

# Large Language Models Are Overconfident in Their *Own* Responses

Mario Sanz-Guerrero<sup>∇</sup> Manuel Mager<sup>∇</sup> Katharina von der Wense<sup>∇♣</sup>

<sup>∇</sup>Johannes Gutenberg University Mainz, Germany

<sup>♣</sup>University of Colorado Boulder, USA

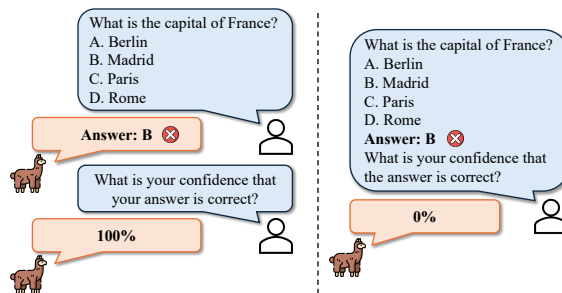
msanz@uni-mainz.de

## Abstract

Prior work has shown that instruction-tuned large language models (LLMs) are less well calibrated than their base pre-trained counterparts. However, little is known about the frequently used chat template’s effect on the calibration of conversational LLMs. In this work, we investigate the mechanisms driving this miscalibration by decoupling the effects of the post-training algorithm and the chat format. We find that, while instruction tuning fundamentally harms calibration, the chat template aggravates the issue through an “ownership bias” – models are significantly more confident in their *own* answers than in identical answers provided by a user. Extensive experiments across six recent open-weight LLMs, three benchmarks, and three confidence elicitation methods show that models assign up to 26% higher confidence to their own responses. Leveraging this insight, we propose a simple inference-time strategy: framing the model’s answer as user input during confidence elicitation. This approach significantly reduces overconfidence and improves calibration by up to 26% without the need for retraining, narrowing the gap between base and instruction-tuned models.

## 1 Introduction

Reliable confidence estimation is crucial for safe and trustworthy AI. As large language models (LLMs) are increasingly employed in high-stakes applications, their ability to accurately assess the certainty of their predictions becomes critical. A well-calibrated model should be highly confident only when it is correct (Guo et al., 2017). However, while base (pre-trained) LLMs are generally well-calibrated, instruction-tuned (post-trained) LLMs exhibit significant miscalibration (Tian et al., 2023; OpenAI et al., 2024). This gap raises a critical question: *what causes instruction-tuned LLMs to be miscalibrated, and are there ways to mitigate the problem that are easy to apply for end users?*



(a) Answer provided by the assistant (model’s own response). (b) Answer provided by the user.

Figure 1: LLMs are overconfident in their *own* answers, regardless of whether they are correct or not, leading to miscalibration. The figure represents real outputs from Llama 3.1 (8B).

We address this question in four steps: 1) We investigate whether the reduced calibration stems from the training algorithm or the prompting style by isolating the effects of instruction tuning and the chat template (frequently introduced during this step). We find that the post-training process is the main source of miscalibration, though the chat template further aggravates the issue. 2) We explore if *explicitly* asking the model for its certainty (as compared to looking at answer probabilities directly) alters the impact of instruction tuning and the chat template. We find similar trends – instruction-tuned LLMs remain significantly worse calibrated than base models, and the post-training algorithm is the root of the problem. 3) We hypothesize that LLMs exhibit an “ownership bias” – i.e., that they are overconfident in their *own* answers – which drives the effects observed in our experiments. Based on this, we propose a straightforward and, importantly, easy to apply strategy to reduce miscalibration and narrow the gap between base and instruction-tuned models: framing the model’s answer as a user input during confidence elicitation. 4) Finally, with an additional analysis, we confirm that the effectiveness of our proposed method does, in fact, stem from reducing the models’ inherent overconfidence in their own answers.

## 2 Related Work

**The Impact of Post-training on LLM Calibration** While pre-trained LLMs generally exhibit well-calibrated probabilities for next-token prediction, the post-training process – usually done through Supervised Fine-tuning (SFT) and Reinforcement Learning from Human Feedback (RLHF) – has been shown to significantly degrade this property (Kadavath et al., 2022; Tian et al., 2023; OpenAI et al., 2024; Nakkiran et al., 2025). Zhu et al. (2023) systematically analyze this phenomenon, finding that instruction tuning and the use of synthetic data act as primary drivers of miscalibration.

**Mechanisms of Miscalibration** Recent studies have sought to understand the underlying mechanisms that lead to miscalibration in instruction-tuned LLMs. Leng et al. (2025) identify that reward models used in Proximal Policy Optimization (PPO; Schulman et al., 2017) often exhibit a bias toward high-confidence responses, regardless of their actual correctness; this reward bias trains the policy model to be overconfident. Similarly, Xiao et al. (2025) attribute miscalibration to “preference collapse,” where the model’s optimization for human preference leads it to ignore alternative, potentially correct answers, thereby artificially concentrating probability mass and increasing confidence.

**Verbalized Confidence and Elicitation** Given the poor calibration of logits in post-trained LLMs, several works have shifted toward prompting models to express their uncertainty in natural language. Lin et al. (2022a) demonstrate that LLMs can be fine-tuned to generate “verbalized probability” that correlates well with empirical accuracy. Tian et al. (2023) further establish that, for instruction-tuned LLMs, these verbalized confidence scores are better calibrated than the model’s conditional probabilities. However, Xiong et al. (2024) shows that verbalized confidence is prone to overconfidence, potentially due to the model imitating human patterns of expressing uncertainty.

**Mitigation Strategies** To address these issues, several mitigation strategies have been proposed recently. Luo et al. (2025) introduce Disagreement-Aware Confidence Alignment, an unsupervised method that leverages the calibrated nature of pre-trained LLMs to guide the calibration of post-trained ones. Other approaches involve complex interventions during training, such as calibrated

reward modeling (Leng et al., 2025) or calibration-aware fine-tuning (Xiao et al., 2025), or the use of auxiliary models to estimate uncertainty (Mielke et al., 2022; Ulmer et al., 2024).

Unlike these resource-intensive methods, our work identifies a fundamental mechanism behind miscalibration in instruction-tuned LLMs – their overconfidence when assessing their *own* answers – and proposes a simple inference-time strategy to alleviate this issue that does not require retraining.

## 3 Finding the Root of Miscalibration

Prior work has shown that *pre-trained* LLMs are well-calibrated but post-training processes harm this calibration (Tian et al., 2023; OpenAI et al., 2024; *inter alia*). Instruction-tuned models are trained to follow a *chat format*, where the model plays the role of an assistant responding to user messages (Ouyang et al., 2022). These models usually start their responses in a conversational manner, e.g., “Certainly” or “Sure,” and their first-token probabilities are poorly aligned with their actual accuracy (Wang et al., 2024). Given these findings, one might hypothesize that the chat format itself is responsible for the miscalibration of instruction-tuned LLMs. However, all prior work has only evaluated instruction-tuned LLMs using the chat format, and little is known about the specific contribution of the chat template to calibration. We begin by asking: *are instruction-tuned LLMs truly miscalibrated, or are they simply miscalibrated when following the chat format?*

### 3.1 Experimental Setup

**Methods** To isolate the effect of instruction tuning and the chat format, we compare three variants of each model: (1) the base, pre-trained model without instruction tuning (prompted as in Figure 2a); (2) the instruction-tuned model invoked *without* the chat template (i.e., as if it was the base model; Figure 2a); and (3) the instruction-tuned model invoked *with* the chat template (Figure 2b).

The following is a multiple-choice question (with answer). Question: {question} A. {option A} B. {option B} C. {option C} D. {option D} Answer: $\rightarrow P("X") \forall X \in \{A, B, C, D\}$	<b>User:</b> Answer the following multiple-choice question. Question: {question} A. {option A} B. {option B} C. {option C} D. {option D} <b>Assistant:</b> Answer: $\rightarrow P("X") \forall X \in \{A, B, C, D\}$
---	--

(a) Without chat template.

(b) With chat template.

Figure 2: Prompts used to evaluate models in Section 3.

**Models** Throughout this work, we run our experiments on 6<sup>1</sup> open-weight LLMs for which both the base (pre-trained) and instruction-tuned (post-trained) versions are available: Llama 3.1 (8B & 70B; Grattafiori et al., 2024), Qwen3 (4B & 30B; Yang et al., 2025), and Gemma 3 (4B & 27B; Gemma Team et al., 2025). These models cover three different families, and we take a smaller and a larger version of each to assess whether the observed trends hold across different scales.

**Datasets** In all our experiments, we use the MMLU dataset (Hendrycks et al., 2021), one of the most widely used benchmarks for evaluating the factual knowledge of LLMs. MMLU contains multiple-choice questions from 57 disciplines with different levels of difficulty. Each question has four answer options, which allows us to easily evaluate model calibration, since we can directly obtain the probability of each option from the model logits and determine whether the answer is correct or not.

**Evaluation** To be consistent with prior work (Tian et al., 2023; OpenAI et al., 2024; Sanz-Guerrero et al., 2025), in this first experiment we extract the model’s confidence from the probability of the answer tokens directly (i.e., logit-based confidence estimation). To evaluate performance, we select the answer with the highest probability as the model’s prediction and report **accuracy**. To evaluate calibration, we use the probabilities assigned to all answer options as the model’s confidence for each of the four choices (where only one is correct). Using these confidences, we compute **Expected Calibration Error (ECE; Pakdaman Naeini et al., 2015)**, which measures the expected difference between model confidence and actual accuracy. We partition predictions into  $M = 10$  bins of equal width (i.e., 10% confidence intervals) and calculate the weighted average of the absolute difference between the accuracy and confidence in each bin:

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

where  $N$  is the total number of samples,  $|B_m|$  is the number of samples in bin  $m$ , and  $\text{acc}(B_m)$  and  $\text{conf}(B_m)$  are the average accuracy and confidence in that bin, respectively. We also report the **Brier score (Brier, 1950)**, which measures the accuracy

<sup>1</sup>Technically, we run experiments on 12 models (total of 6 base and 6 instruct), but we consider the base and instruct versions of each model as a single model for simplicity.

Model	IT	Chat	Acc. (↑)	ECE (↓)	Brier (↓)
Llama 3.1 (8B)	✗	✗	62.81	<b>0.0664</b>	<b>0.1706</b>
	✓	✗	68.77	0.1176	0.1890
	✓	✓	<b>69.12</b>	0.1666	0.2005
Llama 3.1 (70B)	✗	✗	74.04	<b>0.0572</b>	<b>0.1232</b>
	✓	✗	78.60	0.0796	0.1413
	✓	✓	<b>79.30</b>	0.1087	0.1425
Qwen3 (4B)	✗	✗	67.72	<b>0.0425</b>	<b>0.1709</b>
	✓	✗	69.82	0.2395	0.2393
	✓	✓	<b>72.98</b>	0.2415	0.2455
Qwen3 (30B)	✗	✗	81.40	<b>0.0397</b>	<b>0.1137</b>
	✓	✗	82.11	0.1134	0.1279
	✓	✓	<b>83.16</b>	0.1373	0.1429
Gemma 3 (4B)	✗	✗	49.47	<b>0.0619</b>	<b>0.1971</b>
	✓	✗	56.95	0.3835	0.3814
	✓	✓	<b>58.14</b>	0.4214	0.4161
Gemma 3 (27B)	✗	✗	72.98	<b>0.1028</b>	<b>0.1418</b>
	✓	✗	74.39	0.2226	0.2260
	✓	✓	<b>74.49</b>	0.2453	0.2448

Table 1: Accuracy and calibration of models on MMLU. “IT” indicates instruction-tuned models and “Chat” indicates applying the chat template. Instruct models perform better, but are worse calibrated than base models.

of probabilistic predictions as the mean squared error between the predicted probability  $f_i$  and the actual outcome  $o_i$  (1 if correct, 0 otherwise):

$$\text{BS} = \frac{1}{N} \sum_{i=1}^N (f_i - o_i)^2 \quad (2)$$

### 3.2 Results

Our results (Table 1) confirm that instruction tuning improves accuracy but harms calibration, a trade-off observed in prior work (Tian et al., 2023; OpenAI et al., 2024). On average, instruction tuning boosts accuracy by 3.7% but degrades calibration, increasing ECE by 13.1% and Brier score by 6.5%. Applying the chat template further aggravates this issue: while it yields a modest accuracy gain (+1.1%), it adds another 2.74% to ECE and 1.5% to Brier score. Overall, the combination of instruction tuning and chat templates leads to a total ECE increase of 15.8% compared to base models.

Our findings indicate that, while instruction tuning is the main driver of miscalibration, the chat template also contributes to this effect. Since the general public typically interacts with LLMs in this exact configuration – instruction-tuned and using chat templates (e.g., ChatGPT) – addressing miscalibration in this setting is critical. This motivates the exploration of techniques to obtain model confidence more reliably in these widely used systems.

## 4 Does Explicitly Asking for Confidence Alter the Miscalibration Trends?

Prior work has shown that verbalized confidence estimation methods (i.e., explicitly asking the model for its confidence) lead to better calibration than logit-based methods for instruction-tuned LLMs (Tian et al., 2023). To further investigate the effect of instruction tuning and the chat format on calibration, we explore whether *explicitly* asking the model for its confidence changes the observed miscalibration trends.

### 4.1 Experimental Setup

**Methods** We compare the same three model variants as in Section 3 (base model, instruction-tuned without chat template, and instruction-tuned with chat template) using three confidence estimation methods well established in the literature: (1) **P(True)**, a logit-based method that computes the probability of the “true” token after asking the model whether the provided answer is correct (Kadavath et al., 2022); (2) **Verbalized Percentage**, where we ask the model to provide a percentage score of confidence (0–100%) in the answer being correct (Tian et al., 2023; Lin et al., 2022a); and (3) **Verbalized Linguistic**, where the model selects a qualitative confidence score from a defined set of linguistic expressions forming a Likert scale (Likert, 1932) ranging from “very low” to “very high” (Tian et al., 2023; Lin et al., 2022a).

**Evaluation** As in our previous experiment, we assess calibration not only on the top answer (i.e., the model’s prediction) but on *all* possible answers – ideally, the model should be confident when an answer is correct and uncertain otherwise. To achieve this, we “force” the model to consider each of the four options in turn, asking for its confidence in that specific answer. This allows us to assess the model’s calibration in low-confidence scenarios as well (i.e., when the model is asked about incorrect answers). As in the previous section, we measure calibration using **ECE** and **Brier score**. For “P(True)” and “Verbalized Percentage,” we use the direct probability values obtained from the model (from the logits or the generated tokens, respectively). For “Verbalized Linguistic,” we map the linguistic categories to numerical scores between 0 and 1 using equal intervals.<sup>2</sup> Since we now ask for

<sup>2</sup>We consider seven options (“very low,” “low,” “somewhat low,” “medium,” “somewhat high,” “high,” “very high”) and map them to 0.0, 0.17, 0.33, 0.5, 0.67, 0.83, and 1.0.

confidence scores for every option independently, we do not compute standard accuracy – in this setup, the model is not making a single prediction, but rather providing a confidence estimate for each possible answer.

### 4.2 Results

Our results, reported in Table 2, are consistent with those in Section 3: by explicitly asking for confidence in all three methods, instruction-tuned models remain significantly worse calibrated than base models, regardless of the chat template. This further confirms that the post-training process is the root cause of miscalibration.

## 5 Are LLMs Overconfident in Their Own Answers?

Our results so far indicate that instruction-tuning does harm the calibration of LLMs, regardless of the chat format. However, given the conversational nature of these models, it seems plausible to think that they should be confident in the answers they give. As recently shown by Kalai et al. (2025), LLMs are inherently trained to predict answers – even incorrect ones – over recognizing “I don’t know.” Here, we hypothesize that LLMs are overconfident in their *own* answers, regardless of whether they are correct or not (see Figure 1).

To test this hypothesis, we prompt LLMs with a question and a possible answer, showing the answer as part of either (1) the assistant message (i.e., the model’s own response – the normal usage of LLMs); or (2) the user message. We then ask the model to provide its confidence in the answer being correct. If LLMs are indeed overconfident in their own answers, we should observe higher confidence when the answer is presented as the assistant’s response compared to the user message.

### 5.1 Experimental Setup

**Methods** We run our experiments using two different prompts, where the only difference is whether the answer is provided as part of the assistant message or the user message (see Figure 3). We compare the effect of *who* provides the answer across three confidence estimation methods well established in the literature, as described in Section 4.1: **P(True)**, **Verbalized Percentage**, and **Verbalized Linguistic**.

**Evaluation** As in our previous experiments, we measure calibration using **ECE** and **Brier score**,

Model	IT	Chat	ECE ( $\downarrow$ )			Brier ( $\downarrow$ )		
			P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.
Llama 3.1 (8B)	$\times$	$\times$	<b>0.1910</b>	<b>0.3483</b>	<b>0.3161</b>	<b>0.1930</b>	<b>0.3174</b>	<b>0.2872</b>
	$\checkmark$	$\checkmark$	0.2796	0.6490	0.3767	0.2315	0.5992	0.3252
Llama 3.1 (70B)	$\times$	$\times$	<b>0.0267</b>	<b>0.1733</b>	<b>0.2745</b>	<b>0.0992</b>	<b>0.1405</b>	<b>0.1983</b>
	$\checkmark$	$\checkmark$	0.0744	0.2135	0.3151	0.1030	0.1739	0.2288
Qwen3 (4B)	$\times$	$\times$	<b>0.1635</b>	<b>0.4539</b>	<b>0.4123</b>	<b>0.1732</b>	<b>0.4182</b>	<b>0.3598</b>
	$\checkmark$	$\checkmark$	0.2268	0.5850	0.5021	0.2186	0.5544	0.4241
Qwen3 (30B)	$\times$	$\times$	<b>0.1326</b>	<b>0.4037</b>	<b>0.3589</b>	<b>0.1415</b>	<b>0.3377</b>	<b>0.2806</b>
	$\checkmark$	$\checkmark$	0.2340	0.5835	0.4565	0.1829	0.5001	0.3650
Gemma 3 (4B)	$\times$	$\times$	<b>0.0209</b>	<b>0.4558</b>	<b>0.3452</b>	<b>0.1882</b>	<b>0.3930</b>	<b>0.3062</b>
	$\checkmark$	$\checkmark$	0.5779	0.7132	0.5418	0.5645	0.6929	0.4731
Gemma 3 (27B)	$\times$	$\times$	<b>0.0428</b>	<b>0.5213</b>	<b>0.3356</b>	<b>0.1798</b>	<b>0.4465</b>	<b>0.2942</b>
	$\checkmark$	$\checkmark$	0.2395	0.5465	0.5256	0.2366	0.4785	0.4520

Table 2: ECE and Brier score on MMLU when explicitly asking for confidence using three confidence elicitation methods.

Answer from “Assistant”	Answer from “User”
<p><b>User:</b> Answer the following multiple-choice question. Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Assistant:</b> Answer: {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> Is your answer true or false? Respond only with “True” or “False”.</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{“True”})</math></p>	<p><b>User:</b> The following is a multiple-choice question (with answer). Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Assistant:</b> {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> Is the answer true or false? Respond only with “True” or “False”.</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{“True”})</math></p>

Figure 3: Prompts used to measure confidence in an answer provided by the assistant (left) and by the user (right) using the P(True) method (see Appendix A for the other two methods).

considering the model’s confidence in all answer options (not only the top-1 answer). Moreover, we measure the average **raw confidence** assigned by the model to each answer option to directly compare the levels of certainty between the two prompt formats.

**Statistical Significance** To assess whether the difference in confidence when the answer is provided by the user rather than the assistant is statistically significant, we use the Wilcoxon signed-rank test (Wilcoxon, 1945) for the Brier scores and raw confidence values. For the ECE, which is an aggregate metric computed over the entire dataset, we determine significance using a paired bootstrap

resampling test ( $K = 1000$  iterations; Efron and Tibshirani, 1994) to estimate the 95% confidence interval of the difference in ECE.

## 5.2 Results

Our results, reported in Table 3, demonstrate that LLMs are indeed more miscalibrated when they provide the answer themselves. Notably, all deltas are positive – indicating worse calibration for assistant-generated answers – and the vast majority are statistically significant. For both calibration metrics (ECE and Brier score), the confidence estimation method with the smallest differences is P(True); nevertheless, we still observe average differences of 9.8% in ECE and 8.8% in Brier score. The method with the largest difference is consistently linguistic expressions, reaching average differences above 25% in both metrics.

For illustrating the calibration across confidence levels, Figure 4 shows the reliability diagrams of all models using the three confidence estimation methods. In all cases, we observe an overall better calibration when the answer is provided by the user compared to when it is provided by the assistant.

Looking at the difference in raw confidence assigned by the model (Table 3, right columns), we observe that all models show a significant increase in confidence when the answer is provided by the assistant. The method with the smallest differences is P(True), with a notable average increase of

Model	$\Delta$ ECE			$\Delta$ Brier Score			$\Delta$ Confidence		
	P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.
Llama 3.1 (8B)	0.072*	0.076*	0.235*	0.032*	0.071*	0.216*	0.134*	0.081*	0.237*
Llama 3.1 (70B)	0.091*	0.020	0.025	0.067*	0.022*	0.029*	0.110*	0.015*	0.025*
Qwen3 (4B)	0.002	0.193*	0.411*	0.011*	0.227*	0.416*	0.164*	0.195*	0.416*
Qwen3 (30B)	0.045*	0.347*	0.436*	0.040*	0.354*	0.428*	0.070*	0.354*	0.461*
Gemma 3 (4B)	0.252*	0.121*	0.268*	0.253*	0.149*	0.246*	0.312*	0.121*	0.273*
Gemma 3 (27B)	0.126*	0.315*	0.191*	0.126*	0.345*	0.175*	0.158*	0.321*	0.198*
Average	0.098	0.179	0.261	0.088	0.195	0.252	0.158	0.181	0.268

Table 3: Difference in ECE, Brier score, and raw confidence on MMLU between answers provided by the assistant (model itself) and the user. Values represent  $\Delta = \text{Assistant} - \text{User}$ , so positive values indicate higher metrics for the assistant. \* indicates significant difference ( $p < 0.01$ ) according to the Wilcoxon signed-rank test (for Brier and confidence) or paired bootstrap resampling test (for ECE).

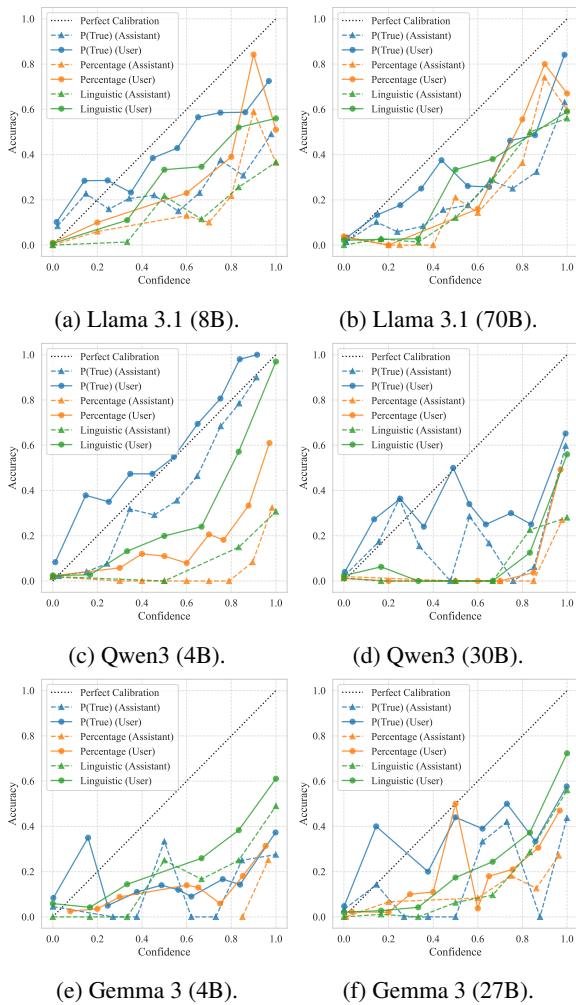


Figure 4: Reliability diagrams of all models using three confidence estimation methods. In all cases, we observe a better calibration when the answer is provided by the user compared to when it is provided by the assistant.

15.8%. The method with the largest differences is linguistic expressions, reaching an average increase of 26.8%. These differences are further illustrated in the confidence distributions (Figure 5). For all

three confidence estimation methods, we observe a clear shift towards higher confidence levels when the answer is provided by the assistant.

Our results demonstrate that LLMs are significantly worse calibrated when the answer is provided by themselves compared to when it is provided by the user and that the reason for this miscalibration is an overconfidence in their own answers. This finding contrasts with the phenomenon of *sycophancy*, where LLMs typically bias their responses to align with user inputs or beliefs (Sharma et al., 2024; Wei et al., 2024; Perez et al., 2023). If sycophancy were the dominant factor in confidence estimation, we would expect models to validate the user’s authority by assigning higher confidence to answers provided by the user. Instead, we observe an “ownership bias” where models are significantly more confident in their own outputs. We interpret this as a form of *self-consistency*: the model implicitly trusts its own generation process under the assumption that if it generated a specific answer, it must be confident in it; otherwise, it would have predicted a different response. This leads to inflated confidence scores for assistant-generated answers, which do not translate to improved calibration.

Therefore, we propose that, to obtain a more trustworthy confidence estimate from an LLM, the answer should be presented as part of the user message rather than as the model’s own response, as this leads to a more objective assessment of confidence – not guided by the bias that the model itself has provided the answer. This simple yet effective strategy significantly reduces overconfidence and recovers a calibration comparable to that of base models, sometimes even surpassing it (see Appendix B for full results), without requiring any changes to the model or additional training.




Figure 5: Distribution of confidence scores for answers provided by the assistant and by the user, aggregated across all models. P(True) and Verbalized Percentage are continuous values, while Verbalized Linguistic is discrete.

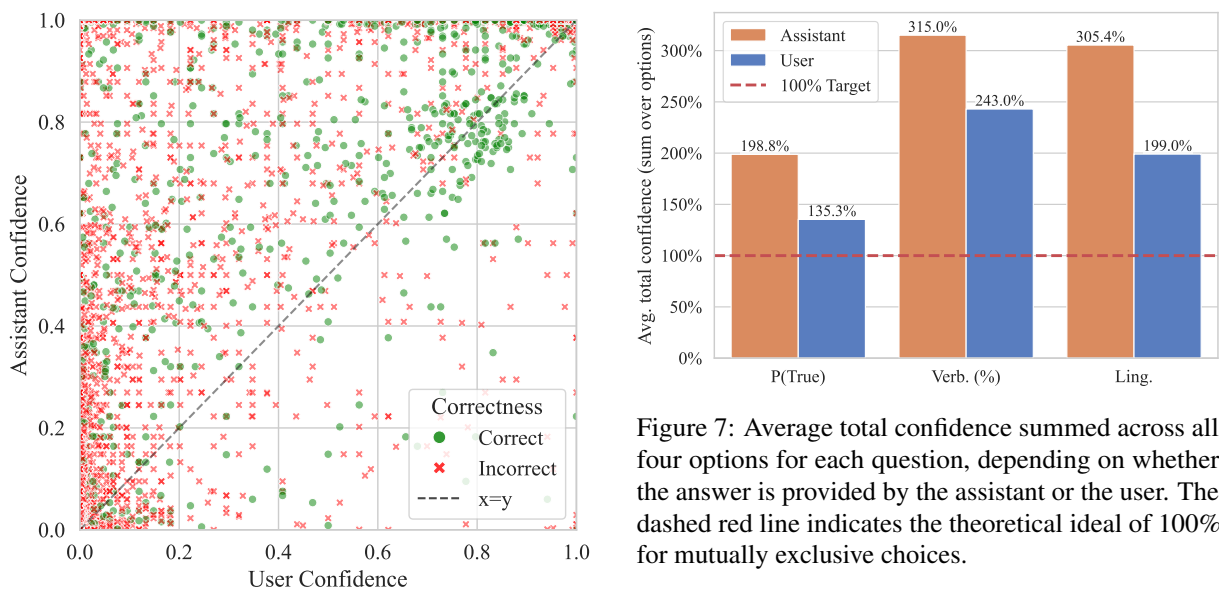


Figure 6: Confidence to answers provided by the user (x-axis) and by the assistant (y-axis), broken down by answer correctness (colors). Each point represents a single question-answer pair, aggregated across all models. The dashed diagonal line indicates equal confidence in both settings.

### 5.3 Analysis

#### 5.3.1 Assessment by Answer Correctness

The concept of calibration is closely tied to the correctness of answers – a well-calibrated model should be more confident when its answer is correct and less confident when it is incorrect. Our previous results indicate that LLMs are generally overconfident in their own answers, regardless of whether they are correct or not. To further analyze this phenomenon, we break down the confidence assigned by the model depending on whether the answer is correct or incorrect in Figure 6. We generally do not observe a clear distinction between

Figure 7: Average total confidence summed across all four options for each question, depending on whether the answer is provided by the assistant or the user. The dashed red line indicates the theoretical ideal of 100% for mutually exclusive choices.

the distributions of correct and incorrect answers, suggesting that the model’s confidence behavior is largely independent of accuracy. However, most data points are located above the diagonal, showing once again that the assistant is consistently more confident in its own answers than in the user’s. This is most pronounced in the leftmost cluster of incorrect answers, where the confidence assigned to user’s answers is correctly low (close to 0), yet the assistant’s self-confidence is erroneously high (up to 60%).

#### 5.3.2 Total Confidence

Since we experiment with a multiple-choice setting, the model’s total confidence across the four options for a single question should ideally add up to 100%. Figure 7 shows the total confidence (sum of confidences assigned to the four options) given by the model – averaged across all models – depending on whether the answer is provided by the assistant

	Model	$\Delta$ ECE			$\Delta$ Brier Score			$\Delta$ Confidence		
		P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.
GSM8K	Llama 3.1 (8B)	0.150*	0.241*	0.287*	0.127*	0.267*	0.348*	0.170*	0.302*	0.289*
	Llama 3.1 (70B)	0.248*	0.123*	0.091*	0.187*	0.124*	0.103*	0.366*	0.170*	0.038*
	Qwen3 (4B)	0.294*	0.093*	0.221*	0.282*	0.094*	0.233*	0.363*	0.100*	0.239*
	Qwen3 (30B)	0.011	0.003	0.009*	0.014	0.003*	0.007	0.085*	0.007*	0.011
	Gemma 3 (4B)	0.143*	0.012*	0.204*	0.155*	0.014*	0.228*	0.162*	0.012*	0.204*
	Gemma 3 (27B)	0.004	0.009	0.020*	0.005	0.006*	0.007	0.022	0.016*	0.031*
	Average	0.142	0.080	0.139	0.128	0.085	0.154	0.195	0.101	0.135
TruthfulQA	Llama 3.1 (8B)	0.053*	0.038*	0.078*	0.048*	0.035*	0.035	0.109*	0.193*	0.283*
	Llama 3.1 (70B)	0.046*	0.074*	0.051*	0.048*	0.086*	0.051*	0.068*	0.012	0.008*
	Qwen3 (4B)	0.087*	0.029*	0.164*	0.089*	0.024*	0.070*	0.079*	0.036*	0.138*
	Qwen3 (30B)	0.003	0.097*	0.044*	0.001	0.101*	0.054*	0.093*	0.108*	0.137*
	Gemma 3 (4B)	0.092*	0.115*	0.090*	0.092*	0.117*	0.067*	0.018*	0.098*	0.051*
	Gemma 3 (27B)	0.146*	0.027*	0.055*	0.145*	0.045*	0.052*	0.199*	0.034*	0.034*
	Average	0.071	0.063	0.080	0.070	0.068	0.055	0.094	0.080	0.109
MMLU (Open-ended)	Llama 3.1 (8B)	0.031*	0.254*	0.330*	0.044*	0.245*	0.388*	0.061*	0.296*	0.342*
	Llama 3.1 (70B)	0.046*	0.013	0.013	0.046*	0.002	0.020	0.042*	0.029	0.028*
	Qwen3 (4B)	0.064*	0.100*	0.179*	0.060*	0.136*	0.224*	0.077*	0.105*	0.163*
	Qwen3 (30B)	0.002	0.232*	0.242*	0.006	0.241*	0.225*	0.006	0.256*	0.283*
	Gemma 3 (4B)	0.176*	0.138*	0.232*	0.180*	0.149*	0.225*	0.203*	0.142*	0.246*
	Gemma 3 (27B)	0.056*	0.152*	0.092*	0.055*	0.150*	0.089*	0.063*	0.161*	0.113*
	Average	0.063	0.148	0.181	0.065	0.154	0.195	0.075	0.165	0.196

Table 4: Difference in ECE, Brier score, and raw confidence between answers provided by the assistant (model itself) and the user on GSM8K, TruthfulQA, and open-ended MMLU. Values represent  $\Delta = \text{Assistant} - \text{User}$ , so positive values indicate higher metrics for the assistant. \* indicates significant difference ( $p < 0.01$ ) according to the Wilcoxon signed-rank test (for Brier and confidence) or paired bootstrap resampling test (for ECE).

or the user. We observe that, with all confidence estimation methods, the average total confidence is always higher than 100%. This further indicates that LLMs are generally overconfident and are not consistent when assigning confidence to mutually exclusive options. Additionally, the total confidence is much higher when the answer is provided by the assistant (ranging from roughly 198% to 315%) compared to the user (ranging from roughly 135% to 243%), reinforcing our hypothesis that LLMs are overconfident in their *own* answers.

### 5.3.3 Generalization to Other Tasks

So far, our analysis has focused on multiple-choice questions from the MMLU dataset, where the model is evaluated over a closed set of options. While this format enables precise calibration measurement, one could argue that explicitly listing candidate answers affects confidence behavior. To test whether our findings generalize beyond this setup, we extend the evaluation to GSM8K (Cobbe et al., 2021), TruthfulQA (Lin et al., 2022b), and an open-ended version of MMLU without predefined options. These datasets cover diverse challenges: mathematical reasoning with open numeric answers (GSM8K), robustness to common

misconceptions (TruthfulQA), and free-form factual answering (open-ended MMLU). As shown in Table 4, the same pattern holds across all tasks: models are consistently better calibrated and less overconfident when judging answers framed as user input rather than as their own output. On GSM8K, self-generated answers yield up to 19.5% higher confidence and a 14.2% increase in ECE. On TruthfulQA, the confidence gap reaches up to 10.9% with similarly higher calibration error. Open-ended MMLU shows the same behavior, with up to 19.6% higher confidence and 18.1% higher ECE for assistant-framed answers. Together, these results show that the ownership bias is not a multiple-choice artifact but a broader phenomenon across domains and answer formats.

### 5.3.4 Generalization to Proprietary Models

To test whether this behavior also appears in proprietary models, we further run an experiment with GPT-5.2 using MMLU. Table 5 shows the same directional trend observed in open-weight models: assistant-framed answers produce higher ECE, Brier score, and raw confidence than user-framed answers across all three confidence elicitation methods. ECE and confidence differences are statisti-

Method	$\Delta$ ECE	$\Delta$ Brier	$\Delta$ Confidence
P(True)	0.077*	0.042	0.076*
Verb. (%)	0.087*	0.047	0.112*
Ling.	0.113*	0.095*	0.222*

Table 5: Difference in ECE, Brier score, and raw confidence between answers provided by the assistant (model itself) and the user for GPT-5.2 on MMLU. Values represent  $\Delta = \text{Assistant} - \text{User}$ , so positive values indicate higher metrics for the assistant. \* indicates significant difference ( $p < 0.01$ ) according to the Wilcoxon signed-rank test (for Brier and confidence) or paired bootstrap resampling test (for ECE).

cally significant for all methods, while Brier score also degrades and reaches significance for the linguistic method. These results indicate that ownership bias extends beyond open-weight LLMs and, thus, reinforce the generality of our findings as well as the practical relevance of our inference-time mitigation.

## 6 Conclusion

In this work, we investigate the specific impact of the chat template on the miscalibration of instruction-tuned LLMs. While we find that the chat format itself is not the main driver of miscalibration, it plays a critical role in how confidence is perceived: models exhibit an inherent “ownership bias,” showing significantly higher confidence in their own responses than in identical answers provided by a user, regardless of correctness. Building on this insight, we propose a simple inference-time strategy that turns the chat format into an advantage: framing the model’s answer as a user input during confidence elicitation. This method significantly reduces overconfidence and improves calibration across diverse benchmarks, narrowing the gap between base and instruction-tuned models. Our findings suggest that reliable confidence estimation requires decoupling the generation of an answer from its evaluation, forcing the model into a more objective “observer” role.

## Limitations

Our findings highlight the existence of overconfidence in LLMs’ own answers and offer a simple yet effective strategy to reduce it. However, there are several limitations to our study.

First, although we extend our analysis to GPT-5.2 (Section 5.3.4), the majority of our experiments focus on open-weight LLMs due to the high cost

of experiments and our commitment to open science and reproducibility. While our results cover a range of sizes, families, and include a proprietary model, we cannot fully guarantee that the same overconfidence behavior persists to the same degree in all proprietary, closed-source models where the post-training recipes (RLHF strategies and data mixtures) may differ.

Moreover, our proposed method – framing model outputs as user inputs – is an inference-time mitigation strategy. While it reduces miscalibration and narrows the gap between base and instruction-tuned models by leveraging the chat format, it does not alter the model weights or address the root causes of overconfidence introduced during the alignment process.

Finally, while we extended our analysis beyond multiple-choice questions to other tasks, our scope remains limited to objective question answering. It remains an open question whether similar overconfidence occurs in more subjective tasks (e.g., open-ended generation), where “correctness” is less clearly defined and confidence estimation is more ambiguous.

## Acknowledgments

We thank the anonymous reviewers for their helpful feedback. This work was supported by the Carl Zeiss Foundation through the MAINCE project (grant number P2022-08-009).

## References

- Glenn W. Brier. 1950. [Verification of forecasts expressed in terms of probability](#). *Monthly Weather Review*, 78(1):1–3.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. [Training verifiers to solve math word problems](#). *Preprint*, arXiv:2110.14168.
- Bradley Efron and Robert J Tibshirani. 1994. *An introduction to the bootstrap*. Chapman and Hall/CRC.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, Louis Rouillard, Thomas Mesnard, Geoffrey Cideron, Jean bastien Grill, Sabela Ramos, Edouard Yvinec, Michelle Casbon, Etienne Pot, Ivo Penchev, and 197 others. 2025. [Gemma 3 technical report](#). *Preprint*, arXiv:2503.19786.

- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Korenev, Arthur Hinsvark, and 542 others. 2024. [The Llama 3 herd of models](#). *Preprint*, arXiv:2407.21783.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. [On calibration of modern neural networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. [Measuring massive multitask language understanding](#). In *International Conference on Learning Representations*.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, Scott Johnston, Sheer El-Showk, Andy Jones, Nelson Elhage, Tristan Hume, Anna Chen, Yuntao Bai, Sam Bowman, Stanislaw Fort, and 17 others. 2022. [Language models \(mostly\) know what they know](#). *Preprint*, arXiv:2207.05221.
- Adam Tauman Kalai, Ofir Nachum, Santosh S. Vempala, and Edwin Zhang. 2025. [Why language models hallucinate](#). *Preprint*, arXiv:2509.04664.
- Jixuan Leng, Chengsong Huang, Banghua Zhu, and Jiaxin Huang. 2025. [Taming overconfidence in LLMs: Reward calibration in RLHF](#). In *The Thirteenth International Conference on Learning Representations*.
- Rensis Likert. 1932. A technique for the measurement of attitudes. *Archives of Psychology*, 22(140):55.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022a. [Teaching models to express their uncertainty in words](#). *Transactions on Machine Learning Research*.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022b. [TruthfulQA: Measuring how models mimic human falsehoods](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3214–3252, Dublin, Ireland. Association for Computational Linguistics.
- Beier Luo, Shuoyuan Wang, Sharon Li, and Hongxin Wei. 2025. [Your pre-trained LLM is secretly an unsupervised confidence calibrator](#). In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- Sabrina J. Mielke, Arthur Szlam, Emily Dinan, and Y-Lan Boureau. 2022. [Reducing conversational agents’ overconfidence through linguistic calibration](#). *Transactions of the Association for Computational Linguistics*, 10:857–872.
- Preetum Nakkiran, Arwen Bradley, Adam Goliński, Eugene Ndiaye, Michael Kirchhof, and Sinead Williamson. 2025. [Trained on tokens, calibrated on concepts: The emergence of semantic calibration in LLMs](#). *Preprint*, arXiv:2511.04869.
- OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, and 262 others. 2024. [GPT-4 technical report](#). *Preprint*, arXiv:2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 35, pages 27730–27744. Curran Associates, Inc.
- Mahdi Pakdaman Naeini, Gregory Cooper, and Milos Hauskrecht. 2015. [Obtaining well calibrated probabilities using bayesian binning](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 29(1).
- Ethan Perez, Sam Ringer, Kamile Lukosiute, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Benjamin Mann, Brian Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela Amodei, and 44 others. 2023. [Discovering language model behaviors with model-written evaluations](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 13387–13434, Toronto, Canada. Association for Computational Linguistics.
- Mario Sanz-Guerrero, Minh Duc Bui, and Katharina von der Wense. 2025. [Mind the gap: A closer look at tokenization for multiple-choice question answering with LLMs](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 19573–19583, Suzhou, China. Association for Computational Linguistics.
- Mario Sanz-Guerrero and Katharina von der Wense. 2025. [Mitigating label length bias in large language models](#). In *Proceedings of the 14th International Joint Conference on Natural Language Processing and the 4th Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics*, pages 1404–1420, Mumbai, India. The Asian Federation of Natural Language Processing and The Association for Computational Linguistics.
- John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. [Proximal policy optimization algorithms](#). *Preprint*, arXiv:1707.06347.

- Sagi Shaiyer, Mario Sanz-Guerrero, and Katharina von der Wense. 2025. [Asking again and again: Exploring llm robustness to repeated questions](#). *Preprint*, arXiv:2412.07923.
- Mrinank Sharma, Meg Tong, Tomasz Korbak, David Duvenaud, Amanda Askell, Samuel R. Bowman, Esin Durmus, Zac Hatfield-Dodds, Scott R Johnston, Shauna M Kravec, Timothy Maxwell, Sam McCandlish, Kamal Ndousse, Oliver Rausch, Nicholas Schiefer, Da Yan, Miranda Zhang, and Ethan Perez. 2024. [Towards understanding sycophancy in language models](#). In *The Twelfth International Conference on Learning Representations*.
- Katherine Tian, Eric Mitchell, Allan Zhou, Archit Sharma, Rafael Rafailov, Huaxiu Yao, Chelsea Finn, and Christopher Manning. 2023. [Just ask for calibration: Strategies for eliciting calibrated confidence scores from language models fine-tuned with human feedback](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 5433–5442, Singapore. Association for Computational Linguistics.
- Dennis Ulmer, Martin Gubri, Hwaran Lee, Sangdoon Yun, and Seong Oh. 2024. [Calibrating large language models using their generations only](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15440–15459, Bangkok, Thailand. Association for Computational Linguistics.
- Xinpeng Wang, Bolei Ma, Chengzhi Hu, Leon Weber-Genzel, Paul Röttger, Frauke Kreuter, Dirk Hovy, and Barbara Plank. 2024. [“My answer is C”: First-token probabilities do not match text answers in instruction-tuned language models](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 7407–7416, Bangkok, Thailand. Association for Computational Linguistics.
- Jerry Wei, Da Huang, Yifeng Lu, Denny Zhou, and Quoc V. Le. 2024. [Simple synthetic data reduces sycophancy in large language models](#). *Preprint*, arXiv:2308.03958.
- Frank Wilcoxon. 1945. [Individual comparisons by ranking methods](#). *Biometrics Bulletin*, 1(6):80–83.
- Jiancong Xiao, Bojian Hou, Zhanliang Wang, Ruochen Jin, Qi Long, Weijie J Su, and Li Shen. 2025. [Restoring calibration for aligned large language models: A calibration-aware fine-tuning approach](#). In *Forty-second International Conference on Machine Learning*.
- Miao Xiong, Zhiyuan Hu, Xinyang Lu, Yifei Li, Jie Fu, Junxian He, and Bryan Hooi. 2024. [Can LLMs express their uncertainty? An empirical evaluation of confidence elicitation in LLMs](#). In *The Twelfth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Chiwei Zhu, Benfeng Xu, Quan Wang, Yongdong Zhang, and Zhendong Mao. 2023. [On the calibration of large language models and alignment](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 9778–9795, Singapore. Association for Computational Linguistics.

## A Confidence Elicitation Prompts

Figure 3 in the main text shows the prompts used to measure confidence in an answer provided by the model itself (assistant) and by the user using the P(True) method. Here, we provide the full prompts used for the other two confidence estimation methods: Verbalized Percentage (Figure 8) and Verbalized Linguistic (Figure 9).

Answer from "Assistant"	Answer from "User"
<p><b>User:</b> Answer the following multiple-choice question. Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Assistant:</b> Answer: {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> On a scale from 0% to 100%, how confident are you that your answer is correct?</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{"X\%"}) \quad \forall X \in \{0, 1, \dots, 100\}</math></p>	<p><b>User:</b> The following is a multiple-choice question (with answer). Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Answer:</b> {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> On a scale from 0% to 100%, how confident are you that the answer is correct?</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{"X\%"}) \quad \forall X \in \{0, 1, \dots, 100\}</math></p>

Figure 8: Prompts used to measure confidence in an answer provided by the model itself (left) and by the user (right) using the Verbalized Percentage method.

Answer from "Assistant"	Answer from "User"
<p><b>User:</b> Answer the following multiple-choice question. Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Assistant:</b> Answer: {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> How confident are you that your answer is correct? Respond with one of the following options: "Very low", "Low", "Somewhat low", "Medium", "Somewhat high", "High", "Very high"</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{"X"}) \quad \forall X \in \{\text{"Very low"}, \text{"Low"}, \text{"Somewhat low"}, \text{"Medium"}, \text{"Somewhat high"}, \text{"High"}, \text{"Very high"}\}</math></p>	<p><b>User:</b> The following is a multiple-choice question (with answer). Question: {question} A. {option A} B. {option B} C. {option C} D. {option D}</p> <p><b>Answer:</b> {X} <math>\forall X \in \{A, B, C, D\}</math></p> <p><b>User:</b> How confident are you that the answer is correct? Respond with one of the following options: "Very low", "Low", "Somewhat low", "Medium", "Somewhat high", "High", "Very high"</p> <p><b>Assistant:</b> <math>\rightarrow P(\text{"X"}) \quad \forall X \in \{\text{"Very low"}, \text{"Low"}, \text{"Somewhat low"}, \text{"Medium"}, \text{"Somewhat high"}, \text{"High"}, \text{"Very high"}\}</math></p>

Figure 9: Prompts used to measure confidence in an answer provided by the model itself (left) and by the user (right) using the Verbalized Linguistic method.

In all cases, the model’s confidence is extracted from the completion with the highest log-probability among the options. To account for label length bias (where shorter labels are favored), we apply Normalized Contextual Calibration (Sanz-Guerrero and von der Wense, 2025) by subtracting the log-probability of each option in a neutral context (without the question) from the log-probability obtained in the original prompt.

Regarding the prompts used in open-ended tasks (Section 5.3.3), we use the same templates as in

the multiple-choice setting, but we replace the list of options with a single answer (the one provided by the assistant or the user, depending on the condition) and ask the model to judge its confidence. This represents a standard closed-book, open-ended, zero-shot question answering scenario (Shaier et al., 2025).

## B Full Results

Table 6 shows the full calibration results (ECE and Brier score) on MMLU when explicitly asking for confidence using the three confidence elicitation methods. For each model, we report results for the base version (no instruction-tuning), the instruction-tuned version without chat format, and the instruction-tuned version with chat format using both answer positions (assistant and user). This Table provides a more detailed view of the gap bridged by our proposed method between base and instruction-tuned models. In some cases, presenting the answer as part of the user message in the chat template not only recovers the calibration of the base model but even surpasses it.

Model	IT	Chat	Answer	ECE ( $\downarrow$ )			Brier ( $\downarrow$ )		
				P(True)	Verb. (%)	Ling.	P(True)	Verb. (%)	Ling.
Llama 3.1 (8B)	$\times$	$\times$	N/A	0.1910	<b>0.3483</b>	0.3161	0.1930	<b>0.3174</b>	0.2872
	$\checkmark$	$\times$	N/A	0.2796	0.6490	0.3767	0.2315	0.5992	0.3252
	$\checkmark$	$\checkmark$	Assistant	0.2329	0.5133	0.4414	0.2236	0.4500	0.3707
	$\checkmark$	$\checkmark$	User	<b>0.1609</b>	0.4373	<b>0.2064</b>	<b>0.1916</b>	0.3790	<b>0.1547</b>
Llama 3.1 (70B)	$\times$	$\times$	N/A	0.0267	<b>0.1733</b>	<b>0.2745</b>	<b>0.0992</b>	<b>0.1405</b>	<b>0.1983</b>
	$\checkmark$	$\times$	N/A	0.0744	0.2135	0.3151	0.1030	0.1739	0.2288
	$\checkmark$	$\checkmark$	Assistant	0.1130	0.4137	0.4511	0.1825	0.2500	0.3818
	$\checkmark$	$\checkmark$	User	<b>0.0220</b>	0.3937	0.4261	0.1155	0.2280	0.3528
Qwen3 (4B)	$\times$	$\times$	N/A	<b>0.1635</b>	<b>0.4539</b>	0.4123	<b>0.1732</b>	0.4182	0.3598
	$\checkmark$	$\times$	N/A	0.2268	0.5850	0.5021	0.2186	0.5544	0.4241
	$\checkmark$	$\checkmark$	Assistant	0.2140	0.6535	0.5445	0.2145	0.6026	0.4762
	$\checkmark$	$\checkmark$	User	0.2120	0.4605	<b>0.1335</b>	0.2035	<b>0.3756</b>	<b>0.0602</b>
Qwen3 (30B)	$\times$	$\times$	N/A	<b>0.1326</b>	0.4037	0.3589	<b>0.1415</b>	0.3377	0.2806
	$\checkmark$	$\times$	N/A	0.2340	0.5835	0.4565	0.1829	0.5001	0.3650
	$\checkmark$	$\checkmark$	Assistant	0.1935	0.6759	0.5674	0.1902	0.6368	0.4916
	$\checkmark$	$\checkmark$	User	0.1485	<b>0.3289</b>	<b>0.1314</b>	0.1502	<b>0.2828</b>	<b>0.0636</b>
Gemma 3 (4B)	$\times$	$\times$	N/A	<b>0.0209</b>	<b>0.4558</b>	0.3452	<b>0.1882</b>	<b>0.3930</b>	0.3062
	$\checkmark$	$\times$	N/A	0.5779	0.7132	0.5418	0.5645	0.6929	0.4731
	$\checkmark$	$\checkmark$	Assistant	0.6546	0.7130	0.5477	0.6517	0.6953	0.4959
	$\checkmark$	$\checkmark$	User	0.4026	0.5920	<b>0.2797</b>	0.3987	0.5463	<b>0.2499</b>
Gemma 3 (27B)	$\times$	$\times$	N/A	<b>0.0428</b>	0.5213	0.3356	<b>0.1798</b>	0.4465	0.2942
	$\checkmark$	$\times$	N/A	0.2395	0.5465	0.5256	0.2366	0.4785	0.4520
	$\checkmark$	$\checkmark$	Assistant	0.3179	0.6771	0.4560	0.3171	0.6461	0.3813
	$\checkmark$	$\checkmark$	User	0.1919	<b>0.3621</b>	<b>0.2650</b>	0.1911	<b>0.3011</b>	<b>0.2063</b>
<i>Average</i>	$\times$	$\times$	N/A	<b>0.0963</b>	<b>0.3927</b>	0.3404	<b>0.1625</b>	<b>0.3422</b>	0.2877
	$\checkmark$	$\times$	N/A	0.2720	0.5485	0.4530	0.2562	0.4998	0.3780
	$\checkmark$	$\checkmark$	Assistant	0.2877	0.6078	0.5014	0.2966	0.5468	0.4329
	$\checkmark$	$\checkmark$	User	0.1897	0.4291	<b>0.2404</b>	0.2084	0.3521	<b>0.1813</b>

Table 6: **Full calibration results** (ECE and Brier score) for all models in different settings: (1) base model without instruction tuning (“IT”) and without chat format (“Chat”); (2) instruction-tuned model without chat format; (3) instruction-tuned model with chat format answering as assistant; and (4) instruction-tuned model with chat format answering as user. Best results for each model and metric are highlighted in bold.