

Do Language Models Mirror Human Confidence? Exploring Psychological Insights to Address Overconfidence in LLMs

Chenjun Xu^{1*} Bingbing Wen^{1*}
Bin Han¹ Robert Wolfe¹ Lucy Lu Wang^{1,2} Bill Howe¹

¹University of Washington ²Allen Institute for AI
{bingbw, chenjux, bh193, rwolfe3, lucylw, billhowe}@uw.edu

Abstract

Psychology research has shown that humans are poor at estimating their performance on tasks, tending towards underconfidence on easy tasks and overconfidence on difficult tasks. We examine three LLMs, Llama-3-70B-instruct, Claude-3-Sonnet, and GPT-4o, on a range of QA tasks of varying difficulty, and show that models exhibit subtle differences from human patterns of overconfidence: less sensitive to task difficulty, and when prompted to answer based on different personas—e.g., expert vs layman, or different race, gender, and ages—the models will respond with stereotypically biased confidence estimations even though their underlying answer accuracy remains the same. Based on these observations, we propose Answer-Free Confidence Estimation (AFCE) to improve confidence calibration and LLM interpretability in these settings. AFCE is a self-assessment method that employs two stages of prompting, first eliciting only confidence scores on questions, then asking separately for the answer. Experiments on the MMLU and GPQA datasets spanning subjects and difficulty show that this separation of tasks significantly reduces overconfidence and delivers more human-like sensitivity to task difficulty.¹

1 Introduction

Reliable confidence estimates are essential for effective human-machine collaboration (Guo et al., 2017a). Large language models (LLMs), however, are prone to overconfidence (Xiong et al., 2024), which can result in inaccurate predictions when they should abstain (Wen et al., 2025). As these models are increasingly deployed for real-world tasks such as medical diagnosis (Ríos-Hoyo et al., 2024), legal analysis (Deroy et al., 2024), and decision support systems (Xu et al., 2024), their performance directly impacts outcomes that affect human

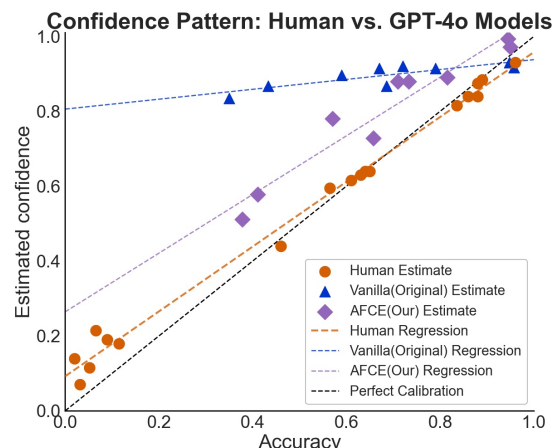


Figure 1: Comparison of Confidence Patterns: Estimated Confidence vs. Actual Accuracy for GPT-4o and Human Participants. This scatter plot highlights the LLM’s tendency to overestimate its own abilities and underestimate the performance of others (referred to as a “average person” in the experiment prompt). The LLM demonstrates more accurate self-assessment compared to its evaluations of other LLM role-playing roles. However, its overestimation and underplacement problems are significantly worse than that of humans. Human data are from Moore and Healy (2008)’s paper.

lives. Overconfidence in LLMs can lead to significant errors (Zhou et al., 2023), reduced trust (Kim et al., 2024), and potentially harmful downstream consequences (Li, 2023; Wen et al., 2024a). Prior work (Aher et al., 2023a; Park et al., 2023) suggests that AI can reflect collective human-like behaviors, while also introducing new risks, such as amplification of misinformation. Therefore, understanding whether LLMs exhibit overconfidence in ways that parallel or exceed human patterns of overconfidence can inform improvements in reliability and safety in real-world applications.

Human overconfidence is recognized as a significant cognitive bias (Kruger and Dunning, 1999). Moore and Healy (2008) reconcile experimental findings that individuals tend to (i) overestimate

*Equal contribution

¹Code: <https://github.com/chenjux/AFCE>

their own abilities on difficult tasks and underestimate them on easy tasks shown in Figure 1, and (ii) estimate that they outperform relative to others on easy tasks (overplacement) and underperform relative to others on hard tasks (underplacement). The authors [Moore and Healy \(2008\)](#) explain these phenomena using an information theoretic model demonstrating individuals’ regressive estimates of their performance and even more regressive estimates of others’ performance: on easy tasks, they underestimate their own success and others’ even more so; on hard tasks, they overestimate their own performance and others’ to an even greater degree.

Current confidence estimation approaches such as vanilla verbalized confidence reveals persistent overconfidence and highlights a disconnect between model-reported confidence and actual task performance as shown in Figure 1. Motivated by these observations, we propose Answer-Free Confidence Estimation (AFCE), which separates confidence estimation and answer generation into separate stages, and find that overconfidence effects as measured by Expected Calibration Error (ECE) are reduced in challenging tasks where overconfidence is pronounced. These results suggest that verbalized confidence methods should not assume human-like behavior in their design.

We further use AFCE to explore whether confidence patterns observed in humans, as described by [Moore and Healy \(2008\)](#), are also present in LLMs. We uncover three distinct phenomena. (i) Models’ confidence scores are comparatively insensitive to task difficulty and exhibit only a weak correlation with actual accuracy, unlike human patterns reported by [Moore and Healy \(2008\)](#). (ii) When prompted with occupational personas, the model reflects stereotypical confidence levels (e.g., “layman” lower than “expert”) regardless of performance. (iii) Adding demographic cues (e.g., gender, race, age) further reduces expressed confidence, even when accuracy stays the same.

Our work investigates three core questions: (i) Sensitivity to Task Difficulty: Is a model’s expressed confidence calibrated to task difficulty, and does it reflect the same over- and underestimation patterns observed in human judgment? (ii) Over- and Underplacement Across Expertise: Do models demonstrate placement biases when estimating the performance of others, particularly when adopting personas with varying levels of expertise? (iii) Demographic Bias in Confidence Expression: Do models exhibit systematic confidence biases when

conditioned on demographic attributes such as race, gender, or age?

We summarize our contributions below:

- We evaluate LLMs’ confidence estimations across tasks of varying difficulty. We observe that, similar to results in human subjects, LLMs exhibit underconfidence in task performance on easy tasks and overconfidence on hard tasks. However, model confidence estimates are less sensitive to task difficulty than human confidence estimates, suggesting a different mechanism mediates self-elicited confidence in LLMs.
- We propose Answer-Free Confidence Estimation (AFCE), a method decoupling confidence estimation from answer generation to improve confidence calibration on challenging tasks and enables comparisons between human and LLM confidence patterns. We demonstrate the effectiveness of this method in three LLMs (LLaMa-3-70B, Claude-3-Sonnet, and GPT-4o).
- We prompt the LLMs with personas from various levels of expertise to investigate over- and underplacement in LLMs and find that models consistently express lower confidence for “Randomly chosen person” persona or “layman” persona and higher confidence for “expert” persona, despite similar task accuracy. These results suggest verbalized confidence is influenced more by persona-based bias than actual performance.
- We assess how demographic personas influence the model’s confidence estimation. We find that LLMs tend to be underconfident when adopting *any* human persona (with the exception of GPT-4o), but that the degree of underconfidence exhibits stereotypical biases: an Asian persona is more confident than other races; a female persona is less confident than other gender identities; a middle-aged persona tends to be more confident than other age groups. These biases highlight the importance of considering demographic factors in confidence calibration, especially as role-playing techniques become more pervasive.

2 Related Work

We review related work on human overconfidence and confidence elicitation methods for LLMs.

Human Overconfidence Overconfidence refers to an unjustified belief in one’s knowledge and abilities ([Kruger and Dunning, 1999](#)), leading to undesirable outcomes in domains such as medicine

(Berner and Graber, 2008), politics (van Prooijen, 2021), and science (Light et al., 2022). Models to explain overconfidence have been broadly considered (e.g., Dunning-Kruger (Kruger and Dunning, 1999), or recent results from Sanchez and Dunning (2024) showing that those with intermediate knowledge may be the most overconfident). In this paper, we focus on the experiments of Moore and Healy (2008), whose influential unifying model explained a variety of previous findings.

In this study, we adopt the overconfidence measure from Moore and Healy (2008) including over(under)estimation and over(under)placement.

Confidence Elicitation in Language Models Previous methods for eliciting confidence have primarily relied on white-box approaches, which have estimated confidence using token likelihoods (Wang et al., 2024a) and internal state-based methods (Kadavath et al., 2022; Kuhn et al., 2023). While effective, these techniques require internal access to the model, making them less applicable to models served over closed APIs, like GPT-4 (Achiam et al., 2023). Verbalized confidence approaches (Tian et al., 2023), primarily the vanilla method, appropriate to such models (*i.e.*, prompting the model to produce confidence estimates in its output) tend to produce uniformly high estimations of model confidence, usually between 80% and 100% (Mielke et al., 2022; Xiong et al., 2024). To improve these estimates, some studies introduce consistency-based methods (Lin et al., 2023; Xiong et al., 2024) to mitigate overconfidence. Other studies (Kumar et al., 2024; Tian et al., 2023) investigated the correlation between verbalized uncertainty and token probability and showed GPT-4o has strongest confidence-probability alignment across variety of tasks.

In this study, we adopt these widely used confidence elicitation methods as baselines, but demonstrate a surprising divergence from human behavior that suggests a problematic decoupling of confidence estimation from answer generation. Indeed, prompting models to estimate confidence without producing answers reduces overconfidence and outperforms baseline methods on hard tasks.

Roly-playing with Language Model LLMs are increasingly employed to simulate human personas. Recent studies (Argyle et al., 2023; Aher et al., 2023b; Park et al., 2022; Wen et al., 2024b) provide empirical evidence that LLM-driven simulations can replicate social science experiments and online forums with consistency comparable to data ob-

tained from human participants. Likewise, Wang et al. (2024b); Jiang et al. (2023) show that advanced role-playing agents exhibit personalities closely aligned with human perceptions, underscoring the effectiveness of role-playing approaches.

Such methodologies underscore the emergence of LLM-based role-play as a versatile, powerful tool across multiple domains (Mbakwe et al., 2023). In this paper, we consider the interaction between role-playing and confidence estimation to investigate whether LLMs exhibit overplacement patterns akin to those observed in humans. We also address associated limitations and ethical considerations, motivating continued research into best practices for employing LLMs in role-playing contexts.

Bias in LLM-Driven Computational Social Science LLMs are known to poorly represent certain groups that make up a significant portion of the population (e.g., age 65+) (Santurkar et al., 2023). People perceive ChatGPT as predominantly male when asked about its gender (Wan et al., 2023), particularly about its core capabilities such as text summarization (Wong and Kim, 2023). Moreover, ChatGPT has been shown to generate gender-biased responses (Wan et al., 2023; Hada et al., 2023). Previous work (Feng et al., 2023) revealed that pretrained LLMs do have political leanings in pretraining corpora, propagating social biases into hate speech predictions and misinformation detectors. Schramowski et al. (2022) showed recent LLMs contain human-like biases about right and wrong behaviors, reflecting existing ethical and moral norms of society. Dong et al. (2024) confirmed a substantial degree of in-group and out-group bias of LLMs across languages and personas.

These latent (and in some applications, explicit) demographic roles can confound confidence estimation: roles associated with over- or underperformance in the training data are associated with verbalized over- or underconfidence estimates, even when actual performance is consistent.

3 Data & Models

Datasets Following the experiment design of Moore and Healy (2008),² which examined question banks spanning various difficulty levels (high school, college, expert) and subjects (physics, chemistry, biology), we use two datasets:

- MMLU (Hendrycks et al., 2021), a collection

²We are unable to repeat the exact questions used in Moore and Healy (2008), as they were not available.

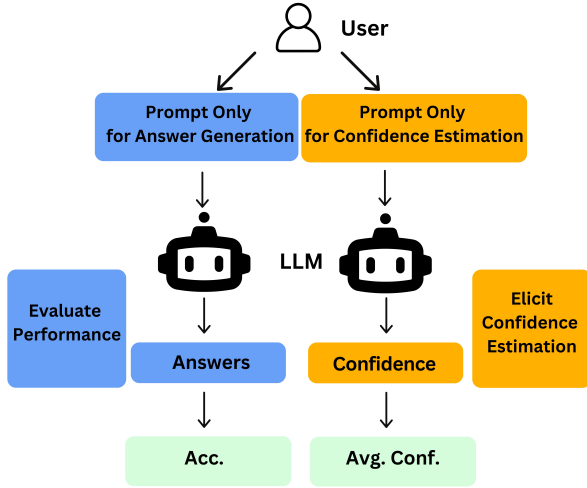


Figure 2: AFCE gathers confidence levels for a set of questions without requiring answers, thereby separating confidence estimation from the answering process. We utilize this approach to simulate human test confidence in psychology and compare the confidence patterns of humans and LLMs.

of domain-specific multiple-choice questions across 57 subjects at multiple educational stages. We select MMLU questions from High School and College in the subject areas of Physics, Chemistry, and Biology.

- GPQA (Rein et al., 2023), a dataset of multiple-choice questions created by experts (*i.e.*, individuals holding or pursuing a PhD). We select questions from physics, chemistry, and biology to represent expert-level difficulty.

Each (subject, difficulty) pair is considered a distinct subtask. To afford a meaningful comparison, we calculate accuracy and confidence in a manner consistent with Moore and Healy (2008), where 10 questions are evaluated in each prompt. Further dataset details can be found in Appendix Table 6.

Models We examine the widely used open-source LLM Meta-Llama-3-70B-Instruct (LLaMA-3-70B), as well as two API-based models, Claude-3-sonnet-20240229 (Claude-3) and GPT-4o-2024-05-13 (GPT-4o). To enhance reproducibility and limit variability in model outputs, we set temperature to 0 and top- p sampling to 1. Model details are provided in Appendix Table 4. Additional experiments on Gemma and Mistral, along with their results, are presented in Appendix Figures 7 and 8.

4 Over(Under) Estimation Across Task Difficulty

In this section, we study the relationship between the LLMs’ confidence estimation and task diffi-

culty. We compare our Answer-Free Confidence Estimation (AFCE) method with other commonly used confidence estimation methods.

4.1 Experiment Setup

Baselines. We consider five baseline methods for model confidence estimation: three verbalized approaches that prompt the model to express confidence directly; one sampling-based method that infers confidence from output variability across generations; and one probability-based method that uses the token probability as confidence proxy.

- Vanilla Verbalized Confidence (Lin et al., 2022), which prompts the model with "Read the question, provide your answer, and report your confidence in this answer".
- Top- k Prompting Verbalized Confidence (Tian et al., 2023), which prompts the model to provide "your K best guesses and the probability that each option is correct (0% to 100%) for the following question".
- Quiz-Like Prompting inspired by Moore and Healy (2008), which prompts the model to "Answer the following 10 questions and estimate how many were answered correctly".
- Sampling-based. Sampling strategy refers to self-random sampling strategy (with three samples) combined with Avg-Conf Aggregation, as proposed by (Xiong et al., 2024).
- Probability-based. We use the probability assigned to the first token as confidence (Wang et al., 2024a).

Our Method. We propose *Answer-Free Confidence Estimation (AFCE)*, which distinguishes task performance evaluation from confidence estimation shown in Figure 2. To evaluate performance, we prompt the model with "Please answer the following 10 questions by selecting only the option letter," and we use the model’s responses to compute its accuracy. We separately obtain the model’s confidence by prompting the model to "Read the questions and estimate how many you can answer correctly (choose a number from 0-10)." We hypothesize that task performance and confidence estimation are mediated by different mechanisms, such that task execution can confound confidence estimation, leading to overconfidence. We include the prompt template in Appendix Table 7.

Evaluation We use Expected Calibration Error (ECE) (Guo et al., 2017b) with 10 bins to evaluate confidence calibration, which quantifies the difference between the confidence and actual accuracy.

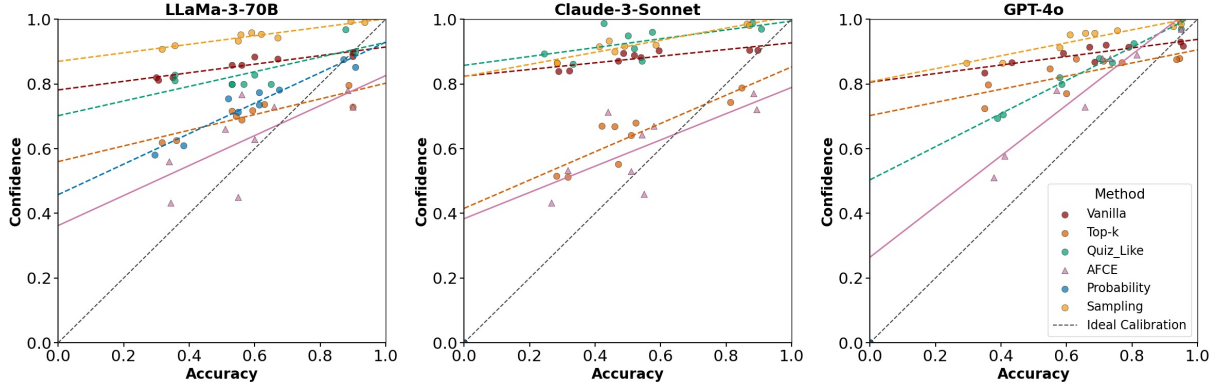


Figure 3: AFCE reduces overconfidence across models and improves sensitivity to task difficulty for GPT-4o. For GPT-4o, AFCE produces a steeper regression slope that aligns more closely with the ideal calibration line. For Claude-3, AFCE and top-k perform comparably, both exceeding other methods. The relatively lower regression line for AFCE for all models suggests reduced overconfidence.

Method	High School						College						Expert						AvE
	Physics		Chemistry		Biology		Physics		Chemistry		Biology		Physics		Chemistry		Biology		
	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	Acc	ECE	
LLaMA-3-70B																			
Vanilla	60.0	28.4	67.0	20.7	90.0	2.1	53.0	32.8	56.0	29.9	90.0	1.4	35.0	47.0	30.6	50.7	54.3	29.9	27.0
Top-K	59.3	12.7	63.0	12.2	88.7	9.3	56.0	16.8	53.0	19.1	87.9	10.8	36.1	26.4	31.7	32.0	54.3	15.7	17.9
Sampling	62.0	33.4	67.0	27.8	89.4	10.0	59.0	36.9	55.0	41.0	93.6	5.5	35.6	56.4	31.7	59.2	55.7	39.5	34.4
Probability	61.3	17.7	67.5	11.5	87.1	5.2	55.0	16.4	52.0	23.6	90.7	7.7	38.3	25.5	29.4	28.6	61.4	14.2	15.1
Quiz	56.7	23.3	65.0	15.0	87.7	9.0	53.0	27.0	53.0	27.0	90.7	5.0	35.6	45.6	35.6	47.2	60.0	22.9	24.7
AFCE	56.0	20.7	66.0	11.0	88.4	10.3	51.0	15.0	55.0	6.0	90.0	17.1	34.4	16.7	33.9	22.2	60.0	11.4	14.5
Claude-3-sonnet																			
Vanilla	48.7	40.9	59.5	30.9	89.7	2.3	52.0	37.0	54.0	35.9	87.1	4.2	28.9	55.6	32.2	51.8	47.1	39.9	23.2
Top-K	42.0	25.2	52.5	15.5	84.8	6.9	46.0	20.9	51.0	13.2	81.4	8.1	31.7	19.5	28.3	24.4	47.1	12.6	21.7
Sampling	41.3	50.2	58.5	33.7	87.1	11.5	46.0	44.0	49.0	42.7	86.4	12.2	28.3	62.3	28.3	58.2	44.3	49.1	29.0
Quiz	42.7	56.0	57.5	38.5	88.1	11.0	50.0	45.0	52.0	39.0	90.7	7.9	24.4	65.0	33.3	52.8	54.3	32.9	17.9
AFCE	44.0	27.3	58.0	9.0	88.4	11.3	51.0	2.0	55.0	9.0	89.3	17.1	26.7	16.7	31.7	21.7	54.3	10.0	13.2
GPT-4o																			
Vanilla	72.0	20.6	79.0	12.5	94.8	2.9	67.0	25.3	59.0	31.1	95.7	4.0	43.3	43.4	35.0	48.5	68.6	20.1	33.2
Top-K	71.3	15.7	77.0	10.9	94.5	6.7	63.0	26.7	55.0	30.6	93.6	6.0	36.1	44.3	35.0	37.4	60.0	17.1	16.3
Sampling	68.7	28.0	75.5	21.7	92.6	5.3	61.0	34.3	58.0	35.2	95.0	3.6	40.6	45.9	29.4	57.0	65.7	30.0	40.4
Quiz	74.0	12.7	80.5	12.0	94.8	4.8	70.0	18.0	58.0	24.0	95.0	5.0	40.6	32.2	38.9	30.6	58.6	21.4	38.7
AFCE	73.3	14.7	81.5	9.5	94.5	6.1	71.0	17.0	57.0	21.0	95.0	3.6	41.1	21.1	37.8	16.1	65.7	10.0	13.8

Table 1: Confidence elicitation and performance comparison across models, subjects and difficulty levels. AFCE significantly reduces overconfidence and achieves better calibration performance compared to other baseline methods, especially for challenging tasks. Acc: Accuracy, ECE: Expected Calibration Error. AvE: Average ECE. All values are percentages.

4.2 Results

AFCE can produce confidence estimation that is sensitive to task difficulty in GPT-4o. Figure 3 shows that GPT-4o’s calibration slope becomes significantly steeper relative to the ideal line when applying our method, indicating enhanced sensitivity despite not achieving perfect calibration. We interpret difficulty through two lenses: (1) stated education level (high school, college, expert) and (2) actual task performance. While LLaMA-3-70B and Claude-3 exhibit relatively flat confidence curves regardless of accuracy, GPT-4o’s confidence estimation is more sensitive to performance. This

flatness suggests that LLaMA-3-70B and Claude-3 rely on a “standard” confidence level, limiting the effectiveness of verbalized confidence-elicitation strategies. Figure 6 in the Appendix further illustrates the relationship between confidence estimation and (subject, difficulty).

AFCE consistently outperforms other baseline approaches in calibration performance across models on challenging tasks. As shown in Table 1, we found while LLMs exhibit lower accuracy on expert-level tasks, they sometimes exhibit stronger performance on college-level subjects than on high school-level subjects, breaking

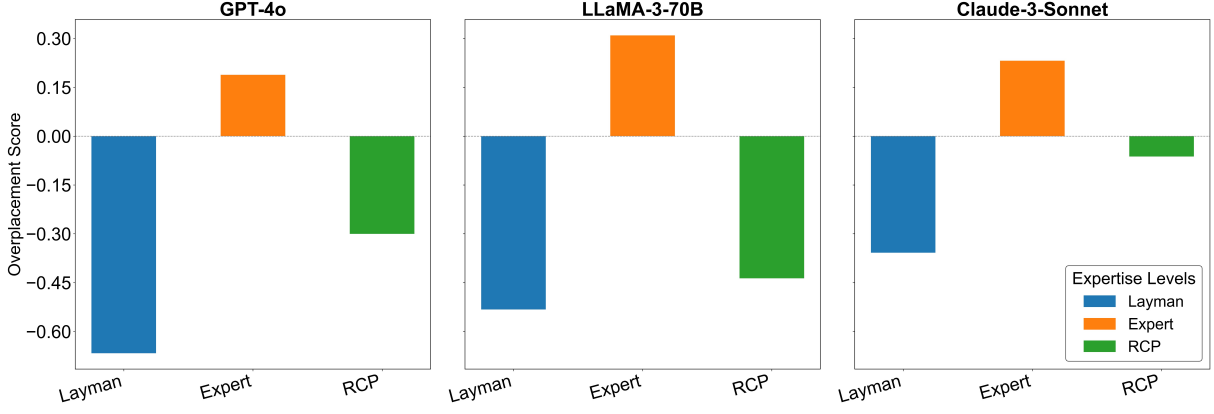


Figure 4: **Overplacement Score** quantifies the degree of overplacement, calculated as: $(\text{Confidence}_{\text{Estimate Others}} - \text{Accuracy}_{\text{Others}}) - (\text{Confidence}_{\text{Self-Estimate}} - \text{Accuracy}_{\text{Self}})$. All models exhibit overplacement towards *Expert* persona and underplacement towards *Layman* persona. GPT-4o and LLaMA-3-70B also show underplacement in the *Randomly Chosen Person* role, whereas Claude-3 demonstrates better calibration.

with typical human judgments of task difficulty. As the difficulty increases—particularly at the expert level—AFCE demonstrates consistently stronger improvements in ECE compared to all baselines. Table 1 illustrates this trend: the progression from high school to college to expert difficulty highlights bold (i.e., best) results for AFCE at higher levels. For GPT-4o, applying AFCE decreases average ECE by 58.4% compared to the vanilla prompt method, 63.8% relative to the quiz-like prompt, 65.8% against the sampling-based method, and 15.3% compared to top- k prompts, outperforming all other baselines.

Figure 3 demonstrates that the AFCE method substantially alleviates overconfidence across models. The corresponding regression line for the AFCE approach (purple line) is consistently lower than those associated with other methods. We posit that task-insensitive overconfidence in the vanilla case results from the epistemically intensive process of generating factual information (Teplica et al., 2025) dominating the reasoning, leading the models to default to a typical confidence answer regardless of question difficulty. But when the tasks are separated, the model is able to fully attend to confidence estimation and become more accurate, more sensitive to task difficulty, and therefore more human-like. Additionally, omitting answer generation may prevent the model from “overthinking” (Cuadron et al., 2025) which may lead to overconfidence. We intend to investigate these underlying mechanisms and their relationship to human cognition in future work aimed at expanding the utility of AFCE.

Given that AFCE outperforms baselines, it will be employed in the subsequent overplacement and demographic bias experiments.

5 Over(Under)placement Across Levels of Expertise

In psychology, overplacement is similar to overconfidence but refers to an inaccurate belief about one’s abilities, performance, or qualities *compared to others*, often overestimating one’s relative standing (Moore and Healy, 2008).

5.1 Experiment Setup

We adapt the experiments on human subjects (Moore and Healy, 2008) for LLMs by prompting language models to adopt the personas of other individuals to estimate comparative confidence. Specifically, We prompt the model to adopt the persona with a particular expertise level then to 1) answer the questions and 2) estimate its confidence using AFCE. We specifically instruct the model to adopt the persona of a “randomly chosen person”, an “expert” in the subject under consideration, and a “layman” with regard to the subject. In prompting language models to adopt personas (Shanahan et al., 2023; Wang et al., 2023; Hagendorff, 2023; Shah et al., 2023), we follow recent work on using LLMs for simulation in computational social science (Aher et al., 2023b), as well as assessments of model bias and fairness (Cheng et al., 2023). The template for overplacement prompts is provided in Appendix Table 8. We measure overplacement as the difference between over(under)confidence in others and over(under)confidence in self (see

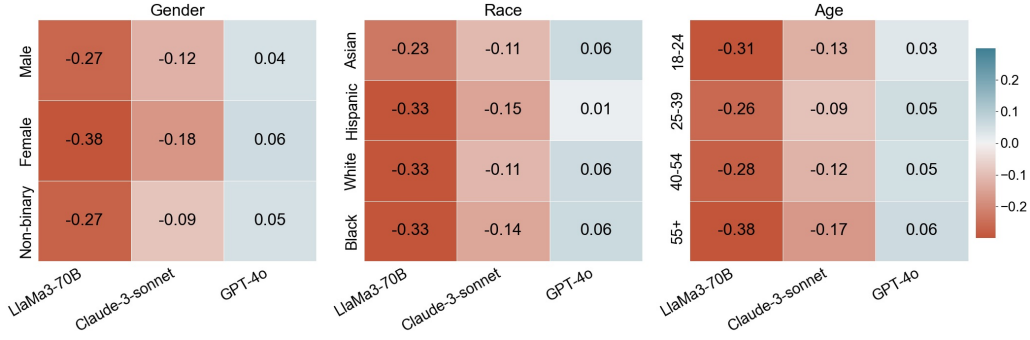


Figure 5: $\Delta_{\text{Demographic}}$ measures the gap between confidence and accuracy (Confidence-Accuracy) across various demographic groups (gender, race, or age). GPT-4o demonstrates more balanced confidence estimations while LLaMA-3-70B and Claude-3 show consistent underconfidence across all demographic groups.

caption of Figure 4).

5.2 Results

LLMs exhibit substantial overplacement toward expert personas, while demonstrating notable underplacement toward average individuals and layperson personas. Figure 4 compares the extent of overplacement across different expertise personas for LLaMA-3-70B, GPT-4o, and Claude-3. All models consistently show pronounced overplacement for expert roles, contrasted with clear underplacement for average and layperson roles. Additionally, models differ significantly in how they assess the performance of a randomly selected individual: LLaMA-3-70B and GPT-4o display marked underplacement biases against this neutral baseline, whereas Claude-3 provides relatively balanced estimates. Despite these overplacement biases, actual model accuracy remains relatively consistent across different personas, indicating a systematic tendency to overestimate confidence in expert roles and underestimate confidence in layperson roles. Overall, the models’ confidence judgments about others seem disconnected from their true capabilities when adopting persona-based prompts. We expect that this mismatch could prove problematic for the range of research that now utilizes persona-prompted LLMs in social scientific simulations (Aher et al., 2023b; Ziems et al., 2023).

6 Bias in Confidence Estimation Across Demographic Personas

Here, we combine the AFCE method and role-playing to analyze variations in confidence estimation across different protected attributes (race, gender, and age).

6.1 Experiment Setup

We utilize LLMs within a structured role-playing framework (following the setup described in §5) to simulate individuals from various demographic groups as follows:

$$\begin{cases} \text{Race} \in \{\text{White, Black, Asian, Hispanic}\} \\ \text{Gender} \in \{\text{Male, Female, Non-binary}\} \\ \text{Age} \in \left\{ \begin{array}{l} \text{Young adult (18–24), Adult (25–39),} \\ \text{Middle-aged (40–54), Senior (55+)} \end{array} \right\} \end{cases}$$

Overplacement prompts are provided in Appendix Table 9.

6.2 Results

LLMs tend to be underconfident when adopting *any* human persona (with the exception of GPT-4o), but that the degree of underconfidence exhibits stereotypical biases. Figure 5 illustrates the accuracy and confidence estimates of LLMs when role-playing different assignments of gender, race, and age. Variations in $\Delta_{\text{Demographic}}$ are mainly driven by confidence estimation, as prediction differences remain minimal, suggesting independent confidence estimation and question-answering mechanisms.

LLaMA3-70B and Claude-3 underestimate confidence when playing female-identifying roles, with $\Delta_{\text{Demographic}}$ ranking as Non-binary \geq Male $>$ Female. GPT-4o demonstrates nearly uniform $\Delta_{\text{Demographic}}$ values, indicating reduced gender biases. Under racial assignments, LLaMA3-70B shows larger negative $\Delta_{\text{Demographic}}$ for Asians, while Claude-3 exhibits higher negative values for White and Asian individuals. GPT-4o displays smaller disparities, indicating more balanced racial confidence estimations. For age groups, LLaMA3-70B and Claude-3 show underestimation for 18–24

NQ-open	Acc	AvC	ECE
Quiz-Like	74.0	78.0	6.0
Vanilla	74.0	77.2	6.0
AFCE	74.0	75.0	4.0
SimpleQA	Acc	AvC	ECE
Quiz-Like	36.0	78.0	42.0
Vanilla	31.0	87.0	56.0
AFCE	36.0	25.0	6.0

Table 2: AFCE’s performance on open-ended generation QA using GPT-4o indicates that its calibration surpasses that of other verbalized baselines. Acc: Accuracy, AvC: Average Confidence, ECE: Expected Calibration Error. All values are percentages.

and 55+ groups but favor 25–39 and 40–54 cohorts while GPT-4o maintains minimal bias.

The LLaMA and Claude models show consistent underconfidence across all demographic groups, with lower negative $\Delta_{\text{Demographic}}$ values, reflecting a misalignment between their confidence levels and actual performance. Despite lower confidence, their accuracy often remains unaffected, especially in complex demographic contexts, highlighting calibration issues. In contrast, GPT-4o demonstrates balanced confidence estimations across races, ages, and genders, effectively mitigating biases. Its consistent performance makes it well-suited for research requiring fairness and inclusivity, paving the way for equitable insights in social science studies.

7 Ablation Study

We investigate whether AFCE generalizes to open-ended QA and whether it remains robust to variations in question order and the number of questions within a prompt. Accordingly, we conducted the following two experiments using GPT-4o.

AFCE Generalizes to Open-Ended Generation QA. We randomly sampled 100 open-ended questions each from NaturalQuestions-open (Park et al., 2023) and SimpleQA (Wei et al., 2024). Experimental results presented in Table 2 indicate that AFCE consistently outperforms other baselines across both datasets. Although our primary focus was on multiple choice questions for cross-domain comparisons and alignment with human subject studies, these findings suggest that AFCE extends beyond structured formats.

AFCE Demonstrates Robustness Under Variations in Question Order and Group Size. We conducted two experiments: (1) reducing the group size to five questions, and (2) randomizing question

order. We observed no substantial differences in performance under either condition. As shown in Table 10, AFCE maintains its effectiveness across these configurations.

8 Discussion

Our study provides insights into LLM confidence calibration informed by the psychology literature, exposing subtle differences between LLM and human behavior with implications for the understanding of confidence estimation and the design of confidence calibration techniques.

While our method AFCE effectively mitigates overconfidence in challenging tasks, it may lead to underconfidence in easier tasks. This effect could stem from the decoupling of answer generation from confidence estimation. However, the underlying mechanism by which this decoupling influences confidence estimation remains unclear. It is possible that AFCE adjusts confidence more aggressively for challenging questions, while its influence is attenuated for simpler ones. Further investigation is needed to clarify how and why this decoupling shapes confidence across varying task difficulties.

We find that LLMs can be easily manipulated by expert personas, leading to inflated confidence scores. This highlights key directions for future work: first, understanding how confidence signals influence human trust and decision-making in real-world settings; second, developing dynamic confidence calibration methods tailored to specific applications or user profiles to better manage risk. Bias in confidence calibration across demographics, such as race, gender, and age, poses additional concerns. While recent models like GPT-4o show progress in reducing these biases, other LLMs still exhibit notable disparities. This warrants further investigation, as addressing confidence calibration alone may be insufficient to ensure fairness and reliability. Further experiments should explore how demographic-aware verbalized confidence estimation can help mitigate bias across different groups.

9 Conclusion

We analyze LLM overconfidence from a psychological perspective, showing that LLMs are less sensitive to task difficulty and can drastically shift their confidence when adopting personas—despite minimal changes in actual performance. This suggests that answer generation and confidence estimation involve distinct processes differently influenced by

prompts and biases. We propose AFCE, which significantly reduces ECE for difficult reasoning questions prone to overconfidence. We also find that most LLMs (except GPT-4o) become underconfident when adopting Asian, age 55+, or female personas, again without affecting task performance. Likewise, a model’s self-assessment is typically overconfident, yet it claims higher confidence as an expert and lower confidence as a novice, with no notable impact on performance.

10 Limitations

In this study, we mainly focus on verbalized confidence elicitation methods which are accessible to all kinds of models. Our investigation shows a promising correlation between a model’s confidence pattern and humans’ confidence pattern, but heavily depends on precise prompting techniques. Besides, although our method outperforms in the ECE metric, ECE has limitations, including its sensitivity to bin definitions and its inability to account for the overall prediction distribution. The analysis was conducted on a select group of LLMs: GPT-4o, Claude-3, LLaMA-3, Mistral and Gemma. These models were chosen for their architectural diversity and representativeness of current state-of-the-art. However, the inclusion of other models, such as Gemini Pro, PerplexityAI or those specialized in specific languages and domains, in future studies would likely reveal further interesting findings.

11 Ethical Considerations

This study highlights key ethical concerns in LLM confidence calibration. Disparities in confidence estimation across demographics, such as race, gender, and age, risk perpetuating inequities, particularly in sensitive areas like healthcare, education, and law. Overconfidence, especially in "Expert" roles, may lead to unwarranted trust in incorrect outputs, requiring careful prompt design to mitigate risks. Transparency is essential to communicate model limitations and build user trust, while biases and confidence misalignment may spread misinformation or disproportionately affect marginalized groups. Addressing these issues is crucial for fairness, reliability, and ethical LLM deployment.

Acknowledgements

This work is partially supported by gift funds from the Allen Institute for AI.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Gati Aher, Rosa I. Arriaga, and Adam Tauman Kalai. 2023a. Using large language models to simulate multiple humans and replicate human subject studies. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org.
- Gati V Aher, Rosa I Arriaga, and Adam Tauman Kalai. 2023b. Using large language models to simulate multiple humans and replicate human subject studies. In *International Conference on Machine Learning*, pages 337–371. PMLR.
- Lisa P Argyle, Ethan C Busby, Nancy Fulda, Joshua R Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351.
- Eta S. Berner and Mark L Graber. 2008. *Overconfidence as a cause of diagnostic error in medicine*. *The American journal of medicine*, 121 5 Suppl:S2–23.
- Myra Cheng, Esin Durmus, and Dan Jurafsky. 2023. *Marked personas: Using natural language prompts to measure stereotypes in language models*. *ArXiv*, abs/2305.18189.
- Alejandro Cuadron, Dacheng Li, Wenjie Ma, Xingyao Wang, Yichuan Wang, Siyuan Zhuang, Shu Liu, Luis Gaspar Schroeder, Tian Xia, Huanzhi Mao, et al. 2025. The danger of overthinking: Examining the reasoning-action dilemma in agentic tasks. *arXiv preprint arXiv:2502.08235*.
- Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2024. *Applicability of large language models and generative models for legal case judgement summarization*. *ArXiv*, abs/2407.12848.
- Wenchao Dong, Assem Zhunis, Hyojin Chin, Jiyoung Han, and Meeyoung Cha. 2024. I am not them: Fluid identities and persistent out-group bias in large language models. *arXiv preprint arXiv:2402.10436*.
- Shangbin Feng, Chan Young Park, Yuhao Liu, and Yulia Tsvetkov. 2023. *From pretraining data to language models to downstream tasks: Tracking the trails of political biases leading to unfair NLP models*. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 11737–11762, Toronto, Canada. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. 2017a. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR.

- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017b. [On calibration of modern neural networks](#). In *Proceedings of the 34th International Conference on Machine Learning, ICML 2017, Sydney, NSW, Australia, 6-11 August 2017*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.
- Rishav Hada, Agrima Seth, Harshita Diddee, and Kalika Bali. 2023. [“fifty shades of bias”: Normative ratings of gender bias in GPT generated English text](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 1862–1876, Singapore. Association for Computational Linguistics.
- Thilo Hagendorff. 2023. Machine psychology: Investigating emergent capabilities and behavior in large language models using psychological methods. *arXiv preprint arXiv:2303.13988*, 1.
- 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*.
- Hang Jiang, Xiajie Zhang, Xubo Cao, Cynthia Breazeal, Deb Roy, and Jad Kabbara. 2023. Personallm: Investigating the ability of large language models to express personality traits. *arXiv preprint arXiv:2305.02547*.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield Dodds, Nova DasSarma, Eli Tran-Johnson, et al. 2022. [Language models \(mostly\) know what they know](#). *ArXiv preprint*, abs/2207.05221.
- Sunnie S. Y. Kim, Q. Vera Liao, Mihaela Vorvoreanu, Steph Ballard, and Jennifer Wortman Vaughan. 2024. [“i’m not sure, but...”: Examining the impact of large language models’ uncertainty expression on user reliance and trust](#). *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*.
- Justin Kruger and David Dunning. 1999. Unskilled and unaware of it: how difficulties in recognizing one’s own incompetence lead to inflated self-assessments. *Journal of personality and social psychology*, 77(6):1121.
- Lorenz Kuhn, Yarin Gal, and Sebastian Farquhar. 2023. [Semantic uncertainty: Linguistic invariances for uncertainty estimation in natural language generation](#). *ArXiv preprint*, abs/2302.09664.
- Abhishek Kumar, Robert Morabito, Sanzhar Umbet, Jad Kabbara, and Ali Emami. 2024. [Confidence under the hood: An investigation into the confidence-probability alignment in large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 315–334, Bangkok, Thailand. Association for Computational Linguistics.
- Zihao (Michael) Li. 2023. [The dark side of chatgpt: Legal and ethical challenges from stochastic parrots and hallucination](#). *ArXiv*, abs/2304.14347.
- Nicholas Light, Philip M Fernbach, Nathaniel Rabb, Mugur V Geana, and Steven A Sloman. 2022. Knowledge overconfidence is associated with anti-consensus views on controversial scientific issues. *Science Advances*, 8(29):eabo0038.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. [Teaching models to express their uncertainty in words](#). *Transactions on Machine Learning Research*.
- Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2023. [Generating with confidence: Uncertainty quantification for black-box large language models](#). *ArXiv preprint*, abs/2305.19187.
- Amarachi B. Mbakwe, Ismini Lourentzou, Leo Anthony Celi, Oren J. Mechanic, and Alon Dagan. 2023. [Chatgpt passing usmle shines a spotlight on the flaws of medical education](#). *PLOS Digital Health*, 2(2):1–3.
- 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.
- Don A Moore and Paul J Healy. 2008. The trouble with overconfidence. *Psychological review*, 115(2):502.
- Joon Sung Park, Joseph O’Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. [Generative agents: Interactive simulacra of human behavior](#). In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, UIST ’23, New York, NY, USA. Association for Computing Machinery.
- Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R Bowman. 2023. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022*.
- Alejandro Ríos-Hoyo, Naing Lin Shan, Anran Li, Alexander T. Pearson, Lajos Pusztai, and Frederick M. Howard. 2024. [Evaluation of large language models as a diagnostic aid for complex medical cases](#). *Frontiers in Medicine*, 11.
- Carmen Sanchez and David Dunning. 2024. Intermediate science knowledge predicts overconfidence. *Trends in Cognitive Sciences*, 28(4):284–285.

- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, pages 29971–30004. PMLR.
- Patrick Schramowski, Cigdem Turan, Nico Andersen, Constantin A Rothkopf, and Kristian Kersting. 2022. Large pre-trained language models contain human-like biases of what is right and wrong to do. *Nature Machine Intelligence*, 4(3):258–268.
- Rusheb Shah, Soroush Pour, Arush Tagade, Stephen Casper, Javier Rando, et al. 2023. Scalable and transferable black-box jailbreaks for language models via persona modulation. *arXiv preprint arXiv:2311.03348*.
- Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role play with large language models. *Nature*, 623(7987):493–498.
- Carter Teplica, Yixin Liu, Arman Cohan, and Tim GJ Rudner. 2025. Sciurus: Shared circuits for interpretable uncertainty representations in language models. In *MINT: Foundation Model Interventions*.
- 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.
- Jan-Willem van Prooijen. 2021. [Overconfidence in radical politics](#). *The Psychology of Populism*.
- Yixin Wan, George Pu, Jiao Sun, Aparna Garimella, Kai-Wei Chang, and Nanyun Peng. 2023. [“kelly is a warm person, joseph is a role model”: Gender biases in LLM-generated reference letters](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3730–3748, Singapore. Association for Computational Linguistics.
- Xinpeng Wang, Bolei Ma, Chengzhi Hu, Leon Weber-Genzel, Paul Röttger, Frauke Kreuter, Dirk Hovy, and Barbara Plank. 2024a. "my answer is c": First-token probabilities do not match text answers in instruction-tuned language models. *arXiv preprint arXiv:2402.14499*.
- Xintao Wang, Yunze Xiao, Jen-tse Huang, Siyu Yuan, Rui Xu, Haoran Guo, Quan Tu, Yaying Fei, Ziang Leng, Wei Wang, Jiangjie Chen, Cheng Li, and Yanghua Xiao. 2024b. [InCharacter: Evaluating personality fidelity in role-playing agents through psychological interviews](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1840–1873, Bangkok, Thailand. Association for Computational Linguistics.
- Zekun Moore Wang, Zhongyuan Peng, Haoran Que, Jiaheng Liu, Wangchunshu Zhou, Yuhan Wu, Hongcheng Guo, Ruitong Gan, Zehao Ni, Jian Yang, et al. 2023. Rolellm: Benchmarking, eliciting, and enhancing role-playing abilities of large language models. *arXiv preprint arXiv:2310.00746*.
- Jason Wei, Nguyen Karina, Hyung Won Chung, Yunxin Joy Jiao, Spencer Papay, Amelia Glaese, John Schulman, and William Fedus. 2024. [Measuring short-form factuality in large language models](#). *Preprint*, arXiv:2411.04368.
- Bingbing Wen, Bill Howe, and Lucy Lu Wang. 2024a. [Characterizing LLM abstention behavior in science QA with context perturbations](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 3437–3450, Miami, Florida, USA. Association for Computational Linguistics.
- Bingbing Wen, Chenjun Xu, Bin HAN, Robert Wolfe, Lucy Lu Wang, and Bill Howe. 2024b. [From human to model overconfidence: Evaluating confidence dynamics in large language models](#). In *NeurIPS 2024 Workshop on Behavioral Machine Learning*.
- Bingbing Wen, Jihan Yao, Shangbin Feng, Chenjun Xu, Yulia Tsvetkov, Bill Howe, and Lucy Lu Wang. 2025. Know your limits: A survey of abstention in large language models. *Transactions of the Association for Computational Linguistics (ACL)*.
- Jared Wong and Jin Kim. 2023. Chatgpt is more likely to be perceived as male than female. *arXiv preprint arXiv:2305.12564*.
- 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*.
- Zequ Xu, Lingfeng Guo, Shuwen Zhou, Runze Song, and Kaiyi Niu. 2024. [Enterprise supply chain risk management and decision support driven by large language models](#). *Applied Science and Engineering Journal for Advanced Research*, 3(4):1–7.
- Kaitlyn Zhou, Dan Jurafsky, and Tatsunori Hashimoto. 2023. [Navigating the grey area: Expressions of overconfidence and uncertainty in language models](#). *ArXiv*, abs/2302.13439.
- Caleb Ziems, William B. Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2023. [Can large language models transform computational social science?](#) *Computational Linguistics*, 50:237–291.

A Appendix / supplemental material

A.1 Dataset and Model Details

This section provides comprehensive details on the datasets and language models. Table 3 summarizes all the datasets used in this paper, which are drawn from four sources: **MMLU**, **GPQA**, **SimpleQA**, and **NQ-Open**. Among them, MMLU and GPQA are explicitly structured to include questions of varying difficulty levels, allowing us to analyze how model confidence and performance are affected by task hardness. Table 4 lists the specific versions of the models used in this study.

Dataset	Hardness	Subject	Test Size
MMLU	High School	Physics	170
MMLU	High School	Chemistry	230
MMLU	High School	Biology	340
MMLU	High School	Math	300
MMLU	High School	Computer Science	110
MMLU	College	Physics	110
MMLU	College	Chemistry	110
MMLU	College	Biology	160
MMLU	College	Math	110
MMLU	College	Computer Science	110
MMLU	College	Medicine	200
MMLU	Expert	Medicine	300
GPQA	Expert	Physics	220
GPQA	Expert	Chemistry	210
GPQA	Expert	Biology	100
SimpleQA	general	general	100
NQ-open	general	general	100

Table 3: Dataset statistics.

Model ID	Date Release	Developer
gemma2-9b-it	Jun. 27, 2024	Google
Meta-Llama-3-70B-Instruct	Apr. 18, 2024	Meta
Meta-Llama-3-8B-Instruct	Apr. 18, 2024	Meta
Llama-3.2-90B-Vision-Instruct	Sept. 25, 2024	Meta
Mixtral-8x7B-Instruct-v0.1	Dec. 11, 2023	Mistral
GPT-4o	May 13, 2024	OpenAI
claude-3-sonnet-20240229	Feb. 29, 2024	Anthropic

Table 4: Model statistics.

A.2 Overconfidence under Subject and Difficulty

Figure 6 shows the performance of three models—**GPT-4o**, **LLaMA-3-70B**, and **Claude-3-Sonnet**—across different subjects and difficulty levels. This visualization highlights how both subject area and task hardness influence model confidence, offering additional insights into the calibration behavior of these large language models. Table 5 extends Table 1 by including results from four additional models.

A.3 Overplacement Analysis

Figure 7 presents overplacement results across seven models under different role-playing conditions. It visualizes the gap between confidence and accuracy more intuitively for each role. In comparison, the main text’s Figure 4 focuses on the overall degree of overplacement exhibited by LLMs towards various level of expertise roles, providing a complementary perspective.

A.4 Demographic Bias Results

Figure 8 supplements our demographic bias experiments with results from four additional models. These results provide a more comprehensive view of how demographic attributes—such as gender, race, and age—impact model calibration and bias across different LLM architectures.

A.5 LLM Prompt Templates

This section details the exact prompt templates used in our experiments across various settings, including Answer-Free Confidence Estimation (AFCE), overplacement studies, and demographic bias analyses. Table 6 presents example questions at three different difficulty levels. Table 7 provides a sample prompt for self-estimation, which is used in both the overconfidence and overplacement experiments. Additionally, Table 8 shows an example prompt used to test LLMs’ confidence and accuracy when role-playing different personas. Finally, Table 9 illustrates the prompt format used in demographic experiments involving gender, age, and race.

B Math Formula for ECE

The Expected Calibration Error (ECE) is defined as:

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

Explanation of symbols:

- M : The number of bins or groups into which predictions are divided.
- B_m : The set of predictions in the m -th bin.
- $|B_m|$: The number of predictions in the m -th bin.
- n : The total number of predictions across all bins.

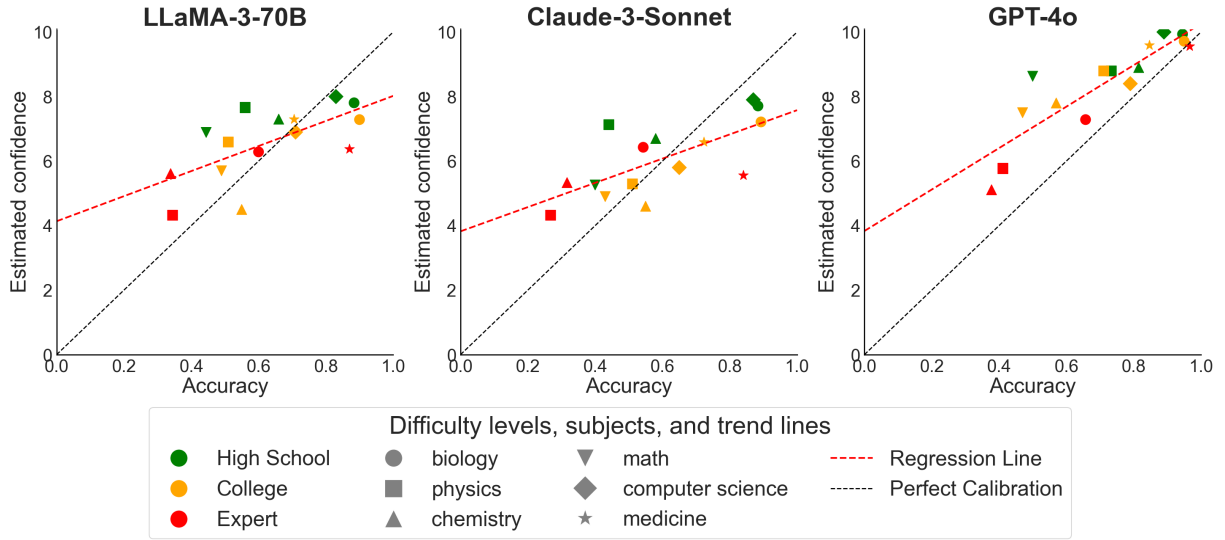


Figure 6: Confidence estimation using our method on tasks with various levels of difficulty. GPT-4o appears especially more sensitive to task difficulty than LLaMA-3-70B and Claude-3-Sonnet.

Method	High School									College									Expert								
	Physics			Chemistry			Biology			Physics			Chemistry			Biology			Physics			Chemistry			Biology		
	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE	Acc	AvC	ECE
Gemma2-9B																											
Vanilla	56.0	98.3	43.1	64.0	98.8	35.1	91.3	99.2	7.9	48.0	99.6	51.6	49.0	99.0	50.0	87.1	99.2	12.1	34.4	99.0	64.5	30.0	98.0	68.3	50.0	96.6	46.6
Top-K	49.3	91.4	42.8	56.5	91.3	35.9	84.5	93.3	10.4	47.0	93.3	48.7	49.0	94.0	45.0	85.0	93.5	9.9	30.6	91.9	62.1	28.3	89.3	62.7	42.9	88.8	46.0
Quiz-like	58.7	96.0	37.3	67.5	95.0	27.5	87.7	99.4	12.3	48.0	93.0	45.0	55.0	96.0	41.0	86.4	99.3	12.9	36.7	92.8	56.1	29.4	96.7	67.2	41.4	97.1	55.7
Ours	<u>57.3</u>	86.0	28.7	<u>66.0</u>	90.0	24.0	<u>87.7</u>	89.7	7.7	49.0	87.0	38.0	55.0	86.0	31.0	87.1	88.6	14.3	<u>34.4</u>	71.7	37.2	30.0	79.4	49.4	<u>42.9</u>	81.4	38.6
LLaMA-3-8B																											
Vanilla	43.3	83.3	40.0	52.0	83.2	31.7	79.7	84.6	5.5	44.0	83.2	39.2	43.0	83.3	40.3	76.4	84.7	9.0	32.2	82.1	49.9	28.3	81.8	53.5	35.7	80.9	45.1
Top-K	33.3	77.2	44.0	44.5	75.9	31.4	70.3	79.3	11.4	46.0	78.8	32.9	43.0	74.9	<u>35.3</u>	73.6	78.6	<u>8.2</u>	28.3	71.2	43.5	24.4	66.6	<u>42.2</u>	30.0	70.5	40.5
Quiz-like	37.3	80.0	42.7	47.5	80.5	33.0	73.5	85.5	11.9	41.0	83.0	42.0	42.0	83.0	41.0	<u>72.9</u>	85.0	12.1	34.4	81.1	46.7	<u>29.4</u>	84.4	55.0	28.6	84.3	55.7
Ours	<u>39.3</u>	78.7	39.3	47.0	78.0	39.3	<u>73.9</u>	80.0	<u>6.1</u>	42.0	78.0	36.0	42.0	72.0	30.0	<u>73.6</u>	80.0	6.4	<u>33.9</u>	53.3	19.4	31.1	61.1	30.0	<u>31.4</u>	68.6	37.1
LLaMA-3.2-90B																											
Vanilla	65.3	92.4	27.8	<u>74.5</u>	88.2	14.7	92.6	93.3	<u>3.3</u>	61.0	90.9	29.9	58.0	85.8	28.8	95.7	91.3	<u>6.0</u>	40.0	83.2	43.2	<u>33.9</u>	81.6	47.8	<u>52.9</u>	83.1	<u>30.3</u>
Top-K	60.7	88.4	28.7	73.5	79.7	9.6	90.6	84.2	12.1	63.0	81.1	21.1	53.0	83.4	31.6	93.6	82.8	11.5	39.4	73.7	<u>35.4</u>	35.6	68.6	34.5	55.7	69.6	18.1
Quiz-like	61.3	86.7	<u>25.3</u>	73.0	87.0	14.0	91.9	97.4	6.1	48.0	83.0	35.0	<u>57.0</u>	83.0	26.0	93.6	95.7	2.1	42.8	81.1	38.3	30.6	83.3	52.8	47.1	93.3	37.1
Ours	<u>62.7</u>	80.0	17.3	75.0	80.0	5.0	80.3	80.6	1.0	49.0	80.0	31.0	<u>57.0</u>	80.0	23.0	<u>95.0</u>	80.0	15.0	<u>42.2</u>	73.9	31.7	33.3	80.0	<u>46.7</u>	44.3	80.0	35.7
Mixtral-8x7B																											
Vanilla	44.0	93.1	49.5	54.5	91.3	36.8	82.6	90.8	8.3	51.0	91.6	40.6	53.0	88.9	35.9	84.3	89.5	6.0	32.8	87.5	54.7	23.3	86.4	63.6	<u>44.3</u>	84.5	40.8
Top-K	<u>39.3</u>	85.1	47.7	43.0	84.3	42.3	65.5	85.0	20.2	43.0	85.0	43.8	40.0	79.8	39.8	67.1	83.0	21.0	23.3	74.0	<u>51.8</u>	22.2	67.5	45.8	31.4	69.8	<u>40.4</u>
Quiz-like	<u>39.3</u>	90.0	50.7	34.5	86.5	52.0	25.2	82.9	57.7	25.0	85.0	60.0	<u>52.0</u>	90.0	38.0	28.6	82.1	53.6	<u>35.0</u>	89.4	54.4	25.0	89.4	64.4	27.1	81.4	54.3
Ours	38.7	87.3	<u>48.7</u>	59.0	90.0	31.0	<u>82.3</u>	90.3	8.1	<u>47.0</u>	85.0	38.0	46.0	81.0	35.0	<u>78.6</u>	87.9	<u>10.7</u>	36.1	70.6	34.4	20.0	68.3	<u>48.3</u>	51.4	74.3	22.9

Table 5: Confidence elicitation and performance comparison for Gemma2-9B, LLaMA-3-8B, LLaMA-3.2-90B, Mixtral-8x7B on Physics, Chemistry, and Biology across three difficulty levels. Acc: Accuracy, AvC: Avg Confidence, ECE: Expected Calibration Error. All values are percentages.

- $\text{acc}(B_m)$: The accuracy of the predictions in the m -th bin.
- $\text{conf}(B_m)$: The average confidence of the predictions in the m -th bin.

// High school

Question: A rigid, solid container of constant volume holds an ideal gas of volume v_1 and temperature T_1 and pressure P_1 . The temperature is increased in an isochoric process. Which of the following is NOT true?

- A. The average speed of the molecules increases.
- B. The pressure increases.
- C. The kinetic energy of the system increases
- D. The volume increases.

Answer: D

// College

Question: A uniform solid disk starts from rest and rolls down an inclined plane without slipping. After some time, what fraction of the disk's total kinetic energy is rotational kinetic energy?

- A. $1/4$
- B. $1/3$
- C. $1/2$
- D. $2/3$

Answer: B

// Expert

In order to calculate the necessary beam to produce a particle X, we ran the following simulation in a High Energy software $e^+e^- \rightarrow X$, where the electron e^- is at rest. We found that the beam energy has to be at least equal to 9.6×10^6 GeV. is the mass of the particle X used for this simulation?

- A. 1091 GeV
- B. 99 GeV
- C. 3132 GeV
- D. 542 GeV

Answer: B

Table 6: Examples from dataset in three task hardness level.

// Prompt for accuracy calculation

prompt = ""Please answer the following questions by selecting only the option letter. Do not provide any explanation. Questions:

- Q1: <>
- Q2: <>
- Q3: <>
- Q4: <>
- Q5: <>
- Q6: <>
- Q7: <>
- Q8: <>
- Q9: <>
- Q10: <>

Now, please answer these questions with the following format:

- 1. <Only Option Letter >
 - 2. <Only Option Letter >
 - 3. <Only Option Letter >
 - 4. <Only Option Letter >
 - 5. <Only Option Letter >
 - 6. <Only Option Letter >
 - 7. <Only Option Letter >
 - 8. <Only Option Letter >
 - 9. <Only Option Letter >
 - 10. <Only Option Letter >
- ""
-

// Prompt for confidence calculation

prompt = f""Read the question, and estimate how many questions you can answer correctly (choose a number from 0-10). Do not provide any explanation.

Questions:

- Q1: <>
- Q2: <>
- Q3: <>
- Q4: <>
- Q5: <>
- Q6: <>
- Q7: <>
- Q8: <>
- Q9: <>
- Q10: <>

Estimate how many questions you can answer correctly: <choose a number from 0-10>.""

Table 7: Prompts used for self-estimation experiments in our approach, independent of the AFCE method.

// Prompt for accuracy calculation
prompt = ""You are a {role}. Please answer the following questions by selecting only the option letter. Do not provide any explanation.
Questions:
Q1: <>
Q2: <>
Q3: <>
Q4: <>
Q5: <>
Q6: <>
Q7: <>
Q8: <>
Q9: <>
Q10: <>
Now, please answer these questions with the following format:
1. <Only Option Letter >
2. <Only Option Letter >
3. <Only Option Letter >
4. <Only Option Letter >
5. <Only Option Letter >
6. <Only Option Letter >
7. <Only Option Letter >
8. <Only Option Letter >
9. <Only Option Letter >
10. <Only Option Letter >
""

// Prompt for confidence calculation
prompt = ""A is an {role} in {subject}. Read the question, and after considering A's ability, estimate how many questions A can answer correctly (choose a number from 0-10). Do not provide any explanation.
Questions:
Q1: <>
Q2: <>
Q3: <>
Q4: <>
Q5: <>
Q6: <>
Q7: <>
Q8: <>
Q9: <>
Q10: <>
Estimate how many questions you think A can answer correctly: <choose a number from 0-10>.""

Table 8: Prompts used for overplacement experiments in our approach, independent of the AFCE method.

// Prompt for accuracy calculation
prompt = ""You are a {race} {gender} in the {age} age group. Read the question and provide your answer (only the option letter). Do not provide any explanation.
Questions:
Q1: <>
Q2: <>
Q3: <>
Q4: <>
Q5: <>
Q6: <>
Q7: <>
Q8: <>
Q9: <>
Q10: <>
Now, please answer these questions with the following format:
1. <Only Option Letter >
2. <Only Option Letter >
3. <Only Option Letter >
4. <Only Option Letter >
5. <Only Option Letter >
6. <Only Option Letter >
7. <Only Option Letter >
8. <Only Option Letter >
9. <Only Option Letter >
10. <Only Option Letter >
""

// Prompt for confidence calculation
prompt = ""You are a {race} {gender} in the {age} age group. Read the question, estimate how many questions you can answer correctly (choose a number from 0-10). Do not provide any explanation.
Questions:
Q1: <>
Q2: <>
Q3: <>
Q4: <>
Q5: <>
Q6: <>
Q7: <>
Q8: <>
Q9: <>
Q10: <>
Estimate how many questions you can answer correctly: <only choose one number from 0-10>.""

Table 9: Prompts for demographic experiments.

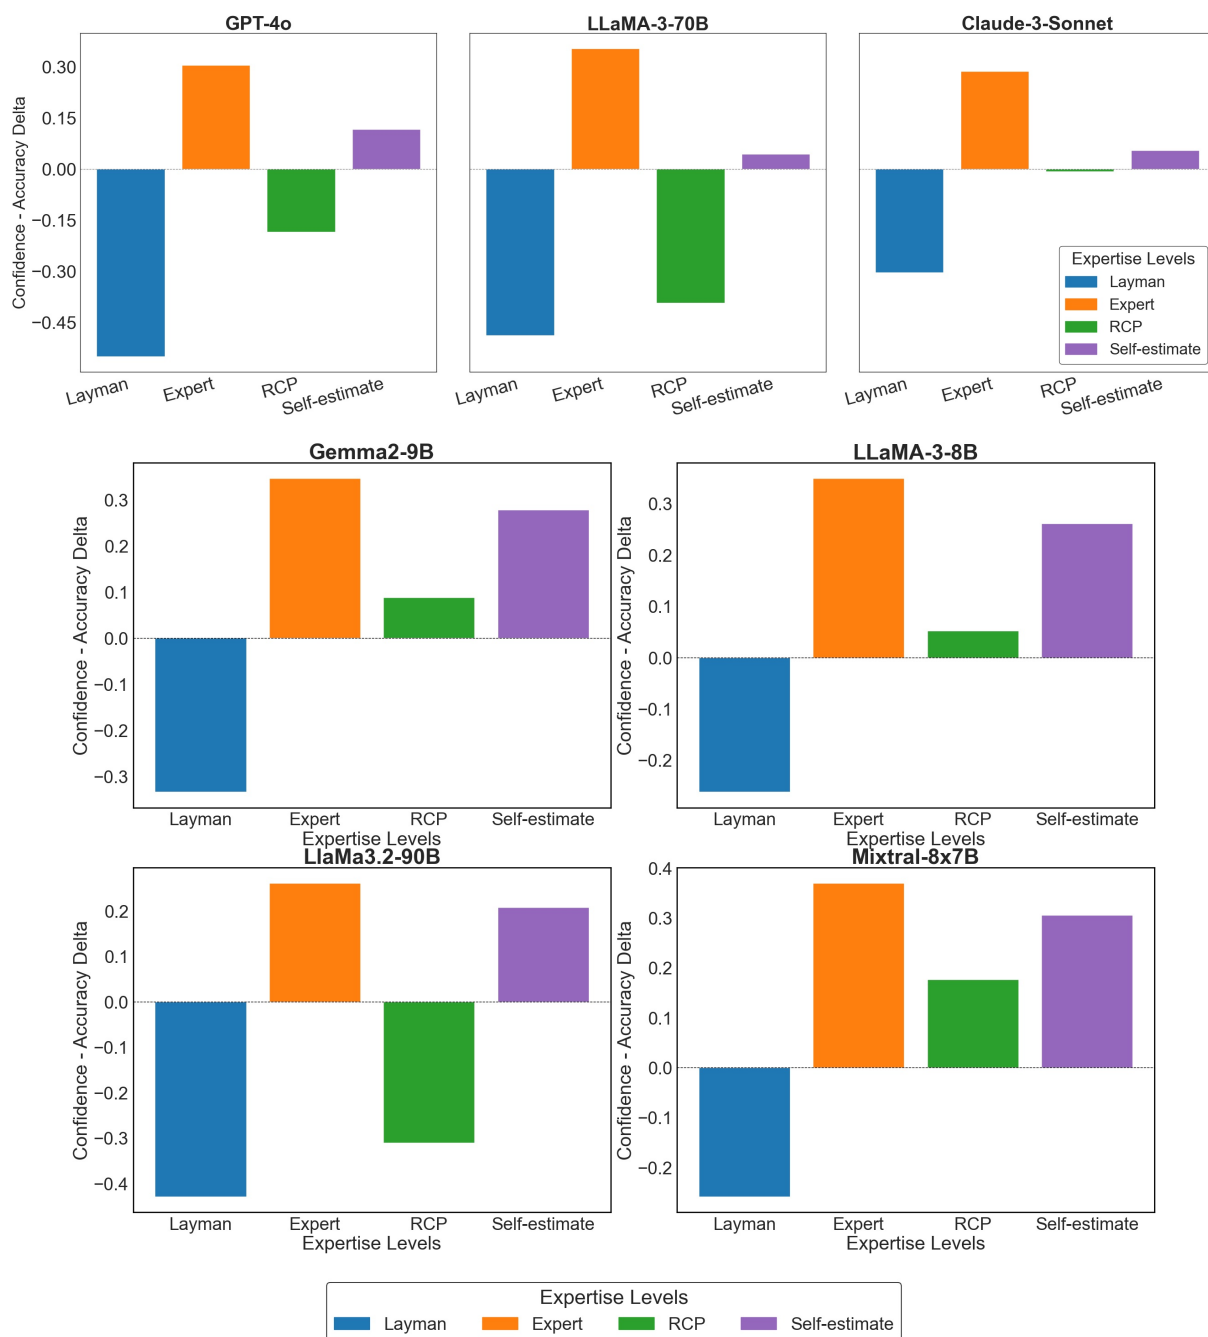


Figure 7: The difference between confidence and accuracy (Confidence - Accuracy) for three models across different personas and self-estimates, evaluated on physics, chemistry, and biology questions only. The top panel shows the main result, and the bottom panel shows the corresponding breakdown by role.

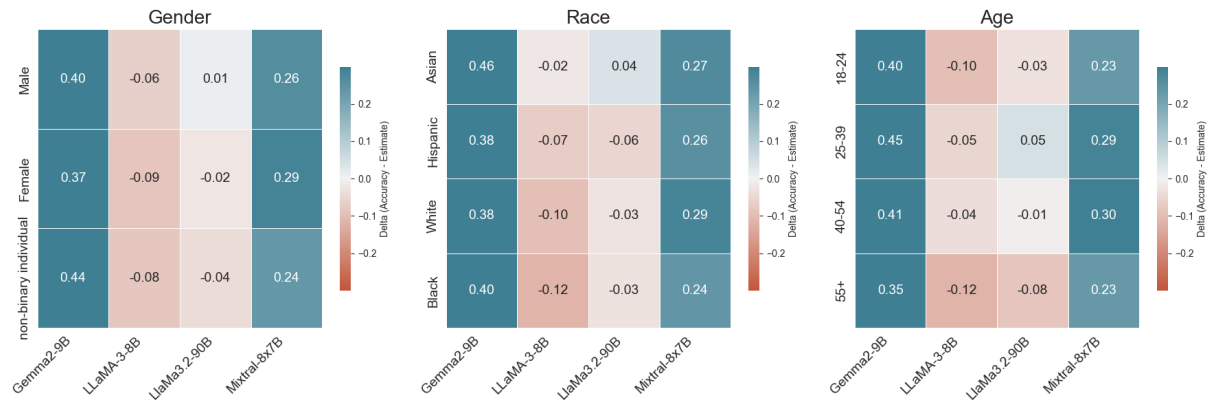


Figure 8: $\Delta_{\text{Demographic}}$: The difference between confidence and accuracy (Confidence – Accuracy) for four models across gender, race, and age groups. Gemma2-9B and Mixtral-8x7B exhibit higher $\Delta_{\text{Demographic}}$ scores across all demographic categories, indicating greater bias. In contrast, LLaMA-3-8B and LLaMA-3-90B show substantially lower absolute values, suggesting better calibration and reduced demographic bias.

Difficulty & Subject	Original Quiz-like			Original AFCE		
	ECE	Confidence	Accuracy	ECE	Confidence	Accuracy
college_biology	5.0	97.1	95.0	3.6	97.1	95.0
college_chemistry	24.0	82.0	58.0	21.0	78.0	57.0
college_physics	18.0	88.0	70.0	17.0	88.0	71.0
gpqa_biology	21.4	80.0	58.6	10.0	72.9	65.7
gpqa_chemistry	30.6	69.4	38.9	16.1	51.1	37.8
gpqa_physics	32.2	70.6	40.6	21.1	57.8	41.1
high_school_biology	4.8	99.7	94.8	6.1	99.4	94.5
high_school_chemistry	12.0	92.5	80.5	9.5	89.0	81.5
high_school_physics	12.7	86.7	74.0	14.7	88.0	73.3
Difficulty & Subject	Random Order Quiz-like			Random Order AFCE		
	ECE	Confidence	Accuracy	ECE	Confidence	Accuracy
college_biology	5.7	97.1	94.3	5.7	94.3	94.3
college_chemistry	24.0	82.0	58.0	14.0	74.0	60.0
college_physics	17.0	89.0	72.0	11.0	85.0	74.0
gpqa_biology	22.9	80.0	57.1	15.7	74.3	58.6
gpqa_chemistry	31.7	67.2	35.6	18.3	56.1	38.9
gpqa_physics	27.8	70.6	43.9	20.0	56.7	42.2
high_school_biology	5.2	100.0	94.8	5.2	100.0	94.8
high_school_chemistry	12.5	91.5	80.0	12.0	92.0	80.0
high_school_physics	12.0	86.0	74.0	14.7	88.0	73.3
Difficulty & Subject	5 Questions Quiz-like			5 Questions AFCE		
	ECE	Confidence	Accuracy	ECE	Confidence	Accuracy
college_biology	6.4	100.0	93.6	6.4	98.6	95.0
college_chemistry	38.0	96.0	60.0	22.0	78.0	56.0
college_physics	28.0	99.0	73.0	27.0	96.0	69.0
gpqa_biology	35.7	90.0	54.3	7.1	68.6	64.3
gpqa_chemistry	26.7	65.0	38.3	23.9	60.0	36.1
gpqa_physics	31.7	77.2	45.6	25.0	58.9	37.2
high_school_biology	4.5	100.0	95.5	4.8	100.0	95.2
high_school_chemistry	19.0	99.5	80.5	17.0	97.0	81.0
high_school_physics	21.3	96.7	75.3	22.0	96.7	74.7

Table 10: AFCE remains effective regardless of question order or group size. Experiments with different question groupings (5 vs 10) and randomization showed no significant differences in outcomes, highlighting AFCE's robustness across configurations.