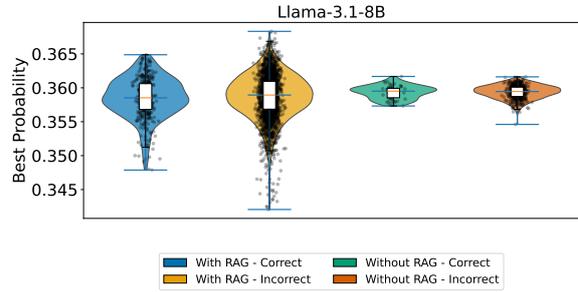
Table 8: The result of ECE using MedMCQA

Model	Pattern	ECE ↓			
		None	Ans1	Ans1-Oth2	Oth3
Llama2 (70B)	w/o RAG	0.12	-	-	-
	Pre-Q	-	0.30	0.31	0.14
	Aft-Q	-	0.30	0.31	0.14
	Aft-C	-	0.14	0.14	0.12
Llama3.1 (70B)	w/o RAG	0.02	-	-	-
	Pre-Q	-	0.10	0.10	0.10
	Aft-Q	-	0.10	0.10	0.10
	Aft-C	-	0.10	0.10	0.10
Llama3.1 (8B)	w/o RAG	0.24	-	-	-
	Pre-Q	-	0.23	0.22	0.23
	Aft-Q	-	0.22	0.23	0.23
	Aft-C	-	0.23	0.23	0.23
Meditron (70B)	w/o RAG	0.18	-	-	-
	Pre-Q	-	0.10	0.17	0.09
	Aft-Q	-	0.08	0.07	0.06
	Aft-C	-	0.04	0.04	0.07
PMC-Llama (13B)	w/o RAG	0.17	-	-	-
	Pre-Q	-	0.09	0.01	0.01
	Aft-Q	-	0.09	0.01	0.02
	Aft-C	-	0.02	0.04	0.05
Gemma2 (2B)	w/o RAG	0.07	-	-	-
	Pre-Q	-	0.04	0.02	0.02
	Aft-Q	-	0.02	0.01	0.01
	Aft-C	-	0.09	0.05	0.07
Phi-3.5 (3.8B)	w/o RAG	0.48	-	-	-
	Pre-Q	-	0.11	0.11	0.58
	Aft-Q	-	0.10	0.28	0.49
	Aft-C	-	0.15	0.32	0.46
Qwen2.5 (14B)	w/o RAG	0.16	-	-	-
	Pre-Q	-	0.04	0.05	0.18
	Aft-Q	-	0.03	0.05	0.27
	Aft-C	-	0.07	0.03	0.24
Qwen2.5 (72B)	w/o RAG	0.06	-	-	-
	Pre-Q	-	0.11	0.03	0.25
	Aft-Q	-	0.08	0.11	0.18
	Aft-C	-	0.09	0.07	0.14

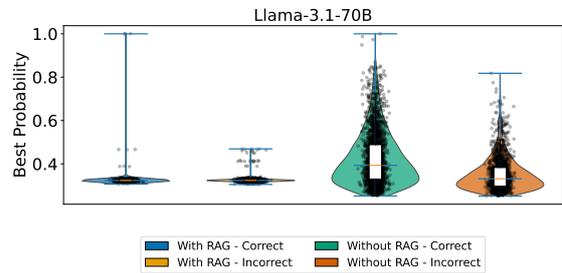
Table 9: The result of ECE using PubMedQA



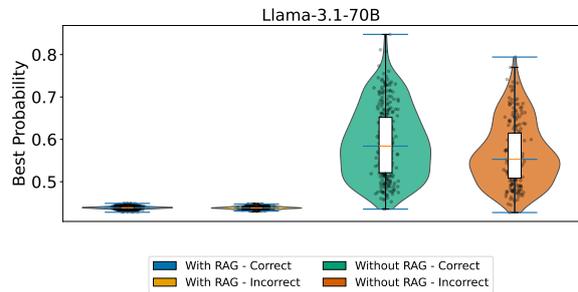
(a) The result of Llama3.1 (8B) using MedMCQA.



(b) The result of Llama3.1 (8B) using PubMedQA.



(c) The result of Llama3.1 (70B) using MedMCQA.



(d) The result of Llama3.1 (70B) using PubMedQA.

Figure 5: Results from MedMCQA and PubMedQA using Llama3.1.

## B.1 Error Analysis

Figure 6 presents a plot illustrating the types of errors made on incorrectly answered questions. In PubMedQA, the answer choices consist of three options: yes, no, and maybe, allowing for detailed error analysis. Each bin represents the gold answer, and the plot visualizes the distribution of incorrect predictions for each question. The colors within the plot indicate how the model misclassified the answers.



Figure 6: Error analysis on PubMedQA: This figure illustrates how the model misclassified answers in relation to the correct ones. We gather the questions the model got wrong, group the items that actually had the correct answer (gold answer) into bins, and use colors to show how the model erred.

## B.2 Prompts

Below are examples of prompts with and without RAG. When RAG is applied, three patterns—Pre-Question, After-Question, and After-Choice—are used in our study.

### Prompt without RAG

You are a helpful medical expert, and your task is to answer a multi-choice medical question using the relevant documents. Please first think step-by-step and then choose the answer from the provided options. Your responses will be used for research purposes only, so please have a definite answer.

Here is the question:  
{question}

Here are the potential choices:  
{choice0}  
{choice1}  
{choice2}  
{choice3}

Answer:

### Prompt with RAG

Here are the relevant documents: (Pre-Question)  
{context}

You are a helpful medical expert, and your task is to answer a multi-choice medical question using the relevant documents. Please first think step-by-step and then choose the answer from the provided options. Your responses will be used for research purposes only, so please have a definite answer.

Here are the relevant documents: (After-Question)  
{context}

Here is the question:  
{question}

Here are the relevant documents: (After-Choice)  
{context}

Here are the potential choices:  
{choice0}  
{choice1}  
{choice2}  
{choice3}

Answer: