

# Susu Box or Piggy Bank: Assessing Cultural Commonsense Knowledge between Ghana and the U.S.

Christabel Acquaye   Haozhe An   Rachel Rudinger  
University of Maryland, College Park  
{cacquaye, haozhe, rudinger}@umd.edu

## Abstract

Recent work has highlighted the culturally-contingent nature of commonsense knowledge (Shen et al., 2024). We introduce AMAMMERε (/ɑː.mɑː.muː.reɪ/), from the Akan word meaning ‘culture.’ This test set of 525 multiple-choice questions is designed to evaluate the commonsense knowledge of English LLMs, relative to the cultural contexts of Ghana and the United States. To create AMAMMERε, we select a set of multiple-choice questions (MCQs) from existing commonsense datasets and rewrite them in a multi-stage process involving surveys of Ghanaian and U.S. participants. In three rounds of surveys, participants from both pools are solicited to (1) write correct and incorrect answer choices, (2) rate individual answer choices on a 5-point Likert scale, and (3) select the best answer choice from the newly-constructed MCQ items, in a final validation step. By engaging participants at multiple stages, our procedure ensures that participant perspectives are incorporated both in the *creation* and *validation* of test items, resulting in high levels of agreement within each pool. We evaluate several off-the-shelf English LLMs on AMAMMERε.<sup>1</sup> Uniformly, models prefer answers choices that align with the preferences of U.S. annotators over Ghanaian annotators. Additionally, when test items specify a cultural context (Ghana or the U.S.), models exhibit some ability to adapt, but performance is consistently better in U.S. contexts than Ghanaian. As large resources are devoted to the advancement of English LLMs, our findings underscore the need for culturally adaptable models and evaluations to meet the needs of diverse English-speaking populations around the world.

## 1 Introduction

Many datasets have been developed to train and test language models’ knowledge and reasoning

<sup>1</sup>Data is available at <https://huggingface.co/datasets/Christabel/AMAMMERε>

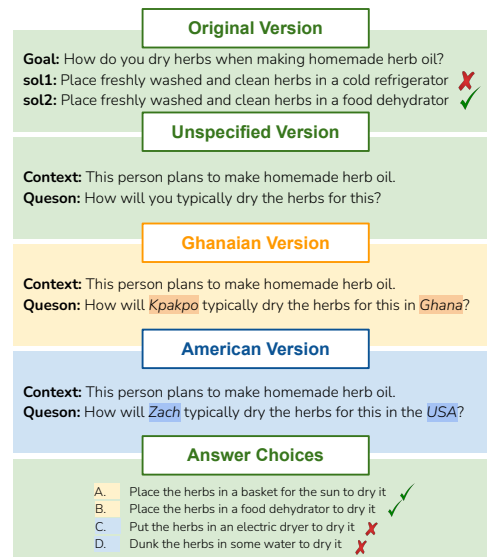


Figure 1: An example from our dataset which asks about drying herbs. The answer choices show two culturally distinct correct responses where the Ghanaian consensus aligns more with sun-drying, while the American consensus leans more towards using a food dehydrator.

abilities; however they struggle to account for cultural differences (Hershcovich et al., 2022), particularly for low-resource cultures—those that lack substantial online data and resources representation in computational research (Duong et al., 2015). We hypothesize that these datasets, created largely by researchers more familiar with Western cultures, contain implicit Western cultural biases. Consequently, models trained on these datasets may underperform for underrepresented cultures, such as Ghanaian culture.

In this paper, we study the implications of cultural differences between the United States and Ghana for developing commonsense benchmark evaluations of English LLMs. Our approach is informed by cultural consensus theory, which posits that the culturally “correct” answer lies in agreement with the shared consensus of the group and

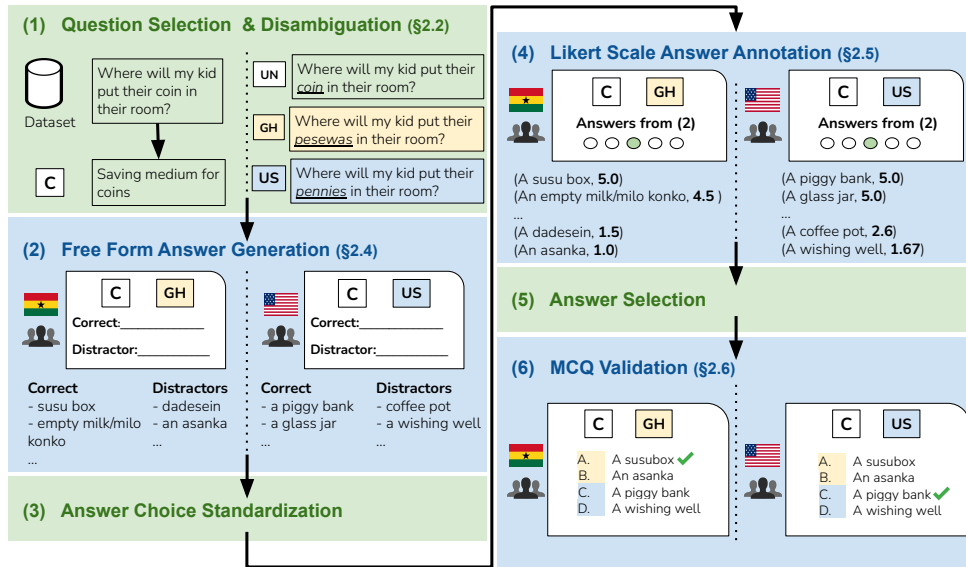


Figure 2: Overall pipeline for our test set generation shown from left to right. It starts with disambiguating sampled questions (§2.2) from our select dataset for unspecified (UN), Ghanaian (GH), and American (US) cultural settings for context (C). Participants provide free-form answers in the Generation Stage (§2.4). Annotators rate these answers in the Likert Scale Answer Annotation task (§2.5). The most and least favored answers for the two cultures is selected and used to construct the MCQs for the human baseline annotations (§2.6).

reflects what is commonly known by the members of that group (Boster, 1985; Romney et al., 1986; Weller, 2007). This can be understood as a culturally-specific refinement of traditional AI accounts of commonsense knowledge as, roughly, “the stuff everyone knows.”

In contrast to works that primarily use human annotators as post-hoc validators on synthetically generated data (Zhang et al., 2017; Rao et al., 2024), this work involves human annotators throughout the entire test set creation process: writing free form answer choices (§2.4), rating the plausibility of answer choices (§2.5), and validating the final multiple-choice test items (§2.6). Though this process is costly, we believe it is well-motivated for our goal of creating a test set to probe comparative cultural knowledge in LLMs, while minimizing (to the extent possible) sources of cultural biases or inaccuracies in the test set itself. Our process ensures that all answers choices representing a culture are (1) written by a member of that culture, and (2) represent a consensus among other members of the culture. We present a dataset of 525 multiple choice questions (MCQ) containing three question versions shown in Figure 1 for a downstream question-answering task.

For our work, we define a culturally adaptable model as one that can recognize implicit and explicit cultural cues and respond in a culturally ap-

propriate manner, e.g., selecting a culturally appropriate answer choice. Models with high degrees of cultural adaptability are desirable as they would be useful to a wider set of users and could better serve a variety of multi- or cross-cultural applications. We conduct experiments to assess whether two encoder models and six LLMs exhibit cultural biases in two ways: (1) by showing a higher preference for answers consistent with one culture over another when presented with culturally unspecified questions, and (2) by showing higher accuracy for some cultural contexts over others when cultural cues are provided.

Our main contributions are as follows:

1. A multi-stage participatory procedure to guide our dataset creation, inspired by theories of cultural consensus.
2. AMAMMER: A test dataset of 525 multiple-choice commonsense reasoning questions that highlight cultural differences between Ghana and the U.S., created through a process involving three separate stages of annotation with volunteer participants from both countries.
3. An examination of the performance of state-of-the-art encoder models and large language models (LLM) in the context of Ghanaian and American cultures.

4. An analysis and discussion of the implications of our experimental results.

## 2 Constructing the AMAMMERe Dataset

We construct the AMAMMERe dataset through a multi-stage process outlined in Figure 2, and described in detail throughout §2. In the first stage we select and disambiguate MCQ questions (§2.2) from existing commonsense datasets (§2.1), covering a range of topics. Next, we collect free-form answers (correct and incorrect) to these questions (§2.4) by surveying separate pools of Ghanaian and US annotators (Appendix H). In a second survey, annotators from each pool rate MCQ answers on a 5-point plausibility Likert scale (§2.5), used to select a “correct” answer and a set of distractors. The newly constructed MCQ is then presented in full to annotators from each pool for final validation (§2.6). Our use of Likert ratings to select correct and incorrect choices in MCQ construction is motivated by our finding (in parallel work) that commonsense MCQs test items exhibit low agreement when the gap in plausibility ratings between the “correct” and distractor answers is small (Palta et al., 2024).

### 2.1 Data Sources

To begin, we select a sample of questions from each of the following publicly-available commonsense benchmarks:

1. Commonsense Question Answering Dataset, CSQA (Talmor et al., 2019)
2. Social Intelligence Question Answering Dataset, SIQA (Sap et al., 2019)
3. Physical Interaction: Question Answering Dataset, PIQA (Bisk et al., 2020)

We select these benchmarks as they cover different aspects of commonsense reasoning and hence may be culturally divergent in different ways. CSQA evaluates general commonsense, while SIQA evaluates social interactions and reasoning about human behavior, relevant for probing cultural norms and social practices, and PIQA tests physical commonsense knowledge which also probes approaches to everyday tasks. We observe that for any of these domains, what is considered a “correct” answer may be culturally dependent.

### 2.2 Question Selection and Disambiguation

**Question Selection** 200 questions were initially selected from these datasets by the first author, who is Ghanaian, based on an expectation that Ghanaian annotators would disagree with the answer labeled as correct. A total of 77, 53, and 70 questions are selected from CSQA, SIQA, and PIQA, respectively. We manually categorized the 200 questions by themes, or cultural facets; the resulting statistics are presented in Table 2.

Each item in CSQA includes a question concept, question, and five answer choices; SIQA items contain a context, question, and three answer choices; PIQA provides a goal and pair of solution choices.<sup>2</sup> To ensure consistency in our final dataset structure and format, we manually created a context for CSQA and PIQA that contained general information relevant to a question (see shared context between the three question versions in Figure 1). For SIQA, which contains a context, we removed culturally specific cues like names and places. We also created an unspecified version of each selected question. In this unspecified version, we removed any obvious cultural markers for the US or Ghana, resulting in a culturally ambiguous question. In Figure 2, we show how we created a general context for a question sampled from CSQA, providing information on saving mediums for coins.

**Disambiguation** To provide Ghanaian and US annotators with questions that accurately reflect their respective cultures, we rewrote questions from existing commonsense reasoning datasets to disambiguate the cultural contexts. We did this by including cultural markers that are US-specific or Ghana-specific. These markers may be explicit (e.g., the phrase “in Ghana”) or implicit (e.g., culturally-associated names of people, places and objects). For example, as detailed in Figure 2, the term *coins* is disambiguated as *pennies* and *pesewas* in the US and Ghana versions, respectively.<sup>3</sup> This cultural disambiguation approach is similar to methods used in the bias benchmarks BBQ (Parrish et al., 2022) and FORK (Palta and Rudinger, 2023). In addition to the insertion of these cultural markers, we also modified pronoun usage to reflect different cultural norms between the U.S. and Ghana, using gender-specific pronouns in the Ghanaian

<sup>2</sup>For the initial formats of questions from the three datasets, see the original version blocks in Figures 13, 14, and 15 of the appendix F.

<sup>3</sup>See specific dataset disambiguated examples in Figures 13, 14, and 15.

questions and gender-neutral terms in the US questions, in response to preliminary survey feedback that highlighted the potential confusion caused by non-traditional pronoun use in Ghana (Hambleton and Patsula, 1998). This process culminated in 200 context-question pairs created for each culture.

### 2.3 Annotation

Our recruitment process was conducted with full approval from an Institutional Review Board (IRB), ensuring that all potential risks of our study were fully addressed. Due to the inaccessibility of major crowdsourcing platforms in Ghana, we advertised our study directly to individuals through social media and to institutions via email and successfully recruited 140 Ghanaian volunteers and 101 American volunteers. More details about participants are in Appendix H.

After our initial set of context-question pairs was generated, we used a three-stage annotation process to create our test set: (1) free form answer generation, (2) likert scale answer annotation, and (3) multiple answer choice annotation. We set up online surveys using Qualtrics and invited participants to help generate diverse and culturally appropriate answers for our study.

### 2.4 Answer Generation

**Free Form Answer Generation** Our two annotator pools (Ghanaians and Americans) were given the context-question pairs to respond to in a free-form QA task, as shown in Figure 2.<sup>4</sup> The survey was divided into 8 sets of 25 questions for Ghanaian and American participants. For each context-question pair, annotators provided three different responses: (1) a culturally correct answer, (2) a distractor answer, that was not too obvious, and would require some thought to discount; and (3) feedback on the formulation of the context-question pair. This feedback with some detailed free-form answers, also provided 25 additional context and questions for a subsequent survey set (Set 9).<sup>5</sup>

**Answer Choice Standardization** We manually standardized the answers to correct grammatical errors and ensure consistency in the answer format as this will be shared by the unspecified and specified questions in our final test set. This standardization

<sup>4</sup>The survey screenshots show the interface in Figure 16 for Ghanaians and Figure 17 for Americans for the free form answer generation.

<sup>5</sup>Participants suggested alternative phrasings and provided extra answers, highlighting ambiguities that led to us creating this new survey set that allowed us to refine the questions for clarity and context.

process is appropriate because our objective is to measure cultural knowledge rather than variations in spelling, typing, or grammar. By lightly editing the text, we preserve the underlying meaning while minimizing potential statistical artifacts that could inadvertently influence a model and undermine our experiments. Our modifications included anonymizing names, incorporating relevant aspects from the questions into all answers for uniform relevance, standardizing the starting structure of responses (e.g. starting with a verb) to remove obvious answer hints, and aligning tenses for uniformity. We add the “a/an” determiners for answers missing these for uniformity, in Figure 2. These light-weight edits help reduce the chance that stylistic artifacts could later be exploited by LLMs during evaluation. If an edit inadvertently changes the meaning of an answer, this can be rectified in the second (§2.5) or third (§2.6) round of annotations.

### 2.5 Answer Plausibility Rating

**Likert Scale Answer Annotation** The collected answers were randomized and presented to participants who rated them on a 5-point Likert scale<sup>6</sup> of subjective likelihood assessments, ranging from “Very Likely” (5) to “Very Unlikely” (1), following Zhang et al. (2017). We compute the average Likert score, which allows us to rank and compare the best and worst answer choices per question. These scores are used to then construct multiple choice questions (§2.6). Figure 2 shows the average ratings with a susu box and a piggy bank rated 5.0 (high agreement). We use Spearman’s rank correlation coefficient (Spearman, 1904) to assess the agreement on the rankings derived from the annotators’ ratings, following the approach used by Kiritchenko and Mohammad (2016). We treat each annotator’s rankings as an independent variable. For every possible pair of annotators who rated the same question, we calculate the Spearman correlation coefficient between their rankings and average these coefficients for each question to get an overall agreement score. Finally, we averaged these per-question scores across the entire dataset to derive a single average coefficient for each group. The analysis yielded average coefficients of 0.72 for the US and 0.67 for Ghana, indicating moderate to substantial agreement among annotators.

**Answer Selection** Using the Likert-rated an-

<sup>6</sup>Figures 18 and 19 show how the answers are presented to the annotators in a survey format so they can provide their ratings with the answer choices in the Likert Annotation Stage.



swer choices, we then select four candidate answer choices: a high-scoring and low-scoring answer choice based on Ghanaian annotator scores, and a high-scoring and low-scoring answer choice based on US annotator scores. By default, we selected the highest- and lowest-scoring answer choices from both pools in a greedy fashion. However, if the highest-scoring answers for both Ghanaian and US annotators were thematically similar, we instead selected the second-highest scoring answer choice, and so on, until contrasting answer choices were obtained. We followed an analogous procedure for selecting the two lowest-scoring answer choices. This process ensured that for every sample, we had a Ghana correct answer, and a Ghana distractor answer as well as a US correct answer and a US distractor answer as shown in Figure 2. For example, a “susu box” and a “piggy bank” were chosen as correct answers with high ratings of 5.0, while an “asanka” and a “wishing well” served as the distractors with the lowest scores.

**Multiple Answer Choice Construction** Our MCQs were created using the most favored and least favored answers from each culture. The Answer Choice block in Figure 1 shows these answers, with Ghanaian answers indicated by orange shade and US answers indicated by the blue shade. “Correct” answers and “Distractors” are defined through consensus among annotators, based on cultural alignment. This consensus-based approach reflects cultural-specific perceptions of correctness.

## 2.6 MCQ Validation

Annotators were presented with a context-question pair and four selected answer choices, two each from the Ghana and US culture. They evaluated the answers by selecting the most appropriate choice for their cultural background and provided feedback, their confidence levels, and reasoning behind their choices.<sup>7</sup>

We assessed the responses for the total 225 questions per culture across 9 survey sets, selecting 175 per culture after excluding 50 samples where the majority-selected answers did not align with the highest-rated answers from earlier Likert score evaluations, which required a 75% agreement threshold (at least three-quarters of all annotators agreed on the same answer choice in our survey.) Each question was evaluated by 4 to 7 annotators, with a larger number from the Ghanaian pool due

<sup>7</sup>The final validation survey interface for Ghanaian and US annotators is shown in Figure 20 and Figure 21 respectively.

to more volunteers. This filtering process helped ensure the reliability of our MCQs. The validation outcomes showed a Krippendorff’s alpha of 0.78 for Ghanaian annotators and 0.86 for US annotators, indicating a higher consensus among US participants potentially due to the familiarity with the US-centric content.

## 2.7 Test set

The test set consists of 525 questions divided into three types, each containing 175 questions:

**Unspecified:** Lacks cultural markers specific to any culture.

**Ghana Specified:** Includes explicit or implicit Ghanaian cultural markers.

**US Specified:** Includes explicit or implicit US cultural markers.

See the Unspecified, Ghanaian, and American-specified versions in Figure 1. Table 5 in the appendix, breaks these 525 samples into different cultural facets.

## 3 Experiment

We experiment with BERT-base (Devlin et al., 2019), RoBERTA-base (Liu et al., 2019) and other generative LLMs across the setups below. We finetuned these encoder models on a consolidated version of the original datasets and run our evaluations. We also evaluated, in zero-shot setting, six state-of-the-art instruction-tuned LLMs: Llama2 (Touvron et al., 2023) with three different model sizes (7b, 13b and 70b), Llama3-70B, Mistral-7B-Instruct-v0.2 (Jiang et al., 2023) and Google’s gemma-7b-it (Team, 2024). We used three different prompt types for these LLMs as detailed in Appendix B. See more experiment details in Appendix E.

### 3.1 Answer-Only Baselines

We run experiments that builds on the hypothesis-only idea from Poliak et al. (2018). This setup ignores the context and question, which contains indicators for the cultural groups and provides only the four answer choices to the models (Balepur et al., 2024). It explores the models’ inherent preference to ascertain if it relies on some artifacts that may be present in the answers. The prompts used in this setup assume an unspecified cultural setting with no specific reference to a group.

### 3.2 Experiment Setups

We hypothesize that models will favor answer choices that reflect US culture when the cultural context in a question is unclear. To study this, we

	Answer-Only				Question-and-Answers				$(\Delta = \text{Question-and-Answers} - \text{Answer-Only})$			
	%GH	%US	%GH Dist.	%US Dist.	%GH	%US	%GH Dist.	%US Dist.	%GH	%US	%GH Dist.	%US Dist.
roberta-base	25.14	35.43	18.29	21.14	30.29	51.43	10.28	8.00	+5.15	+16.00	-8.01	-13.14
Llama2-70B	40.47	39.14	10.39	10.00	29.95	59.48	6.34	4.23	-10.52	+20.34	-4.05	-5.77
Llama3-70B	41.77	41.58	9.95	6.70	23.43	68.57	5.33	2.67	-18.34	+26.99	-4.62	-4.03
Gemma-7B	34.86	33.34	16.57	15.23	27.95	57.09	7.19	7.77	-6.91	+23.75	-9.38	-7.46

Table 1: Selected Models’ Performance Across Cultural Settings Without Specified Context. This table assesses the performance of models in an Answer-Only format and a Question-and-Answers format, focusing on correct and distractor answers for both Ghanaian (“%GH” and “%GH Dist.”) and US (“%US” and “%US Dist.”) culture. The “( $\Delta$ )” column quantifies the change in performance when shifting from Answer-Only to Question-and-Answers format. See full results including other models in Table 6.

present models with a general context and questions that have no implicit or explicit references to either Ghana or the US—the unspecified version described in §2.2. This unspecified setup allows us to evaluate whether models inherently prioritize one cultural groups’ correct answers over the other. We measure the frequency of Ghanaian answers selection over US answers, to assess the models’ predisposition towards U.S. culture.

Furthermore, we hypothesize that when cultural cues are included in a question, models will exhibit some degree of cultural adaptability by choosing the appropriate answer. However, there may still exist a gap between model performance in Ghanaian and US cultural contexts. To investigate this, we reuse the culturally disambiguated questions developed during the annotation stage (see §2.2). Models are presented with questions containing cultural markers referencing either Ghanaian or US cultures. This specified setup assesses the models’ ability to adapt to cultural contexts by prioritizing the most culturally appropriate answer based on the provided cues. We measure a models’ accuracy and examine its ability to understand and prioritize culturally relevant answers based on cultural cues.

### 3.3 Answer Choice Configuration

For each experimental setup, we run a combination of these configurations:

**All-Answers** Each question includes all four answers: two Ghanaian (one correct, one distractor) and two US (one correct, one distractor).

**Correct-Only** Only the correct answers are provided—one from Ghanaian and one from US cultural groups. This setup evaluates whether models can accurately select the appropriate correct answer without the presence of distractors.

**Ghana-Only** Each question presents one correct and one distractor answer exclusively from the Ghanaian cultural group. By this restriction,

we can assess the model’s performance within the Ghanaian culture and identify any potential gaps in cultural knowledge.

**US-Only** Each question has one correct and one distractor answer solely from the US cultural group. This setup examines the model’s ability to answer questions within the US cultural context without the influence of the Ghanaian culture answers.

## 4 Results and Discussion

We evaluate models across our experimental setups to understand their inherent biases and cultural sensitivities. We present the models’ preferences for an unspecified cultural setting in the All-answers configuration in Table 1 and their overall preference in their respective specified cultural settings in Figures 3 and 4. We highlight only selected models’ performances in the main text while detailing full model performances in Table 7 in the Appendix. We analyze results for Correct-only configuration in Figure 7 as well as Ghana-only in Figures 11 and US-only in Figure 12 of the appendix. We qualitatively probe models’ responses to a specific test set question and discuss through the lens of Llama3-70B, its performance for the different cultural facets.

### 4.1 Quantitative Analysis of Preferences

#### Models relatively select US correct and Ghana correct equally in Answer-Only settings

In the Answer-Only section of the unspecified experiments run in the All-answers configuration, we do not record any significant differences in selecting US correct and Ghana correct answers for most of the models as seen in Table 1. The relatively equal rates of selecting US correct and Ghana correct answers in the Answer-Only setup demonstrates that there is nothing intrinsic about the Ghanaian correct answers that the model dislikes compared to the US correct answers. This suggests that the




Figure 3: Selected Model Preference Distribution in GH Specified Settings. \* represents values less than 5. See full results including other additional models in Table 7.




Figure 4: Selected Model Preference Distribution in US Specified Settings. \* represents values less than 5.1. See full results including other additional models in Table 7.

disparities we observe under other conditions are due to substantive cultural factors, not procedural or stylistic artifacts.

**Providing a question helps models pick up relevant cues to select correct answers** Compared to the Answer-Only setup, when a question with no cultural reference to either Ghana or the US is provided in the All-answers configuration, it reduces the chances of the models selecting distractor answers, indicated by the negative differences recorded in the ( $\Delta$ ) column of Table 1. This suggests that providing a question helps models pick up relevant cues to select correct answers, thereby reducing their preference for distractors. The reduction in distractor preference could indicate that the test set is constructed with minimal confounding artifacts, so that the inclusion of questions not only helps models focus but also possibly enhances their ability to utilize context effectively.

**Models show some cultural adaptability** One way we try to measure the “cultural adaptability” of a model is by measuring how a model’s selected answer choice changes (or does not change) when the cultural context of the question changes, but the answer choices remain the same. Figures 3 and 4 show that for the same question-answer pairings,




Figure 5: Accuracy of models in Ghanaian settings when conditioned only on Ghanaian correct and distractor answers. See full results including other additional models in Figure 11.




Figure 6: Accuracy of models in US settings when conditioned only on US correct and distractor answers. \* represents values less than 5. See full results including other additional models in Figure 12.

simply modifying the cultural context from US to Ghana increases the probability that the model will choose the Ghanaian correct answer (e.g. a shift from 18% to 60% for Llama3-70B). We observe that the models can to an extent recognize and respond to cultural cues, albeit with varying levels of accuracy across different settings. Llama3-70B scores higher in the US Specified setting at 77% (Figure 4) compared to 60% in the Ghana Specified setting (Figure 3). This performance, with most models achieving over 50% accuracy when cultural cues are provided, highlights their potential for cultural adaptability. Our findings highlight that models like Llama3-70B when forced to select between the culturally specific correct and distractor answers, it realizes a high performance—87.72% in Ghana-only (Figure 5) and 95.62% in US-only configurations (Figure 6). This suggests that most of these models can effectively distinguish between culture-specific correct answers and distractors, indicating some degree of cultural adaptability.

**Higher preference for US correct answers in unspecified settings** We observe positive and higher differences for %US in the ( $\Delta$ ) column of Table 1 compared to the %GH which indicates a higher preference for US correct answers for all models. RoBERTa-base consistently prefers more

US correct answers both in the Answer-Only and the Question-and-Answers. Llama3-70B shows the highest preference for US correct answers in this unspecified cultural setting. In the unspecified cultural setting in the Correct-only configuration, models prefer more US Correct answers as shown in Figure 7. Table 8 in the appendix, shows a similar trend for other models aside Llama2-7B that has a 51% preference for Ghana correct answers. This suggest that the models are more attuned to US culture, likely stemming from the predominance of US based data used in the pre-training of these models.

### Performance discrepancy favoring US culture

For specified settings, in Figure 7, while most models show some ability to adapt, preferring culturally correct answers over 50% of the time, there is a disparity in their performance based on culture. Encoder models show a preference for Ghanaian correct answers less than 50% of the time, compared to a stronger preference for US correct answers. Llama3-70B, the best-performing model in the Ghana specified setting of the Correct-only configuration, favors Ghana correct answers 67% of the time, whereas in the US specified setting, its preference for US correct answers jumps to 82%. This highlights not only a general US bias in unspecified contexts but also a significant performance variability depending on the specified cultural context.

## 4.2 Qualitative Analysis

**Context:** This person is married with two little kids.  
**Question:** How can Kojo make Bronya feel more magical for his family?  
**Options:**  
 A. Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights  
 B. Take them to the North Pole  
 C. Ensure there is enough food, drinks and fun games  
 D. Leave a couple of bite marks in the cookies that his children leave for Santa  
**Correct Answer:** Ensure there is enough food, drinks and fun games

We explore how Christmas or ‘Bronya’ (in Ghana), is celebrated differently for our two cultures, despite their shared religious origins. While in the US, Christmas often involves decorations and presents, in Ghana, the emphasis on food, drinks and games captures the relevant aspects of the Ghanaian culture, which stresses the communal and festive spirit essential to Ghanaian celebrations. Although there are some decorations in Ghana like Christmas trees, these are usually found in institutions rather than individual homes (Domingo, 2022). This question tests models’ knowledge of this relative cultural difference. We use implicit markers like “Kojo” (a common name in Ghana)

Cultural Facet	# Samples	US Accuracy (%)	GH Accuracy (%)
Social Customs and Lifestyle	118	70	52
Architecture/Housing	19	84	42
Food and Drinks	16	88	68
Gender Roles, Marriage, & Family	13	77	69
Geography	9	78	78
	175	74%	55%

Table 2: Accuracy across different cultural facets for US-specified and Ghana (GH)-specified settings. ‘# Samples’ shows the number of questions per facet for a given group.

and “Bronya” (Christmas) to subtly indicate the Ghana cultural setting.

From the models responses in Appendix D, five out of seven models select the US correct answer, suggesting that these models may lack the interpretation of the correct cultural context from the implicit markers to select the relevant response for Ghana. Conversely, Llama2-7B and Gemma-7B selects the correct answer which aligns better with Christmas in Ghana. Albeit, a closer look at their explanations suggests a shallow grasp of the specific cultural practices in Ghana, which may indicate a limited cultural understanding. These variations in models responses underscores the need to diversify cultural data to include more implicit cues in training to serve global audiences.

### 4.3 Analysis of Llama3-70B across facets

We breakdown the performance of our highest performing model, Llama3-70B, for Prompt 1,<sup>8</sup> by cultural facets as shown in the Table 2. We focus our analysis on Llama3-70B because it is one of the most recent and advanced models at the time of this writing, demonstrating a lower preference for distractors. For this, we record a 74.29% preference for the US Correct answers in the US Specified Setting and a 55.42% for the Ghana correct answers in the Ghana Specified Setting. Our analysis shows that Llama3-70B is most familiar with the ‘Food and Drinks’ cultural facet in the US contexts, as it answers 14 out of 16 questions correctly, and ‘Geography’ in the Ghana context, with 7 out of 9 correct answers. For ‘Social Customs and Lifestyle,’ the model answered 83 out of 118 questions correctly (70.34%), compared to 62 out of 118 for Ghana (52.54%), indicating better familiarity with US social customs.

In ‘Architecture/Housing,’ it answered 16 out of 19 correctly for the US (84.21%) compared to 8 out of 19 for Ghana (42.11%), showing a stronger grasp for US culture. ‘Food and Drinks’ being the

<sup>8</sup>Prompt 1 is described in appendix B






Figure 7: Preference distribution for different selected models when conditioned only on correct answers from Ghana and the US. See full results including other models in Table 8.

most familiar facet in the US contexts may be due to enough representation of this facet in the training data from different datasets, while the strong familiarity in Ghana ‘Geography,’ might stem from the stable and consistently documented nature of geographical information, which is less culturally variable and easier for the model to learn. Also, the model struggled the most with ‘Social Customs and Lifestyle’ in both contexts, likely due to the intricacies and nuances of social customs, which vary widely. These speculations highlight the importance of balanced training datasets to improve model performance across different cultural facets.

## 5 Related Work

**Data Collection and Annotation** Researchers often rely on collaborative efforts from diverse cultures at different stages in their work. Huang et al. (2019) and Clark et al. (2020) allow diverse annotators to create commonsense questions, ensuring real-world relevance and promoting natural question generation. Ponti et al. (2020) strategically recruits translators to make cultural and linguistic adjustments, ensuring high-quality translations. Our work builds on this to allow broader participation, benefiting from diverse multi-cultural perspectives. While participants do not explicitly create the initial questions, they have the freedom to suggest refinements to questions and contexts, ensuring cultural accuracy and relevance. Zellers et al. (2018) uses adversarial methods to generate multiple-choice answers validated by crowd workers. Yin et al. (2021) employs human annotators for questions and correct answers, with automated distractors generation. Our approach involves annotators creating, selecting, and validating answers.

**Cultural Sensitivity and Bias Detection in Commonsense Reasoning** Palta and Rudinger (2023),

Zhou et al. (2023), Huang and Yang (2023), Bauer et al. (2023) and Naous et al. (2024) evaluate reasoning through downstream QA tasks, inference, entity recognition and sentiment analysis tasks, and discuss Western cultural biases. Shi et al. (2024) fine-tune LLMs on a dataset of self-narratives from TikTok, and measure cultural awareness through automatic entailment scores and human evaluations. Li et al. (2024) prompts models with culturally relevant sentences and analyses the text for cultural symbols and biases, focusing on diversity and implicit Western biases. Rao et al. (2024) assesses LLM responses to social acceptability questions across cultures, revealing difficulties with non-Western norms. Our study evaluates cultural adaptability with our test set that contains both explicit and implicit cultural references. We assess how well models can interpret and respond to culturally nuanced questions in a MCQ task.

## 6 Conclusion

We create a new test set to investigate cultural biases in models, focusing on Ghana and the US. The focus on the Ghanaian population, a harder-to-study yet important group, underlines the importance of representing diverse cultures in computational research. Our study offers multi-stage annotator participation throughout the dataset creation, ensuring we capture nuanced cultural specificities. Our experiments show that models favor US culture in unspecified cultural contexts and specifying cultural cues improves accuracy for both groups, albeit, the improvement is more pronounced for the US. These findings highlight the need for ongoing efforts to make models more culturally inclusive, particularly for low-resource languages. Mitigation techniques should involve annotators from diverse backgrounds at all stages of dataset creation, and training models tailored to specific cultures.

## Limitations

**Limited dimensions of culture** Culture as a social concept, spans a wide range of practices, beliefs, and expressions. Our research focuses on specific cultural facets for only two select cultural groups, which represents just a small aspect of culture. This focus was necessary due to the scope of this study and the resource constraints our study presents. Future research could expand on this by exploring more cultural dimensions for more low-resource cultures, that will offer a more comprehensive understanding of diverse human cultures in language models.

**Cultural stereotype** There is a risk that some survey questions inadvertently perpetuate cultural stereotypes, particularly towards some ethnic groups, regions, or states. For instance, questions that reference a specific aspect of an ethnic group's food culture, or social custom, that is well-known among other people in the country could create an implicit marker that primes participants to select a correct cultural context answer but potentially reinforce stereotypes about the ethnic group.

**Incomplete representation of ethnic/region demographics** The survey might overemphasize the opinions of the majority ethnic groups. This could occur if minority viewpoints are less recognized, leading to the majority's perspective being perceived as the "correct" answer. As such, our best-voted answers might represent the viewpoints of the majority group while other good but least-voted answers although relatable to other groups are in the minority.

**Limited dataset size and skewed demographics** The size of our dataset is limited because we had to rely on volunteer annotators, as we could not use standard crowdsourcing tools for the Ghanaian population. This population is harder to study but important to understand due to its unique cultural context. Additionally, the participants we recruited skewed towards younger subjects, which may not fully represent the broader demographic distribution of the populations studied. This age skew could influence the cultural insights and relevance of the responses collected. Future research should aim to include a more diverse age range and employ strategies to increase participation from underrepresented groups.

## Ethics Statement

Our research investigates cultural biases in commonsense reasoning models by exploring specific cultural facets. We recognize that exploring cultural nuances presents inherent risks, particularly in inadvertently simplifying or misrepresenting cultures. Our work which focused on two select cultures, inherently limits the breadth of cultural diversity and understanding. We involved annotators from diverse backgrounds from these two groups to validate our findings, but we acknowledge that this does not fully encompass the rich diversity within each culture. We also emphasize that while our findings may highlight the presence of cultural biases in certain models, it is crucial to approach the interpretation of our results with an understanding of their contextual limitations and the complexities of cultural representation in AI. Finally, we utilized AI assistants in our experiments (e.g., paraphrasing prompts, etc.).

## Acknowledgements

We thank the anonymous reviewers for their constructive feedback. We also thank Kwesi Cobbina for their invaluable insights, as well as Nelson Padua-Perez, Julius Amegadzie, Belinda Laryea, Deborah Laryea, Amanda Acquaye and Gatha Adhikari for their support in recruiting annotators. We extend our heartfelt gratitude to the volunteers from Academic City University College, the University of Maryland, as well as all other volunteers from other institutions who contributed to this work. Their support was instrumental in the successful completion of our project. This work was supported by NSF CAREER Award No. 2339746. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

## References

- Nishant Balepur, Abhilasha Ravichander, and Rachel Rudinger. 2024. [Artifacts or abduction: How do LLMs answer multiple-choice questions without the question?](#) In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10308–10330, Bangkok, Thailand. Association for Computational Linguistics.
- Lisa Bauer, Hanna Tischer, and Mohit Bansal. 2023. [Social commonsense for explanation and cultural bias](#)

- discovery. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 3745–3760, Dubrovnik, Croatia. Association for Computational Linguistics.
- Yonatan Bisk, Rowan Zellers, Ronan Le bras, Jianfeng Gao, and Yejin Choi. 2020. **Piqa: Reasoning about physical commonsense in natural language**. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(05):7432–7439.
- James S. Boster. 1985. *‘Requiem for the omniscient informant’: there’s life in the old girl yet*. University of Illinois Press. Book Title: directions in cognitive anthropology Pages: 177-197.
- Jonathan H. Clark, Eunsol Choi, Michael Collins, Dan Garrette, Tom Kwiatkowski, Vitaly Nikolaev, and Jennimaria Palomaki. 2020. **TyDi QA: A benchmark for information-seeking question answering in typologically diverse languages**. *Transactions of the Association for Computational Linguistics*, 8:454–470.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: Pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zaytoen Domingo. 2022. How christmas is celebrated in ghana. <https://www.gviusa.com/blog/how-christmas-is-celebrated-in-ghana/>. Accessed: 2024-06-06.
- Long Duong, Trevor Cohn, Steven Bird, and Paul Cook. 2015. **Low resource dependency parsing: Cross-lingual parameter sharing in a neural network parser**. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 845–850, Beijing, China. Association for Computational Linguistics.
- R. Hambleton and L. Patsula. 1998. **Adapting tests for use in multiple languages and cultures**. *Social Indicators Research*, 45:153–171.
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, Constanza Fierro, Katerina Margatina, Phillip Rust, and Anders Søgaard. 2022. **Challenges and strategies in cross-cultural NLP**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013, Dublin, Ireland. Association for Computational Linguistics.
- Jing Huang and Diyi Yang. 2023. **Culturally aware natural language inference**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 7591–7609, Singapore. Association for Computational Linguistics.
- Lifu Huang, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2019. **Cosmos QA: Machine reading comprehension with contextual commonsense reasoning**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2391–2401, Hong Kong, China. Association for Computational Linguistics.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth e Lacroix, and William El Sayed. 2023. **Mistral 7b**.
- Svetlana Kiritchenko and Saif M. Mohammad. 2016. **Capturing reliable fine-grained sentiment associations by crowdsourcing and best-worst scaling**. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 811–817, San Diego, California. Association for Computational Linguistics.
- Huihan Li, Liwei Jiang, Jena D. Huang, Hyunwoo Kim, Sebastin Santy, Taylor Sorensen, Bill Yuchen Lin, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024. **CULTURE-GEN: Revealing global cultural perception in language models through natural language prompting**. In *First Conference on Language Modeling*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. **Roberta: A robustly optimized bert pretraining approach**.
- Tarek Naous, Michael Ryan, Alan Ritter, and Wei Xu. 2024. **Having beer after prayer? measuring cultural bias in large language models**. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16366–16393, Bangkok, Thailand. Association for Computational Linguistics.
- Shramay Palta, Nishant Balepur, Peter Rankel, Sarah Wiegreffe, Marine Carpuat, and Rachel Rudinger. 2024. **Plausibly problematic questions in multiple-choice benchmarks for commonsense reasoning**. In *Findings of the 2024 Conference on Empirical Methods in Natural Language Processing*, Miami. Association for Computational Linguistics.
- Shramay Palta and Rachel Rudinger. 2023. **FORK: A bite-sized test set for probing culinary cultural biases**



- in commonsense reasoning models. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9952–9962. Association for Computational Linguistics.
- Alicia Parrish, Angelica Chen, Nikita Nangia, Vishakh Padmakumar, Jason Phang, Jana Thompson, Phu Mon Htut, and Samuel Bowman. 2022. **BBQ: A hand-built bias benchmark for question answering**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2086–2105. Association for Computational Linguistics.
- Adam Poliak, Jason Naradowsky, Aparajita Haldar, Rachel Rudinger, and Benjamin Van Durme. 2018. **Hypothesis only baselines in natural language inference**. In *Proceedings of the Seventh Joint Conference on Lexical and Computational Semantics*, pages 180–191, New Orleans, Louisiana. Association for Computational Linguistics.
- Edoardo Maria Ponti, Goran Glavaš, Olga Majewska, Qianchu Liu, Ivan Vulić, and Anna Korhonen. 2020. **XCOPA: A multilingual dataset for causal commonsense reasoning**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2362–2376. Association for Computational Linguistics.
- Abhinav Rao, Akhila Yerukola, Vishwa Shah, Katharina Reinecke, and Maarten Sap. 2024. **Normad: A benchmark for measuring the cultural adaptability of large language models**.
- A Kimball Romney, Susan C Weller, and William H Batchelder. 1986. Culture as consensus: A theory of culture and informant accuracy. *American Anthropologist*, pages 313–338.
- Maarten Sap, Hannah Rashkin, Derek Chen, Ronan Le Bras, and Yejin Choi. 2019. **Social IQa: Commonsense reasoning about social interactions**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4462–4472. Association for Computational Linguistics.
- Siqi Shen, Lajanugen Logeswaran, Moontae Lee, Honglak Lee, Soujanya Poria, and Rada Mihalcea. 2024. **Understanding the capabilities and limitations of large language models for cultural commonsense**. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 5668–5680, Mexico City, Mexico. Association for Computational Linguistics.
- Weiyang Shi, Ryan Li, Yutong Zhang, Caleb Ziems, Chunhua Yu, Raya Horesh, Rogério Abreu de Paula, and Diyi Yang. 2024. **Culturebank: An online community-driven knowledge base towards culturally aware language technologies**.
- C. Spearman. 1904. **The proof and measurement of association between two things**. *The American Journal of Psychology*, 15(1):72–101.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. **CommonsenseQA: A question answering challenge targeting commonsense knowledge**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158, Minneapolis, Minnesota. Association for Computational Linguistics.
- Gemma Team. 2024. Gemma: Open models based on gemini research and technology. *arXiv preprint arXiv:2403.08295v1*.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shrutu Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023. **Llama 2: Open foundation and fine-tuned chat models**.
- Susan C. Weller. 2007. **Cultural consensus theory: Applications and frequently asked questions**. 19(4):339–368. Publisher: SAGE Publications Inc.
- Da Yin, Liunian Harold Li, Ziniu Hu, Nanyun Peng, and Kai-Wei Chang. 2021. **Broaden the vision: Geodiverse visual commonsense reasoning**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 2115–2129, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Rowan Zellers, Yonatan Bisk, Roy Schwartz, and Yejin Choi. 2018. **SWAG: A large-scale adversarial dataset for grounded commonsense inference**. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 93–104, Brussels, Belgium. Association for Computational Linguistics.
- Sheng Zhang, Rachel Rudinger, Kevin Duh, and Benjamin Van Durme. 2017. **Ordinal common-sense inference**. *Transactions of the Association for Computational Linguistics*, 5:379–395.



Li Zhou, Antonia Karamolegkou, Wenyu Chen, and Daniel Hershcovich. 2023. [Cultural compass: Predicting transfer learning success in offensive language detection with cultural features](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 12684–12702. Association for Computational Linguistics.

## A Dataset Statistics

Dataset	No. of samples
CSQA	77
SIQA	53
PIQA	70

Table 3: Information showing the number of samples used from each dataset

Disambiguation Type	No. of samples
Explicit	200
Implicit	150
Unspecified	175

Table 4: Information showing the number of samples for each disambiguation type. Unspecified types do not have any cultural reference(either overt or subtle).

Cultural Facet	No. of samples
Social Customs and Lifestyle	354
Architecture/Housing	57
Food and Drinks	48
Gender, Marriage and Family	39
Geography	27
Total	525

Table 5: Test set breakdown for different cultural facets.

## B Prompts

Three prompts were used, increasing in specificity and detail from Prompt 1 to Prompt 3, referencing specific contexts and settings.

### B.1 Prompt 1

"You will be tasked with answering Multiple Choice Questions (MCQs). For each question, identify the answer that best fits the given context. Despite multiple options potentially being correct, select the one that is most appropriate for the specific scenario presented. Your response should be formatted in JSON, indicating the 'Best Answer' along with its corresponding option number and text."

Figure 8: General MCQ task with no specific context provided.

### B.2 Prompt 2

"You will be tasked with answering Multiple Choice Questions (MCQs). For each question, identify the answer that best fits the given context for *alan [specific setting]* setting. Your response should be formatted in JSON, indicating the 'Best Answer' along with its corresponding option number and text."

Figure 9: MCQ task with a specific setting provided.

### B.3 Prompt 3

"You will be tasked with answering Multiple Choice Questions (MCQs). For each question, answer as *alan [specific culture]* the answer that best fits the given context for *alan [specific setting]*. Your response should be formatted in JSON, indicating the 'Best Answer' along with its corresponding option number and text."

Figure 10: MCQ task with both a specific cultural perspective and setting provided.

## C Additional Results

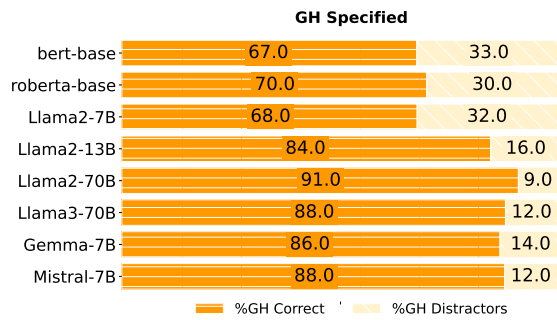


Figure 11: Accuracy of models in Ghanaian settings when conditioned only on Ghanaian correct answers and Ghanaian distractors.




Figure 12: Accuracy of models in US settings when conditioned only on US correct answers and US distractors.

## D Qualitative Results

This section shows the different model responses for one GH-specified test question when all four answer choices are provided.

	Answer-Only				Question-and-Answers				$(\Delta = \text{Question-and-Answers} - \text{Answer-Only})$			
	%GH	%US	%GH Dist.	%US Dist.	%GH	%US	%GH Dist.	%US Dist.	%GH	%US	%GH Dist.	%US Dist.
bert-base	31.43	32.57	17.14	18.86	32.00	45.14	15.43	7.43	+0.57	+12.57	-1.71	-11.43
roberta-base	25.14	35.43	18.29	21.14	30.29	51.43	10.28	8.00	+5.15	+16.00	-8.01	-13.14
Llama2-7B	29.86	24.91	24.91	20.34	36.77	44.38	10.65	8.20	+6.91	+19.47	-14.26	-12.14
Llama2-13B	32.62	30.90	21.19	15.29	30.05	51.96	10.05	7.94	-2.57	+21.06	-11.14	-7.35
Llama2-70B	40.47	39.14	10.39	10.00	29.95	59.48	6.34	4.23	-10.52	+20.34	-4.05	-5.77
Llama3-70B	41.77	41.58	9.95	6.70	23.43	68.57	5.33	2.67	-18.34	+26.99	-4.62	-4.03
Gemma-7B	34.86	33.34	16.57	15.23	27.95	57.09	7.19	7.77	-6.91	+23.75	-9.38	-7.46
Mistral-7B	38.08	35.99	11.16	14.77	30.66	60.19	4.00	5.14	-7.42	+24.20	-7.16	-9.63

Table 6: Models Preference Across Unspecified Cultural Settings. The performance of the models in the Answer-Only evaluates a model preference without any cultural context or questions provided. Columns "%GH" and "%US" show the percentage of correct Ghanaian and US answers, "%GH Dist." the percentage of Ghanaian distractors, and "%US Dist." the percentage of US distractors. The Question-and-Answers columns show similar metrics with unspecified questions provided. The  $(\Delta = \text{Question-and-Answers} - \text{Answer-Only})$  column shows the actual differences between the corresponding values in the Answer-Only and Question-and-Answers columns.

	GH Specified				US Specified			
	%GH	%US	%GH Dist.	%US Dist.	%GH	%US	%GH Dist.	%US Dist.
bert-base	34.0	37.0	18.0	11.0	28.0	44.0	15.0	13.0
roberta-base	33.0	46.0	14.0	7.0	31.0	52.0	10.0	7.0
Llama2-7B	38.0	37.0	15.0	10.0	26.0	53.0	10.0	11.0
Llama2-13B	42.25	37.25	9.25	11.25	29.0	54.0	7.5	9.5
Llama2-70B	55.0	35.0	6.0	4.0	23.0	67.0	4.0	6.0
Llama3-70B	60.0	33.0	4.1	2.9	18.0	77.0	3.0	2.0
Gemma-7B	50.0	34.25	8.75	7.0	24.25	62.0	5.0	8.75
Mistral-7B	48.0	42.0	5.0	5.0	25.0	65.0	5.0	5.0

Table 7: Full Results on Model Preference Distribution in GH and US Specified Settings. Columns "%GH" and "%US" show the percentage of correct Ghanaian and US answers, "%GH Dist." the percentage of Ghanaian distractors, and "%US Dist." the percentage of US distractors.

	UN Specified		GH Specified		US Specified	
	%GH	%US	%GH	%US	%GH	%US
bert-base	41.0	59.0	47.0	53.0	41.0	59.0
roberta-base	37.0	63.0	42.0	58.0	39.0	61.0
Llama2-7B	51.0	49.0	49.0	51.0	37.0	63.0
Llama2-13B	35.0	65.0	58.0	42.0	31.0	69.0
Llama2-70B	37.0	63.0	66.0	34.0	25.0	75.0
Llama3-70B	26.0	74.0	67.0	33.0	18.0	82.0
Gemma-7B	40.0	60.0	59.0	41.0	30.0	70.0
Mistral-7B	34.0	66.0	56.0	44.0	25.0	75.0

Table 8: Full Results on Model Preference Distribution in UN, GH, and US Specified Settings for the Correct-only configuration. Columns "%GH" and "%US" show the percentage of correct Ghanaian and US answers.

### Model Responses:

#### bert-base:

Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights

#### roberta-base:

Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights

#### Llama2-7B:

"Best Answer": "(3) Ensure there is enough food, drinks, and fun games for the whole family."  
 "Explanation": "The context of the question is that Kojo wants to make Bronya feel more magical for his family, and the options provided are all related to Christmas celebrations. Option (3) is the best answer because it suggests ensuring that there is enough food, drinks, and fun games for the whole family, which is a great way to create a magical and enjoyable experience for everyone."

#### Llama2-13B:

"Best Answer": "(1) Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights",  
 "Explanation": "To make Bronya feel more magical for his family, Kojo should focus on creating a festive and joyful atmosphere around the Christmas tree. Decorating the tree with lots of presents underneath it, and the house and front porch with ornaments and lights, will help create a sense of wonder and excitement for the children. This option aligns best with the context of the question, which is focused on creating a magical atmosphere for the family during the holiday season."

#### Llama3-70B:

"Best Answer": "(1) Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights",  
 "Explanation": "This answer is the most suitable because it creates a festive and magical atmosphere for Kojo's family, especially his children, a key aspect of the Christmas season."

#### Gemma-7B:

"Best Answer": "(3) Ensure there is enough food, drinks and fun games",  
 "Explanation": "The text describes a person who is married with two little kids and wants to make Bronya feel more magical for their family."

### Mistral-7B:

"Best Answer": "1. Decorate the Christmas tree with lots of presents underneath it, and decorate the house and the front porch with ornaments and lights."  
"Explanation": "The context mentions that Kojo has a family with two little kids. The question asks how Kojo can make Bronya feel more magical for his family. Option 1 is the best answer as it aligns with the context and the holiday season. Decorating the Christmas tree."

- Gemma: <https://github.com/google-deepmind/gemma/blob/main/LICENSE>

## E Additional Experimental Setup Details

**Encoder Models** We finetuned the encoder models using a consolidated version of the original datasets,(CSQA, SIQA, and PIQA) for 3 epochs and a learning rate of 1e-5 for the base models. The consolidated dataset for the models was created by standardizing the number of answer choices across the different datasets to four answer choices per question. We evaluated these finetuned encoder models on the validation sets of these original datasets to ensure it is consistent with existing baselines. Performance-wise, on the consolidated dataset, BERT-base achieved scores of 60.77% on CSQA, 60.29% on SIQA, and 65.83% on PIQA. Roberta-base performed better with scores of 69.37% on CSQA, 70.37% on SIQA, and 69.37% on PIQA.

**LLMs** For the Llama2 models, Gemma and Mistral, we run the experiments with a random seed of 42. We used a temperature of 0.1 for all LLMs.

### Summary of Experiments

Exp. Setting	Cultural Setting	Exp. Configuration	Answers Provided
Unspecified	No culture No culture	All-answers Correct-only	GH Correct, GH Dist., US Correct, US Dist. GH Correct, US Correct
Specified	GH GH GH	All-answers Correct-only GH-only	GH Correct, GH Dist., US Correct, US Dist. GH Correct, US Correct GH Correct, GH Dist.
Specified	US US US	All-answers Correct-only US-only	GH Correct, GH Dist., US Correct, US Dist. GH Correct, US Correct US Correct, US Dist.
Answer-Only	No culture	All-answers	GH Correct, GH Dist., US Correct, US Dist.

Table 9: This table summarizes all experiments we conduct. GH = Ghana, Dist. = Distractors.

**Terms of use for each model** We carefully follow the guidelines per the terms of usage described by the model authors or company

- Llama2: <https://ai.meta.com/llama/license/>
- Llama3: <https://llama.meta.com/llama3/license/>
- Mistral-Instruct-v0.1: <https://mistral.ai/terms-of-service/>

**Licenses** The CSQA, SIQA and PIQA are used under the MIT<sup>9</sup> and CC-BY<sup>10</sup> licenses.

## F Disambiguation for the three datasets

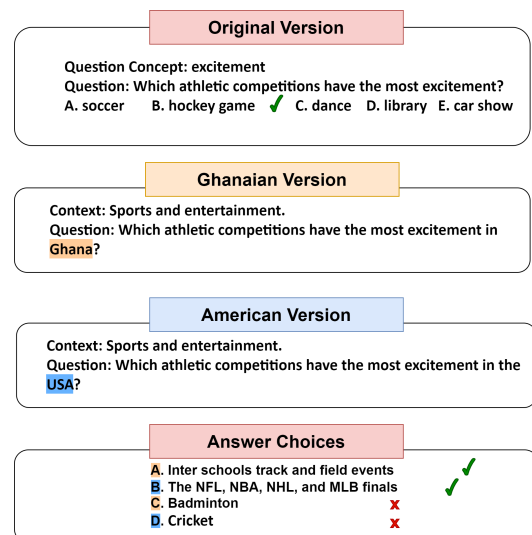


Figure 13: An example of the question-answer format in the original CSQA (Talmor et al., 2019) dataset and the disambiguated version from our dataset, with Ghanaian and American versions as well as the correct and distractor answers for both versions.

**Original Version**

Context: Quinn was unhappy in the marriage and promptly filed for a divorce.  
 Question: What will Quinn want to do next?  
 A. separate from his wife ✓ B. throw a party C. get himself a lawyer

**Ghanaian Version**

Context: Someone was unhappy in their marriage and had promptly filed for divorce.  
 Question: If this unhappy person was Adepa, what might she want to do next in Ghana?

**American Version**

Context: Someone was unhappy in their marriage and had promptly filed for divorce.  
 Question: If this unhappy person was Quinn, what might she want to do next in the USA?

**Answer Choices**

A. talk to their parents or religious leaders ✓  
 B. speak to her partner, if things can't be fixed then divorce ✓  
 C. take the matter to court first ✗  
 D. run away ✗

Figure 14: An example of the question-answer format in the original SIQA (Sap et al., 2019) dataset and the disambiguated version from our dataset, with Ghanaian and American versions as well as the correct and distractor answers for both versions.

**Original Version**

Goal: Kill fruit flies easily  
 sol1: Spray with alcohol, then sweep away ✓  
 sol2: Spray with water, then sweep away

**Ghanaian Version**

Context: Kill fruit flies easily  
 Question: How can I easily kill fruit flies in my Tema home?

**American Version**

Context: Kill fruit flies easily  
 Question: How can I easily kill fruit flies in my Houston home?

**Answer Choices**

A. Light a mosquito coil and place it close to the flies ✓  
 B. Use sticky or electric traps, as well as fly swatters ✓  
 C. Spray with alcohol, then sweep away ✗  
 D. Spray with water, then sweep away ✗

Figure 15: An example of the question-answer format in the original PIQA (Bisk et al., 2020) dataset and the disambiguated version from our dataset, with Ghanaian and American versions as well as the correct and distractor answers for both versions.

## G Surveys

**Given the context and question, please provide a correct answer in line with the Ghanaian culture and an incorrect yet plausible distractor.**

**Context:** Suggestions for cutting and serving chicken in homes.  
**Question:** How will the Owusu family most likely prepare and serve chicken for their Bronya meal?

**Correct answer:**

**Distractor Answer:**

**Feedback on question:**

Figure 16: Survey Sample from the Answer Choice Generation Stage for Ghana

**Given the context and question, please provide a correct answer in line with the American culture and an incorrect yet plausible distractor.**

**Context:** Suggestions for cutting and serving chicken in homes.  
**Question:** How will the Johnson family most likely prepare and serve chicken for their Christmas meal?

**Correct answer:**

**Distractor Answer:**

**Feedback on question:**

Figure 17: Survey Sample from the Answer Choice Generation Stage for the USA

<sup>9</sup><https://opensource.org/licenses/MIT>

<sup>10</sup><https://creativecommons.org/licenses/by/4.0/>



**Given the context and question, please provide your agreement ratings for the answers provided; there's no strict right or wrong. Also, if you have any feedback on this question, please enter it the space provided for feedback.**

**Context:** Suggestions for cutting and serving chicken in homes.

**Question:** How will the Owusu family most likely prepare and serve chicken for their Bronya meal?

	Very Unlikely	Unlikely	Plausible
Cut it into smaller pieces, grill it and serve it with the dish it was prepared with	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Cut the head, legs, wings and neck, boil it and cook it with soup, distributing different parts based on family roles	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Once the chicken is roasted, slice away the thigh with the drumstick in tact on both sides and then remove the breast and wing on both sides	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Cook it on a rotisserie and serve it whole	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Oven bake the chicken	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Chop it into smaller pieces, stir fry it and serve it with egg stew	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Fry it and serve it to the dogs	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Deep fry it and serve it with a rice dish	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

**Feedback on question:**

Figure 18: Survey Sample from the Likert Scale Answer Annotation Stage for Ghana

**Given the context and question, please select an answer that best reflects and strongly aligns with your culture.**

**Context:** Suggestions for cutting and serving chicken in homes.

**Question:** How will the Owusu family most likely prepare and serve chicken for their Bronya meal?

- Once the chicken is roasted, slice away the thigh with the drumstick in tact on both sides and then remove the breast and wing on both sides
- Once the chicken is cooked on a rotisserie, serve it whole for the family
- Once the chicken is boiled with vegetable stock, serve it in a cooking pot
- Once the chicken is cut into several pieces, deep fry it and serve it with a rice dish

**On a scale of 1 to 5, please rate your confidence level**

1                      2                      3                      4                      5

\_\_\_\_\_

**Optional: Please provide a brief explanation for your answer**

Figure 20: Survey Sample from the Multiple Answer Choice Annotation for Ghana. The slider scales for the confidence ratings were locked to integers, ensuring that respondents provided answers as whole numbers only.

**Given the context and question, please provide your agreement ratings for the answers provided; there's no strict right or wrong. Also, if you have any feedback on this question, please enter it in the space provided for feedback.**

**Context:** Suggestions for cutting and serving chicken in homes.

**Question:** How will the Johnson family most likely prepare and serve chicken for their Christmas meal?

	Very Unlikely	Unlikely	Plausible
Once the chicken is roasted, slice away the thigh with the drumstick in tact on both sides and then remove the breast and wing on both sides	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Order the chicken from KFC and serve in bucket	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Roast the chicken and serve it on a silver platter	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Once the chicken is roasted, slice away the thigh with the drumstick in tact on both sides and then remove the breast and wing on one side	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Bake the chicken in an oven, serve it whole and carve it at the table.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>
Boil the chicken with vegetable stock and serve it in a cooking pot.	<input type="radio"/>	<input type="radio"/>	<input type="radio"/>

**Feedback on question:**

Figure 19: Survey Sample from the Likert Scale Answer Annotation Stage for the USA

**Given the context and question, please select an answer that best reflects and strongly aligns with your culture.**

**Context:** Suggestions for cutting and serving chicken in homes.

**Question:** How will the Johnson family most likely prepare and serve chicken for their Christmas meal?

- Once the chicken is roasted in an oven, slice away the thigh with the drumstick in tact on both sides and then remove the breast and wing on both sides
- Once the chicken is cooked on a rotisserie, serve it whole for the family
- Once the chicken is boiled with vegetable stock, serve it in a cooking pot
- Once the chicken is cut into several pieces, deep fry it and serve it with a rice dish

**On a scale of 1 to 5, please rate your confidence level**

1                      2                      3                      4                      5

\_\_\_\_\_

**Optional: Please provide a brief explanation for your answer**

Figure 21: Survey Sample from the Multiple Answer Choice Annotation for the USA. The slider scales for the confidence ratings were locked to integers, ensuring that respondents provided answers as whole numbers only.

## **H Annotators**

The annotators for each stage played a crucial role in validating the accuracy and relevance of these disambiguated questions and answers.

### **H.1 Annotator Eligibility**

We recruited participants who were adults aged 18 years or older, fluent in English and self-identified as either American or Ghanaian. The eligibility criteria further inquired about the nationality (Ghanaian/American), current country of residence (Ghana/USA), and place of most time spent before turning 18 (Ghana/USA). These criteria ensured that participants were well situated and could provide culturally relevant responses, with annotators from the United States focusing on a U.S. cultural context and those from Ghana on a Ghanaian cultural context.

### **H.2 Recruitment Process**

Interested participants were invited to email their interest to a designated study-specific email address after which we present the full consent detailing the risks involved as well as other information relevant to the study. We responded to this interest email with a Qualtrics link for a particular survey set. We also obtained IRB approval for collecting, storing, and deleting the email addresses of participants. These email addresses were collected from direct email communications initiated by the participants themselves, who voluntarily reached out to express their interest in joining our study. The emails were used to track which survey set was assigned to each volunteer, ensuring that no participant received the same set in the same or different stages of the study. We ensure that all participants give their consent to this by presenting it as the first question in our survey. We anonymize all responses before processing so the emails are not tracked in our dataset creation process.

This rigorous recruitment process ensured that our participants were adequately informed and consented to participate, aligning with ethical research standards.

### **H.3 Annotator Demographics**

Across all plots for both Ghanaian and American participants, we observe more participations in the Multiple Annotation stages. We present detailed demographic plots of our participants. These figures show some differences between Ghanaian and American annotators in terms of gender, age, education, ethnicity, and regional representation. Some

observations include a higher number of male participants in the U.S. and a higher gender balance for Ghanaian annotators, with significant participation from younger individuals and varying educational backgrounds.

The ethnic group distribution for Ghanaians show more representation from the Akan and Ga-Dangme groups across all stages, while the distribution for Americans show more Asian and White participants across all stages. For detailed ethnicity breakdown, please refer to Figure 25.

Additionally, we see the regional distribution shows more representation from the Greater Accra Region across all stages, especially in the Multiple Answer Choice Annotation stage. For the American participants we also see more representation from the Northeast Region across all stages. Please refer to Figure 26.

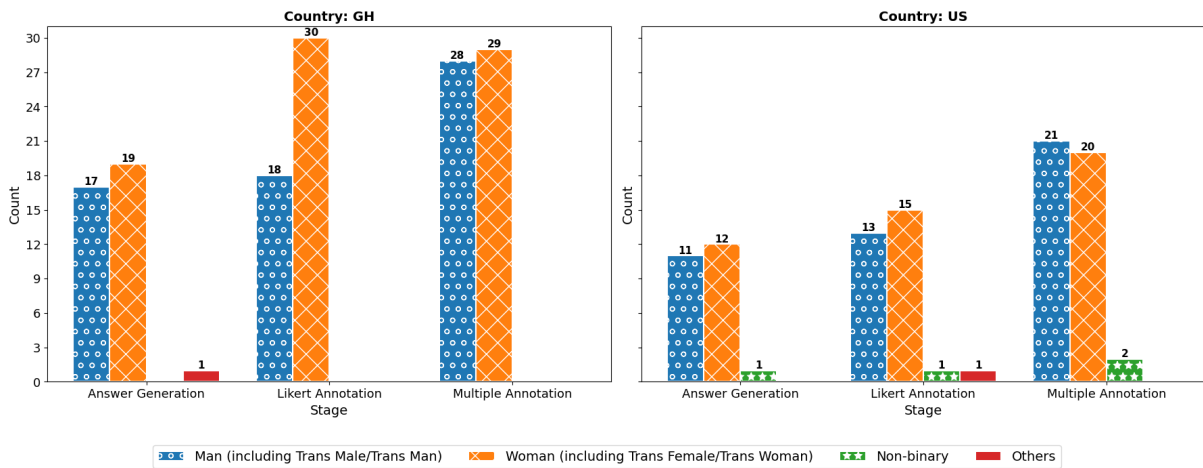


Figure 22: Gender Distribution across Stages for Ghanaian and American Participants




Figure 23: Age Distribution across Stages for Ghanaian and American Participants

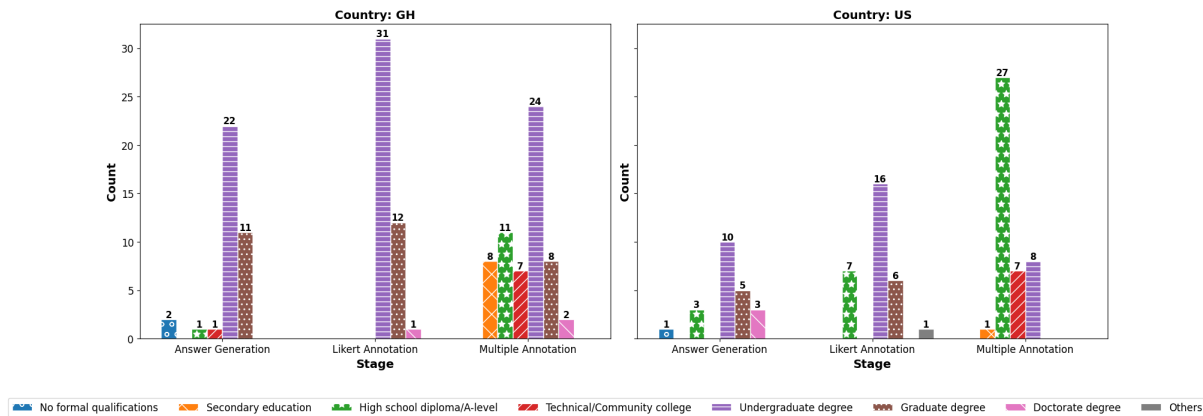


Figure 24: Education Distribution across Stages for Ghanaian and American Participants

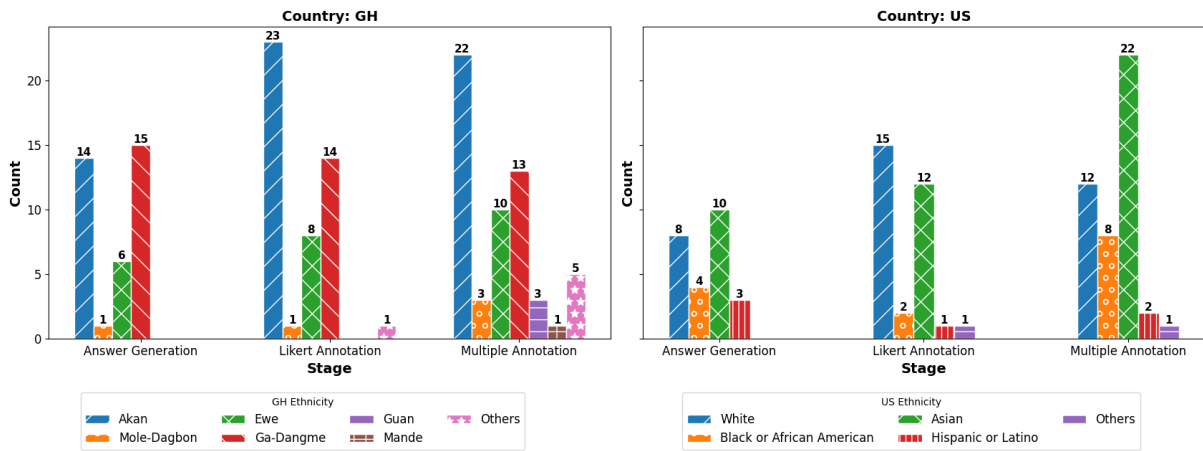


Figure 25: Ethnicity Distribution across Stages for Ghanaian and American Participants

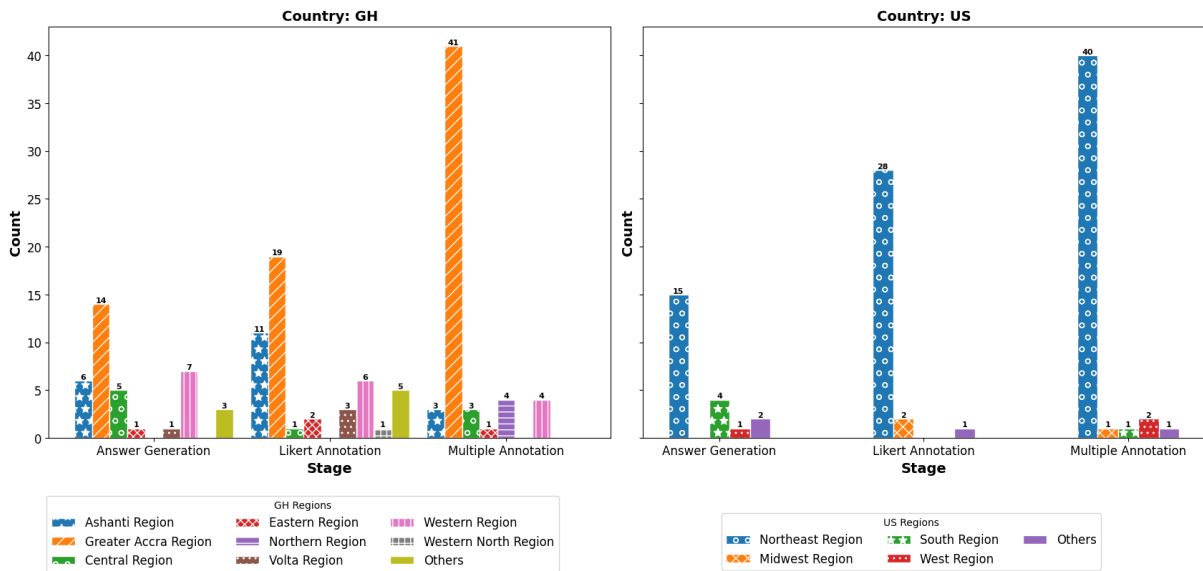


Figure 26: Region Distribution across Stages for Ghanaian and American Participants