

---

# Just rephrase it! Uncertainty estimation in closed-source language models via multiple rephrased queries

---

Anonymous Author(s)

Affiliation

Address

email

## Abstract

1 We explore estimating the uncertainty of closed-source LLMs via multiple rephrasings of an original base query. Specifically, we ask the model, multiple rephrased questions, and use the similarity of the answers as an estimate of uncertainty. We  
2 diverge from previous work in i) providing rules for rephrasing that are simple to memorize and use in practice ii) proposing a theoretical framework for why  
3 multiple rephrased queries obtain calibrated uncertainty estimates. Our method  
4 demonstrates significant improvements in the calibration of uncertainty estimates compared to the baseline and provides intuition as to how query strategies should  
5 be designed for optimal test calibration.  
6  
7  
8  
9

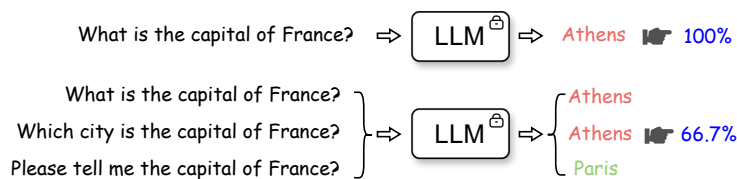


Figure 1: **Multiple rephrased queries for uncertainty estimation.** Querying a closed-source LLM only once with a base query may yield an incorrect top-1 prediction with 100% confidence to this singular prediction. Querying the model multiple times with rephrased versions of the base query produces different answers equivalent to 66.6% confidence.

## 10 1 Introduction

11 Closed-source LLMs are prone to generating highly convincing but false information, a problem  
12 known as "hallucinating" (Huang et al., 2023; Ji et al., 2023). It is folk wisdom that one approach  
13 for estimating LLM uncertainty, even with such limited access to the model, is to query it multiple  
14 times (Wang et al., 2022; Xiong et al., 2023). This approach is based on the premise that LLM-  
15 generated text is frequently stochastic by design, as the next generated token is chosen through  
16 nucleus sampling (Holtzman et al., 2019) or top-k decoding (Fan et al., 2018; Radford et al., 2019).  
17 Wang et al. (2022) and Xiong et al. (2023) proposed to use the consistency of multiple answers as an  
18 estimate of uncertainty. Xiong et al. (2023) furthermore proposed to add "noise" to the base query at  
19 each repetition, through misleading hints.

20 In this work, we delve deeply in, refine, and theoretically analyze multiple queries for uncertainty  
21 estimation. Given a base query, we restrict ourselves to submitting rephrased versions of the base




Figure 2: We plot the AUROC averaged over all models for each dataset and for each uncertainty estimation method. We observe that top-k improves over the naive top-1 decoding. Furthermore, the best rephrasing method (denoted as rephrase\*) improves the AUROC significantly in all cases.

22 query to an LLM, checking the consistency of the answers, and using the result as an estimate of  
 23 uncertainty. Concretely our contributions are the following:

- 24 • We test four simple strategies for creating multiple rephrased queries, and find that they  
 25 result in significant calibration gains over baselines.
- 26 • We propose a theoretical model for multiple rephrased queries on a simplified top-1 and top-  
 27 k (Holtzman et al., 2019) decoding setting. Given multiple rephrased queries, our analysis  
 28 shows that i) it is possible to recover the probability of the answer under the inaccessible  
 29 categorical distribution of the LLM ii) top-k decoding then simply tempers our uncertainty.  
 30 Crucially we show empirically that our uncertainty estimates are close to what could be  
 31 obtained when having access to the last layer logits.

## 32 2 Rephrasing drastically improves calibration for top-1 decoding

33 Let  $f : \mathcal{X} \rightarrow \mathcal{Y}$  be an LLM which takes  $\mathbf{x}$  an input query in the form of a multiple choice  
 34 question, and outputs  $y$ , an answer. We first consider top-1 decoding such that the answers of the  
 35 LLM are deterministic. We consider randomized transformations of the base query  $\mathcal{T}(\mathbf{x}) \sim \tau$   
 36 in the form of rephrasings of the query, and the most probable answer under the transformations  
 37  $A = \operatorname{argmax}_i \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = i)$ . In a multiple choice question setting (which can be seen as a  
 38 multi-class classification problem), we will use  $A$  as the predicted class and

$$p_A(\mathbf{x}) = \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A),$$

39 as our confidence about this prediction (here the predicted class coincides with a predicted token  
 40 denoting this class). We consider four types of rephrasings, with an increasing level of modification to  
 41 the original query: (1) reword: replacing words with synonyms; (2) rephrase: modifies the structure of  
 42 the original query; (3) paraphrase: reconstructs the original query; (4) expansion: elaborate the query.  
 43 In general, we perform the rephrasings with a separate instance of the same model that responds  
 44 to the queries. We estimate  $p_A(\mathbf{x})$  using Monte Carlo sampling with 10 draws from  $\mathcal{T}(\mathbf{x}) \sim \tau$  to  
 45 estimate uncertainty with our method unless stated otherwise.

46 We used three different models, the Llama-2 7B model, the Llama-2 13B model (Touvron et al., 2023)  
 47 and the Mistral 7B model (Jiang et al., 2023). We tested our framework on three multiple choice  
 48 tasks: ARC-Challenge, ARC-Easy (Clark et al., 2018), and Openbookqa (Mihaylov et al., 2018).  
 49 Following Kojima et al. (2022), we extract the answer from LLM-generated texts by looking at the  
 50 first appearance of A/B/C/D. To test for calibration we used standard calibration metrics, including  
 51 the ECE and TACE (Naeini et al., 2015), Brier score (Murphy, 1973) and AUROC (Murphy, 2012).

52 We plot the AUROC results of all methods averaged over all models for each dataset in Figure 2. We  
 53 see that the best rephrasing method outperforms top-1 (naive) and top-k decoding as well as the hint  
 54 based rephrasing approach. In Appendices E and D we show that we also match or outperform Chain  
 55 of Thought (CoT) prompting and Temperature Sampling Wei et al. (2022).

## 56 3 Rephrasing works as well as having access to the last layer logits

57 We now derive a proposition that elucidates why  $p_A(\mathbf{x})$  results in calibrated estimates of uncertainty.

Table 1: Comparisons between our rephrasing methods and white-box logit uncertainty estimation. We see that our rephrasing methods achieve similar calibration to what would be achieved if we had access to last layer logits. This is evident both in the AUROC and TACE as well as the Brier score, which also accounts for accuracy.

Dataset	Model	Method	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$
ARC-C	Mistral-7B	logits	0.742	0.252	0.075	0.503	0.741
		expansion	0.602	0.133	0.099	0.509	0.847
	Llama-2-7B	logits	0.483	0.362	0.168	0.853	0.621
		expansion	0.373	0.112	0.153	0.778	0.687
	Llama-2-13B	logits	0.508	0.132	0.141	0.704	0.669
		reword	0.445	0.084	0.119	0.714	0.721
ARC-E	Mistral-7B	logits	0.866	0.128	0.037	0.264	0.818
		reword	0.753	0.045	0.062	0.297	0.931
	Llama-2-7B	logits	0.672	0.190	0.098	0.493	0.779
		rephrase	0.535	0.131	0.117	0.603	0.830
	Llama-2-13B	logits	0.617	0.060	0.094	0.498	0.763
		expansion	0.524	0.078	0.12	0.552	0.893
OBQA	Mistral-7B	logits	0.655	0.298	0.085	0.602	0.705
		reword	0.552	0.105	0.102	0.592	0.796
	Llama-2-7B	logits	0.478	0.277	0.147	0.758	0.642
		expansion	0.362	0.083	0.138	0.775	0.678
	Llama-2-13B	logits	0.418	0.168	0.135	0.723	0.650
		rephrase	0.428	0.095	0.14	0.729	0.73

58 **Proposition 3.1.** Let  $f : \mathcal{X} \rightarrow \mathcal{Y}$  be an LLM,  $\mathbf{x}$  is a base query and  $\mathcal{T}(\mathbf{x}) \sim \tau$  is some randomized  
59 transformation of the base query. Let

$$p_A(\mathbf{x}) = \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A), \quad (1)$$

60 be the probability of sampling the most probable answer  $A \in \mathcal{Y}$  under transformations  $\mathcal{T}(\mathbf{x}) \sim \tau$ .  
61 Let  $\mathbf{z}_{mean} + \epsilon_{rephrase}$  be the latent representation of  $\mathbf{x}$  under  $\mathcal{T}(\mathbf{x})$  at the final LLM layer, where  
62  $\mathbf{z}_{mean}$  is the mean representation and  $\epsilon_{rephrase}$  is some additive noise. Let  $\mathbf{w}$  be the separating  
63 hyperplane between the most probable answer  $A$  and the second most probable answer  $B$ . Assuming  
64 that  $\mathbf{w}^\top \epsilon_{rephrase} \sim \rho$  follows a logistic distribution with  $\mu = 0$  and  $s = 1$  then

$$p_A(\mathbf{x}) = p(A|\mathbf{z}_{mean}, f) \quad (2)$$

65 where  $p(A|\mathbf{z}_{mean}, f)$  is the probability of  $A$  given  $\mathbf{z}_{mean}$  under the categorical distribution of the  
66 final layer.

67 We prove the above for the binary case of two classes  $A$  and  $B$  in Appendix C, but expect that it should  
68 be sufficiently informative in multi-class settings when  $A, B$  are much more probable than other  
69 classes. A crucial assumption for recovering well-calibrated predictions is that  $\mathbf{w}^\top \epsilon_{rephrase} \sim \rho$   
70 follows a logistic distribution with  $\mu = 0$  and  $s = 1$ . We test this assumption by computing the  
71 cumulative of  $\rho$  for our different experimental setups. In Figure 3c we find and plot the empirical  
72 cumulative using a Kolmogorov-Smirnov test (Smirnov, 1948) and  $S = 100$  MC samples of  $\rho$  for  
73 Mistral-7B, ARC-Challenge, and the ‘‘expansion’’ rephrasing method. We see that the indeed the  
74 cumulative is approximately logistical validating our prediction (the confidence bands cover different  
75 queries  $\mathbf{x}$ ). In Table 1 we use the logits of the answers as an oracle white-box uncertainty estimate.  
76 Specifically, we apply the softmax function and use the probability of the most probable class as our  
77 estimate of uncertainty. We compare the results of this method with the best rephrasing method (in  
78 terms of Brier). We observe that our uncertainty estimates that are similar to what we would get if we  
79 had access to the last layer logits.

#### 80 4 For top-k decoding, rephrasing tempers predictive uncertainty

81 In practice, the assumptions of the above proposition are too restrictive. In particular, decoding in  
82 LLMs is performed with top-k decoding or nucleus sampling instead of top-1 decoding. Furthermore




Figure 3: We plot the distribution of  $p_A(\mathbf{x})$  for the case of top-k decoding with and without rephrasing, for all datasets, models, and rephrasing methods. We see that rephrasing primarily acts to temper the probability of the most probable class  $A$ , thus making the model less confident and possibly better calibrated. We also plot the logistic (blue), and empirical cdf (red) for  $\mathbf{w}^\top \epsilon_{\text{rephrase}} \sim \rho$  for Mistral-7B, ARC-Challenge, and the “expansion” rephrasing method for top-1 decoding.  $\rho$  is often close to a logistic distribution.

83 while for an oracle choice of the rephrasing intensity the modeling assumption that  $\mathbf{w}^\top \epsilon_\eta \sim \rho$  follows  
 84 a logistic distribution with  $\mu = 0$  and  $s = 1$  might be correct, in general, the variance of the noise  
 85 in latent space is unknown. It is thus illustrative to consider an extension of our toy model. The  
 86 following proposition explores these extensions.

87 **Proposition 4.1.** *Let  $g : \mathbb{R}^{d_\eta} \rightarrow \mathcal{Y}$  be the final encoder layer of an LLM,  $\mathbf{x}$  is a base query and*  
 88  *$\mathcal{T}(\mathbf{x}) \sim \tau$  is some randomized transformation of the base query. Let*

$$p_A(\mathbf{x}) = \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A), \quad (3)$$

89 *be the probability of sampling the most probable answer  $f(\mathbf{x}) = A \in \mathcal{Y}$  under transformations*  
 90  *$\mathcal{T}(\mathbf{x}) \sim \tau$ . Let  $\mathbf{z}_{\text{mean}} + \epsilon_{\text{topk}} + \epsilon_{\text{rephrase}}$  be the latent representation of  $\mathbf{x}$  under  $\mathcal{T}(\mathbf{x})$  at the*  
 91 *final LLM layer, where  $\mathbf{z}_{\text{mean}}$  is the mean representation and  $\epsilon_{\text{topk}}$  is additive noise resulting from*  
 92 *the top-k decoding and  $\epsilon_{\text{rephrase}}$  is additive noise resulting from the rephrasings  $\mathcal{T}(\mathbf{x})$ . Assuming*  
 93 *that  $\mathbf{w}^\top (\epsilon_{\text{topk}} + \epsilon_{\text{rephrase}}) \sim \rho$  approximately follows a logistic distribution with  $\mu = 0$  and*  
 94  *$s = \sqrt{s_{\text{topk}}^2 + s_{\text{rephrase}}^2}$  then*

$$p_A(\mathbf{x}) \approx 0.5 + \frac{1}{\sqrt{s_{\text{topk}}^2 + s_{\text{rephrase}}^2}} (p(A|\mathbf{z}_{\text{mean}}, f) - 0.5) \quad (4)$$

95 where  $p(A|\mathbf{z}_{\text{mean}}, f)$  is the probability of  $A$  given  $\mathbf{z}_{\text{mean}}$  under the categorical distribution of  $g$ .

96 The approximation relies on linearizing the involved functions, however, it is illustrative of the  
 97 effect of both  $s_{\text{topk}}^2$  and  $s_{\text{rephrase}}^2$ . In particular, we see that both  $s_{\text{topk}}^2$  and  $s_{\text{rephrase}}^2$  act to temper  
 98 the probability  $p(A|\mathbf{z}_{\text{mean}}, f)$  under the categorical distribution of  $g$ . This highlights why using  
 99 rephrasings with an appropriate temperature might improve the calibration in downstream tasks. In  
 100 previous works, tempering of the categorical distribution has been found to significantly improve the  
 101 calibration of deep neural networks (Guo et al., 2017).

102 Figure 3 shows that in accordance with proposition 4.1 rephrasing acts primarily to temper the  
 103 probability of the top class. In our detailed results in Appendix E, this often improves calibration  
 104 significantly in terms of ECE, and AUROC especially for smaller models.

## 105 5 Discussion

106 We conducted a thorough analysis of rephrased queries as a method for obtaining calibrated predic-  
 107 tions from closed-source LLM models. Notably, we found that two simple methods; making the  
 108 query more verbose, and substituting words with their synonyms, provide a straightforward means of  
 109 identifying false positives. The appeal of our approach lies in its practicality, as it requires only basic  
 110 language and arithmetic skills by the end user to obtain meaningful uncertainty estimates. Exciting  
 111 future directions include learning optimal rephrasing rules in a data-driven manner, to be used in  
 112 conjunction with a rephrasing LLM. While we tested on the multiple choice question setting for ease  
 113 of evaluation, we expect our results to also hold for open-ended text generation.

## 114 References

- 115 Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and  
116 Oyvind Tafjord. Think you have solved question answering? try arc, the ai2 reasoning challenge.  
117 *arXiv:1803.05457v1*, 2018.
- 118 Angela Fan, Mike Lewis, and Yann Dauphin. Hierarchical neural story generation. In *Proceedings*  
119 *of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long*  
120 *Papers)*, pp. 889–898, 2018.
- 121 Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural  
122 networks. In *International Conference on Machine Learning*, pp. 1321–1330. PMLR, 2017.
- 123 Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. The curious case of neural text  
124 degeneration. In *International Conference on Learning Representations*, 2019.
- 125 Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong  
126 Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. A survey on hallucination in large language  
127 models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*,  
128 2023.
- 129 Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan Su, Yan Xu, Etsuko Ishii, Ye Jin Bang,  
130 Andrea Madotto, and Pascale Fung. Survey of hallucination in natural language generation. *ACM*  
131 *Computing Surveys*, 55(12):1–38, 2023.
- 132 Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot,  
133 Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al.  
134 Mistral 7b. *arXiv preprint arXiv:2310.06825*, 2023.
- 135 Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. Large  
136 language models are zero-shot reasoners. *Advances in neural information processing systems*, 35:  
137 22199–22213, 2022.
- 138 Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. Can a suit of armor conduct  
139 electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*,  
140 2018.
- 141 Allan H Murphy. A new vector partition of the probability score. *Journal of Applied Meteorology*  
142 *and Climatology*, 12(4):595–600, 1973.
- 143 Kevin P Murphy. *Machine learning: a probabilistic perspective*. MIT Press, 2012.
- 144 Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. Obtaining well calibrated prob-  
145 abilities using Bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*,  
146 2015.
- 147 Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language  
148 models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- 149 Nikolay Smirnov. Table for estimating the goodness of fit of empirical distributions. *The annals of*  
150 *mathematical statistics*, 19(2):279–281, 1948.
- 151 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay  
152 Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation  
153 and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- 154 Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdh-  
155 ery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models.  
156 *arXiv preprint arXiv:2203.11171*, 2022.
- 157 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny  
158 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in*  
159 *neural information processing systems*, 35:24824–24837, 2022.
- 160 Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. Can llms  
161 express their uncertainty? an empirical evaluation of confidence elicitation in llms. *arXiv preprint*  
162 *arXiv:2306.13063*, 2023.

Table 3: Different Prompts to Induce Consistency Confidence

Consistency Confidence Inducing Prompts	
Hint1	I think the answer should be
Hint2	I read online the answer is
Hint3	I vaguely remember the answer is

163 **A Prompt template**

164 We present our prompt template for initiating rephrases with a one-shot example in Table 2. Note  
 165 that we only present and rephrase questions without revealing choices, to reduce unnecessary bias to  
 166 rephrases when presented with answer choices.

Method	Prompt
reword	[INST]Reword the following question: George wants to warm his hands quickly by rubbing them. Which skin surface will produce the most heat? Respond with the reworded question only: [\INST] George seeks to heat his hands swiftly by rubbing them. Which skin area will generate the maximum heat? [INST]Reword the following question: {question} Respond with the reworded question only: [\INST]
rephrase	[INST]Rephrase the following question: George wants to warm his hands quickly by rubbing them. Which skin surface will produce the most heat? Respond with the rephrased question only: [\INST] What type of skin texture on George’s hands would generate the most heat through rapid rubbing to warm them effectively? [INST]Rephrase the following question: {question} Respond with the rephrased question only: [\INST]
paraphrase	[INST]Semantically paraphrase the following question: George wants to warm his hands quickly by rubbing them. Which skin surface will produce the most heat? Respond with the semantically paraphrased question only: [\INST] How can George induce the highest thermal output by briskly rubbing his hands, and which part of the skin would be most effective? [INST]Semantically paraphrase the following question: {question} Respond with the semantically paraphrased question only: [\INST]
expansion	[INST] Expand the following question with additional context: George wants to warm his hands quickly by rubbing them. Which skin surface will produce the most heat? Respond with the expanded question only: [\INST] In the context of seeking immediate relief from the biting cold and understanding the mechanisms behind heat generation through friction, what type of skin texture on George’s hands would most effectively generate heat by rapid rubbing? [INST]Expand the following question with additional context: {question} Respond with the expanded question only: [\INST]

Table 2: Prompt templates for one-shot rephrasing, with rephrasing methods listed on the left and corresponding prompt on the right. The user instructions are colored in blue and surrounded by the instruction token, whereas model response demonstrations are colored in orange.

167 We followed the instructions in Xiong et al. (2023) to generate "hint" based rephrasings. Specifically,  
 168 to generate a rephrased query given a base query, we appended one of the following three weak  
 169 claims (as they found weak claims outperform other types of hints) together with a random class  
 170 from the available ones.

Method	Question
original	What part of the digestive system first causes chemical changes to food? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
reword	Which region of the gastrointestinal tract initiates the initial chemical modifications to food intake? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
rephrase	In what region of the digestive system does the food undergo its initial chemical transformations? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
paraphrase	At what point in the digestive process do initial chemical transformations of food occur and which section of the system carries out this function? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
expansion	Considering the intricate process by which our bodies break down and absorb nutrients from food, which specific organ or region within the digestive system initiates the essential biochemical transformations through enzyme secretion and the beginning of the digestion process? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.

Table 4: Rephrasing examples generated by Mistral-7B, with rephrasing methods listed on the left and corresponding rephrases on the right.

Method	Question
original	What part of the digestive system first causes chemical changes to food? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
reword	What section of the digestive system initiates chemical alterations to sustenance? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
rephrase	Which portion of the digestive system initially catalyzes the biochemical transformation of ingested sustenance? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
paraphrase	Which digestive organ releases enzymes that initiate chemical breakdown within ingested sustenences, leading to nutrient extraction and energy release? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
expansion	In the context of the digestive process and the breakdown of nutrients, which portion of the digestive system initiates the chemical transformations that result in the nutrient absorption and energy production? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.

Table 5: Rephrasing examples generated by Llama2-7B, with rephrasing methods listed on the left and corresponding rephrases on the right.

## 171 B Rephrase generations

172 Here, we present additional generated rephrasings by Mistral-7B, Llama2-7B and Llama2-13B in  
173 Table 4, Table 5 and Table 6.

## 174 C Additional Proofs

175 **Proposition C.1.** *Let  $f : \mathcal{X} \rightarrow \mathcal{Y}$  be an LLM,  $x$  is a base query and  $\mathcal{T}(x) \sim \tau$  is some randomized*  
176 *transformation of the base query. Let*

$$p_A(x) = \mathbb{P}(f(\mathcal{T}(x)) = A), \quad (5)$$

Method	Question
original	What part of the digestive system first causes chemical changes to food? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
reword	Which section of the gastrointestinal tract initiates the chemical transformation of sustenance? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
rephrase	In which section of the digestive system does the initial chemical breakdown of food occur? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
paraphrase	In the digestive process, where do crucial transformations initially occur to break down nutrients? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.
expansion	Taking into account that human digestive system's several organs coordinate to breakdown, absorb, and expel waste, which part of the gastrointestinal system would have the most significant logic-based influence on the breakdown of food into usable components, prior to the nutrient absorption? A. Teeth in the mouth. B. Saliva in the mouth. C. Enzymes in the stomach. D. Enzymes in the small intestine.

Table 6: Rephrasing examples generated by Llama2-13B, with rephrasing methods listed on the left and corresponding rephrases on the right.

177 *be the probability of sampling the most probable answer  $A \in \mathcal{Y}$  under transformations  $\mathcal{T}(\mathbf{x}) \sim \tau$ .*  
178 *Let  $\mathbf{z}_{mean} + \epsilon_{rephrase}$  be the latent representation of  $\mathbf{x}$  under  $\mathcal{T}(\mathbf{x})$  at the final LLM layer, where*  
179  *$\mathbf{z}_{mean}$  is the mean representation and  $\epsilon_{rephrase}$  is some additive noise. Let  $\mathbf{w}$  be the separating*  
180 *hyperplane between the most probable answer  $A$  and the second most probable answer  $B$ . Assuming*  
181 *that  $\mathbf{w}^\top \epsilon_{rephrase} \sim \rho$  follows a logistic distribution with  $\mu = 0$  and  $s = 1$  then*

$$p_A(\mathbf{x}) = p(A|\mathbf{z}_{mean}, f) \quad (6)$$

182 *where  $p(A|\mathbf{z}_{mean}, f)$  is the probability of  $A$  given  $\mathbf{z}_{mean}$  under the categorical distribution of the*  
183 *final layer.*

184 *Proof.* We first analyze the categorical distribution, resulting from applying the softmax on the final  
185 layer logits. In the binary classification case given a top-1 class prediction  $A$ , the softmax probability  
186 of this class is

$$\begin{aligned} p(A|\mathbf{x}, f) &= \frac{e^{\mathbf{w}_A^\top \mathbf{z} + b_A}}{e^{\mathbf{w}_A^\top \mathbf{z} + b_A} + e^{\mathbf{w}_B^\top \mathbf{z} + b_B}} \\ &= \frac{1}{1 + e^{-(\mathbf{w}_A + b_A - \mathbf{w}_B - b_B)^\top \mathbf{z}}} = \frac{1}{1 + e^{-(\mathbf{w}^\top \mathbf{z} + b)}}. \end{aligned} \quad (7)$$

The above simply corresponds to the folk knowledge that a softmax layer with two classes is equivalent to a single separating hyperplane that assigns classes based on the rule  $\text{sign}(\mathbf{w}^\top \mathbf{z} + b)$ , specifically

$$g(\mathbf{z}) = \begin{cases} A & \text{if } (\mathbf{w}^\top \mathbf{z} + b) > 0, \\ B & \text{otherwise.} \end{cases}$$

187 After establishing that the softmax layer is equivalent to this single separating hyperplane, let us  
188 relate  $p_A(\mathbf{x})$  to  $\mathbf{w}^\top \mathbf{z} + b$ . We have

$$\begin{aligned} p_A(\mathbf{x}) &= \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A) \\ &= \mathbb{P}(\mathbf{w}^\top (\mathbf{z}_{mean} + \epsilon_{rephrase}) + b > 0) \\ &= \mathbb{P}(\mathbf{w}^\top \mathbf{z}_{mean} + \mathbf{w}^\top \epsilon_{rephrase} + b > 0) \\ &= \mathbb{P}(Z > -\mathbf{w}^\top \mathbf{z}_{mean} - b) \\ &= 1 - \mathbb{P}(Z < -\mathbf{w}^\top \mathbf{z}_{mean} - b) \\ &= 1 - F(-\mathbf{w}^\top \mathbf{z}_{mean} - b) \end{aligned} \quad (8)$$



189 Then  $F(-\mathbf{w}^\top \mathbf{z}_{mean} - b) = 1 - p_A \iff \mathbf{w}^\top \mathbf{z}_{mean} + b = -F^{-1}(1 - p_A)$ . We substitute this  
 190 result to 7, assume that  $F$  is the cumulative of the logistic distribution with  $\mu = 0$  and  $s = 1$  and get

$$p(A|\mathbf{z}_{mean}, f) = \frac{1}{1 + e^{F^{-1}(1-p_A)}} \quad (9)$$

$$= \frac{1}{1 + e^{-F^{-1}(p_A(\mathbf{x}))}} \quad (10)$$

$$= p_A(\mathbf{x}) \quad (11)$$

191 In the second line we used the fact that the inverse cumulative  $F^{-1}$  of the logistic distribution is  
 192 symmetric around 0.5. In the third line we use the fact that  $\frac{1}{1+e^{-x}}$  is the cumulative of the logistic  
 193 with  $\mu = 0$  and  $s = 1$ . Thus  $p(A|\mathbf{z}_{mean}, f) = F(F^{-1}(p_A(\mathbf{x}))) \iff p(A|\mathbf{z}_{mean}, f) = p_A(\mathbf{x})$

194 A technical point remains. Even though in the previous we can assume that  $g(\mathbf{z}_{mean}) =$   
 195  $A$  (that  $\mathbf{z}_{mean}$  results in the most probable class) by definition, we still need to show that  
 196  $A = \operatorname{argmax}_i \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = i) \iff g(\mathbf{z}_{mean}) = A$ . This means that for a closed-  
 197 source LLM we can identify the (unknown) top-1 class  $A$  through Monte Carlo sampling ( $A =$   
 198  $\operatorname{argmax}_i \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = i)$ ).

$$\begin{aligned} A = \operatorname{argmax}_i \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = i) &\iff \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A) > \frac{1}{2} \\ &\iff \mathbb{P}(\mathbf{w}^\top (\mathbf{z}_{mean} + \epsilon_{rephrase}) + b \geq 0) > \frac{1}{2} \\ &\iff \mathbb{P}(\mathbf{w}^\top \mathbf{z}_{mean} + \mathbf{w}^\top \epsilon_{rephrase} + b \geq 0) > \frac{1}{2} \\ &\iff \mathbb{P}(Z \geq -\mathbf{w}^\top \mathbf{z}_{mean} - b) > \frac{1}{2} \\ &\iff \mathbb{P}(Z \leq \mathbf{w}^\top \mathbf{z}_{mean} + b) > \frac{1}{2} \\ &\iff \mathbf{w}^\top \mathbf{z}_{mean} + b > 0 \\ &\iff g(\mathbf{z}_{mean}) = A \end{aligned} \quad (12)$$

199 where we use the assumption that  $Z$  follows a logistic distribution with  $\mu = 0$  and  $s = 1$ .  $\square$

200 **Proposition C.2.** Let  $g : \mathbb{R}^{d_n} \rightarrow \mathcal{Y}$  be the final encoder layer of an LLM,  $\mathbf{x}$  is a base query and  
 201  $\mathcal{T}(\mathbf{x}) \sim \tau$  is some randomized transformation of the base query. Let

$$p_A(\mathbf{x}) = \mathbb{P}(f(\mathcal{T}(\mathbf{x})) = A), \quad (13)$$

202 be the probability of sampling the most probable answer  $f(\mathbf{x}) = A \in \mathcal{Y}$  under transformations  
 203  $\mathcal{T}(\mathbf{x}) \sim \tau$ . Let  $\mathbf{z}_{mean} + \epsilon_{topk} + \epsilon_{rephrase}$  be the latent representation of  $\mathbf{x}$  under  $\mathcal{T}(\mathbf{x})$  at the  
 204 final LLM layer, where  $\mathbf{z}_{mean}$  is the mean representation and  $\epsilon_{topk}$  is additive noise resulting from  
 205 the top- $k$  decoding and  $\epsilon_{rephrase}$  is additive noise resulting from the rephrasings  $\mathcal{T}(\mathbf{x})$ . Assuming  
 206 that  $\mathbf{w}^\top (\epsilon_{topk} + \epsilon_{rephrase}) \sim \rho$  approximately follows a logistic distribution with  $\mu = 0$  and  
 207  $s = \sqrt{s_{topk}^2 + s_{rephrase}^2}$  then

$$p_A(\mathbf{x}) \approx 0.5 + \frac{1}{\sqrt{s_{topk}^2 + s_{rephrase}^2}} (p(A|\mathbf{z}_{mean}, f) - 0.5) \quad (14)$$

208 where  $p(A|\mathbf{z}_{mean}, f)$  is the probability of  $A$  given  $\mathbf{z}_{mean}$  under the categorical distribution of  $g$ .

209 *Proof.* We first claim that the sum of two logistic distributions  $(\mu_1, s_1)$  and  $(\mu_2, s_2)$  is approximately  
 210 logistic with  $(\mu_1 + \mu_2, \sqrt{s_1^2 + s_2^2})$  by claiming that logistic distributions are approximately Gaussian.  
 211 Then considering that  $p(A|\mathbf{z}_{mean}, f) = \frac{1}{1 + e^{F^{-1}(1-p_A(\mathbf{x}))}}$  we can write

$$\begin{aligned} p(A|\mathbf{z}_{mean}, f) &= \frac{1}{1 + e^{F^{-1}(1-p_A(\mathbf{x}))}} = \frac{1}{1 + e^{-F^{-1}(p_A(\mathbf{x}))}} \\ &= 0.5 + \frac{1}{4} F^{-1}(p_A(\mathbf{x})) = 0.5 + \frac{1}{4} \sqrt{s_{topk}^2 + s_{rephrase}^2} (p_A(\mathbf{x}) - 0.5) \end{aligned} \quad (15)$$

Table 7: Comparisons between our best rephrasing method and CoT. Our rephrasing method obtains comparable results to CoT in terms of Brier score and other calibration metrics.

Dataset	Model	Method	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$
ARC-C	Mistral-7B	CoT	0.725	0.173	0.071	0.439	0.719
		expansion	0.602	0.133	0.099	0.509	0.847
	Llama-2-7B	CoT	0.407	0.205	0.151	0.783	0.696
		expansion	0.373	0.112	0.153	0.778	0.687
	Llama-2-13B	CoT	0.369	0.137	0.148	0.782	0.729
		reword	0.445	0.084	0.119	0.714	0.721
ARC-E	Mistral-7B	CoT	0.857	0.07	0.037	0.211	0.829
		reword	0.753	0.045	0.062	0.297	0.931
	Llama-2-7B	CoT	0.482	0.104	0.116	0.624	0.842
		rephrase	0.535	0.131	0.117	0.603	0.830
	Llama-2-13B	CoT	0.463	0.097	0.124	0.61	0.884
		expansion	0.524	0.078	0.12	0.552	0.893
OBQA	Mistral-7B	CoT	0.662	0.153	0.083	0.501	0.762
		reword	0.552	0.105	0.102	0.592	0.796
	Llama-2-7B	CoT	0.39	0.185	0.145	0.805	0.713
		expansion	0.362	0.083	0.138	0.775	0.678
	Llama-2-13B	CoT	0.37	0.166	0.153	0.801	0.683
		rephrase	0.428	0.095	0.14	0.729	0.73

212 In the first line we first considered that  $F^{-1}$  for the logistic is symmetric thus  $F^{-1}(1 - p_A(\mathbf{x})) =$   
213  $-F^{-1}(p_A(\mathbf{x}))$ . In the second line we first do a first order Taylor expansion around 0 on  $\frac{1}{1+e^{-x}}$  and  
214 then a first order Taylor expansion around 0.5 on  $F^{-1}$ .  $\square$

## 215 D Additional comparisons with CoT

216 We compare with Chain-of-Thought [Wei et al. \(2022\)](#) for uncertainty estimation and plot the results  
217 in Table 7. We find that we get competitive results with CoT. At the same time our method is  
218 significantly easier and more natural to implement for humans interacting via text with an LLM. In  
219 CoT one needs to first obtain a sequence of reasoning steps. These should then be used as additional  
220 context when asking an LLM to answer again the base question. By contrast we propose a simple  
221 one step process of rephrasing the base question.

## 222 E Additional results

223 We used three different models, the Llama-2 7B model, the Llama-2 13B model ([Touvron et al.,](#)  
224 [2023](#)) and the Mistral 7B model ([Jiang et al., 2023](#)). We tested our framework on three multiple  
225 choice tasks of different difficulty namely ARC-Challenge, ARC-Easy ([Clark et al., 2018](#)), and  
226 Openbookqa ([Mihaylov et al., 2018](#)). Following [Kojima et al. \(2022\)](#), we extract the answer from  
227 LLM-generated texts by looking at the first appearance of A/B/C/D. To test for calibration we used  
228 standard calibration metrics, including the ECE and TACE ([Naeini et al., 2015](#)), Brier score ([Murphy,](#)  
229 [1973](#)) and AUROC ([Murphy, 2012](#)). We note that for a fair comparison when the accuracy drops  
230 significantly, we must consult the Brier score which is a proper scoring rule. This is because, the ECE,  
231 TACE and AUROC are not proper scoring rules and can in general trade-off accuracy for calibration.  
232 For a baseline, we assumed 100% confidence for each deterministic prediction. We also tested the  
233 "hint" based approach of [Xiong et al. \(2023\)](#), which we describe in detail in Appendix A.

234 We present the results in Tables 11, 12 and 13. In the majority of cases rephrasing outperforms the  
235 naive baseline by 10 – 40% in AUROC, 10 – 30% in ECE, and 0 – 0.4 in Brier. Our approach  
236 also typically outperforms the "hint" base approach of [Xiong et al. \(2023\)](#) by 10 – 20% in AUROC,  
237 5 – 10% in ECE, and 0.1 in Brier. In particular, the "hint" based approach which more inflexible  
238 than our approach and typically hurts accuracy significantly 10 – 20% compared to 5 – 10% for our

Table 8: Evaluation results on ARC-Challenge with various rephrasing methods applied to three LLMs. In the majority of cases, the rephrasing approach outperforms the naive baseline by 10 – 40% in AUROC, 10 – 30% in ECE and 0 – 0.4 in Brier.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-1	0.742	0.258	<b>0.065</b>	0.517	0.5	-
	hint	0.593	0.201	0.108	0.614	0.695	-
	reword	0.619	0.12	0.103	0.512	<b>0.846</b>	1.0
	rephrase	0.555	0.125	0.103	0.571	0.817	1.5
	paraphrase	0.525	<b>0.102</b>	0.115	0.592	0.827	1.5
	expansion	0.602	0.133	0.099	<b>0.509</b>	0.847	1.0
Llama-2-7B	top-1	0.483	0.517	-	1.034	0.5	-
	hint	0.258	<b>0.071</b>	<b>0.144</b>	0.839	0.562	-
	reword	0.352	0.193	0.176	0.853	0.626	1.5
	rephrase	0.381	0.263	0.173	0.871	0.656	1.5
	paraphrase	0.39	0.287	0.162	0.883	0.67	1.0
	expansion	0.373	0.112	0.153	<b>0.778</b>	<b>0.687</b>	1.5
Llama-2-13B	top-1	0.508	0.492	-	0.983	0.5	-
	hint	0.331	0.147	0.134	0.813	0.57	-
	reword	0.445	<b>0.084</b>	<b>0.119</b>	<b>0.714</b>	0.721	1.5
	rephrase	0.441	0.128	0.134	0.727	0.713	1.5
	paraphrase	0.453	0.092	0.129	0.717	0.697	1.5
	expansion	0.441	0.154	0.142	<b>0.715</b>	<b>0.784</b>	1.2

239 approach. For our method, these accuracy drops are more prevalent in the smaller 7B models, while  
 240 the larger 13B model often shows a much smaller drop.

241 Crucially, the different rephrasing methods exhibit different calibration gains. On average, in terms  
 242 of all calibration metrics the best methods are the "expansion" and "reword" methods, which make  
 243 the queries more verbose, and substitute words with synonyms respectively. In terms of AUROC  
 244 "expansion" outperforms the alternatives by 1 – 5%. In terms of the Brier score it outperforms by  
 245  $\approx 0.05$ . To instantiate our rephrasings we used a prompt with a one-shot example and a temperature  
 246 parameter resulting in greater or smaller varieties of rephrasings. We include this temperature  
 247 parameter in the Tables. Generally, we choose this temperature that balances accuracy and calibration.  
 248 In Figure 4 we plot the behaviour as the number of MC draws increases.

249 In Appendix D, we also compare with Chain-of-Thought Wei et al. (2022) for uncertainty estimation.  
 250 We find that we get competitive results with CoT. At the same time our method is significantly easier  
 251 and more natural to implement for humans interacting via text with an LLM.

252 In Tables 11, 12 and 13 and Figure 3, we present the results for the top-k experiments with and  
 253 without rephrasing, with  $k = 40$ . We also present the relaxed temperature sampling variant Wei et al.  
 254 (2022). We see that the stochasticity of top-40 compared to top-1 decoding from Tables 8, 9 and 10  
 255 results in an improvement in calibration but a drop in accuracy. The Brier score often improves at the  
 256 cost of accuracy. Further stochasticity in answers caused by rephrasings has a similar effect. These  
 257 observations are consistent with the fact that top-k and nucleus sampling (Holtzman et al., 2019)  
 258 make text more human-like but not necessarily more "accurate".

Table 9: Evaluation results on ARC-Easy with various rephrasing methods applied to three LLMs. In the majority of cases, the rephrasing approach outperforms the naive baseline by 10 – 40% in AUROC, 10 – 30% in ECE, and 0 – 0.4 in Brier.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-1	0.866	0.134	<b>0.034</b>	<b>0.269</b>	0.5	-
	hint	0.773,	0.17,	0.076,	0.386,	0.795,	-
	reword	0.753	0.045	0.062	0.297	0.931	1.0
	rephrase	0.678	0.035	0.076	0.357	<b>0.953</b>	1.5
	paraphrase	0.663	0.036	0.08	0.381	0.943	1.5
	expansion	0.742	<b>0.034</b>	0.067	0.31	0.936	1.0
Llama-2-7B	top-1	0.672	0.328	0.082	0.656	0.5	-
	hint	0.231,	<b>0.041,</b>	0.149,	0.827,	0.663,	-
	reword	0.43	0.084	<b>0.119</b>	0.672	0.818	1.5
	rephrase	0.535	0.131	<b>0.117</b>	<b>0.603</b>	0.830	1.5
	paraphrase	0.526	0.184	0.125	0.626	<b>0.831</b>	1.0
	expansion	0.405	0.045	<b>0.119</b>	0.692	0.818	1.5
Llama-2-13B	top-1	0.617	0.383	<b>0.096</b>	0.767	0.5	-
	hint	0.346,	<b>0.089,</b>	0.128,	0.77,	0.673,	-
	reword	0.546	0.07	0.11	0.58	0.814	1.5
	rephrase	0.526	0.07	0.112	0.579	0.842	1.5
	paraphrase	0.518	0.104	0.119	0.604	0.815	1.5
	expansion	0.524	0.078	0.12	<b>0.552</b>	<b>0.893</b>	1.2

Table 10: Evaluation results on OpenBookQA with various rephrasing methods applied to three LLMs. In the majority of cases, the rephrasing approach outperforms the naive baseline by 10 – 40% in AUROC, 10 – 30% in ECE, and 0 – 0.4 in Brier.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-1	0.655	0.345	<b>0.086</b>	0.69	0.5	-
	hint	0.56,	0.265,	0.119,	0.71,	0.606,	-
	reword	0.552	0.105	0.102	<b>0.592</b>	0.796	1.0
	rephrase	0.482	0.107	0.122	0.641	0.809	1.5
	paraphrase	0.49	<b>0.076</b>	0.116	0.622	0.826	1.5
	expansion	0.518	0.087	0.117	0.596	<b>0.837</b>	1.0
Llama-2-7B	top-1	0.478	0.522	<b>0.131</b>	1.045	0.5	-
	hint	0.275,	<b>0.08,</b>	0.142,	0.832,	0.556,	-
	reword	0.388	0.137	0.143	0.786	0.689	1.5
	rephrase	0.39	0.196	0.156	0.806	<b>0.721</b>	1.5
	paraphrase	0.398	0.227	0.159	0.834	0.712	1.0
	expansion	0.362	0.083	0.138	<b>0.775</b>	0.678	1.5
Llama-2-13B	top-1	0.418	0.582	-	1.165	0.5	-
	hint	0.295,	<b>0.069,</b>	<b>0.138,</b>	0.809,	0.613,	-
	reword	0.428	0.117	0.142	0.75	0.676	1.5
	rephrase	0.428	0.095	0.14	<b>0.729</b>	0.73	1.5
	paraphrase	0.41	0.116	0.141	0.759	0.682	1.5
	expansion	0.41	0.143	0.147	0.772	<b>0.702</b>	1.2

Table 11: Evaluation results on ARC-Challenge with various rephrasing methods applied to three LLMs using top-k decoding. In the majority of cases rephrasing + top-k outperforms simple top-k in terms of calibration.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-k	0.746,	0.272,	0.091,	0.511,	0.6,	-
	temp-sampling	0.742	0.272	0.089	0.513	0.605	-
	reword	0.547,	<b>0.05,</b>	0.093,	0.543,	<b>0.864,</b>	1.5
	rephrase	0.64,	0.106,	<b>0.086,</b>	<b>0.485,</b>	0.82,	1.0
	paraphrase	0.631,	0.11,	0.098,	0.495,	0.83,	1.0
	expansion	0.517,	0.061,	0.114,	0.573,	0.859,	1.5
Llama-2-7B	top-k	0.436,	0.201,	<b>0.139,</b>	<b>0.761,</b>	0.602,	-
	temp-sampling	0.441	0.211	0.132	0.757	0.621	-
	reword	0.335,	0.187,	0.166,	0.858,	0.62,	1.5
	rephrase	0.356,	0.314,	0.17,	0.944,	0.627,	1.0
	paraphrase	0.309,	0.185,	0.162,	0.851,	<b>0.69,</b>	1.5
	expansion	0.322,	<b>0.144,</b>	0.155,	0.828,	0.622,	1.5
Llama-2-13B	top-k	0.462,	0.125,	<b>0.115,</b>	<b>0.679,</b>	<b>0.753,</b>	-
	temp-sampling	0.47,	0.122	0.115	0.662	0.766	-
	reword	0.352,	0.087,	0.136,	0.771,	0.687,	1.5
	rephrase	0.398,	<b>0.068,</b>	0.136,	0.725,	0.743,	1.0
	paraphrase	0.364,	0.109,	0.137,	0.738,	0.719,	1.2
	expansion	0.373,	0.124,	0.143,	0.76,	0.669,	1.5

Table 12: Evaluation results on ARC-Easy with various rephrasing methods applied to three LLMs using top-k decoding. In the majority of cases rephrasing + top-k outperforms simple top-k in terms of calibration.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-k	0.868,	0.133,	<b>0.042,</b>	<b>0.255,</b>	0.695,	-
	temp-sampling	0.859	0.131	0.046	0.266	0.677	-
	reword	0.694,	0.054,	0.076,	0.344,	0.941,	1.5
	rephrase	0.789,	0.047,	0.049,	0.274,	0.911,	1.0
	paraphrase	0.753,	<b>0.036,</b>	0.056,	0.3,	0.922,	1.0
	expansion	0.63,	0.042,	0.086,	0.403,	<b>0.942,</b>	1.5
Llama-2-7B	top-k	0.612,	0.25,	0.115,	0.612,	0.73,	-
	temp-sampling	0.619	0.261	0.114	0.617	0.717	-
	reword	0.401,	<b>0.074,</b>	0.121,	0.681,	0.825,	1.5
	rephrase	0.564,	0.145,	<b>0.108,</b>	<b>0.584,</b>	0.819,	1.0
	paraphrase	0.425,	0.08,	0.117,	0.665,	<b>0.835,</b>	1.5
	expansion	0.335,	0.054,	0.138,	0.742,	0.791,	1.5
Llama-2-13B	top-k	0.557,	0.06,	<b>0.098,</b>	<b>0.528,</b>	<b>0.865,</b>	-
	temp-sampling	0.544	0.087	0.107	0.532	0.866	-
	reword	0.412,	0.106,	0.129,	0.72,	0.741,	1.5
	rephrase	0.458,	<b>0.05,</b>	0.12,	0.643,	0.817,	1.0
	paraphrase	0.427,	0.066,	0.126,	0.652,	0.845,	1.2
	expansion	0.366,	0.087,	0.13,	0.74,	0.75,	1.5

Table 13: Evaluation results on OpenBookQA with various rephrasing methods applied to three LLMs using top-k decoding. In the majority of cases rephrasing + top-k outperforms simple top-k in terms of calibration.

Model	Rephrasing	Acc $\uparrow$	ECE $\downarrow$	TACE $\downarrow$	Brier $\downarrow$	AUROC $\uparrow$	temp
Mistral-7B	top-k	0.638,	0.289,	0.101,	0.636,	0.636,	-
	temp-sampling	0.668	0.289	0.098	0.607	0.624	-
	reword	0.528,	0.103,	0.105,	0.606,	0.794,	1.5
	rephrase	0.582,	0.109,	<b>0.093</b> ,	<b>0.542</b> ,	<b>0.821</b> ,	1.0
	paraphrase	0.552,	0.078,	0.101,	0.57,	0.817,	1.0
	expansion	0.445,	<b>0.061</b> ,	0.128,	0.653,	0.818,	1.5
Llama-2-7B	top-k	0.412,	0.208,	<b>0.129</b> ,	<b>0.776</b> ,	0.617,	-
	temp-sampling	0.442	0.235	0.13	0.772	0.599	-
	reword	0.34,	0.14,	0.153,	0.807,	0.696,	1.5
	rephrase	0.408,	0.239,	0.154,	0.815,	0.704,	1.0
	paraphrase	0.355,	0.127,	0.145,	0.783,	<b>0.721</b> ,	1.5
	expansion	0.308,	<b>0.098</b> ,	0.151,	0.807,	0.711,	1.5
Llama-2-13B	top-k	0.43,	0.114,	<b>0.13</b> ,	<b>0.708</b> ,	<b>0.72</b> ,	-
	temp-sampling	0.43,	0.099	0.121	0.702	0.733	-
	reword	0.345,	0.111,	0.144,	0.794,	0.618,	1.5
	rephrase	0.345,	<b>0.062</b> ,	0.141,	0.767,	0.706,	1.0
	paraphrase	0.37,	0.092,	0.141,	0.763,	0.67,	1.2
	expansion	0.36,	0.138,	0.138,	0.799,	0.574,	1.5




Figure 4: The behavior of the Accuracy, ECE, TACE, Brier, and AUROC for all datasets, architectures, and expansion methods, as we increase the number of samples. We plot the average value as well as confidence intervals  $\pm 2\sigma$ . We see that the ECE and the AUROC improve with more samples while the accuracy drops slightly. This might be because the meaning of some queries is completely destroyed by our rephrasings. The Brier score captures this tradeoff by having a minimum at approximately 5 samples. The TACE remains relatively stable with respect to the number of samples.