Location Not Found: Exposing Implicit Local and Global Biases in Multilingual LLMs

Guy Mor-Lan[G*] **Omer Goldman**[BC*†] **Matan Eyal**[G] **Adi Mayrav Gilady**[G]
Sivan Eiger[G] **Idan Szpektor**[G] **Avinatan Hassidim**[G] **Yossi Matias**[G] **Reut Tsarfaty**[BG]

[G]Google Research [B]Bar-Ilan University [C]University of Cambridge

guymorlan@google.com, {omer.goldman, reut.tsarfaty}@gmail.com

Abstract

Multilingual large language models (LLMs) have minimized the fluency gap between languages. This advancement, however, exposes models to the risk of biased behavior, as knowledge and norms may propagate across languages. In this work we aim to quantify models' *inter-* and *intra-lingual biases*, via their ability to answer *locale-ambiguous* questions. To this end, we present LocQA, a test set containing 2,156 questions in 12 languages, referring to various locale-dependent facts such as laws, dates, and measurements. The questions do not contain indications of the locales they relate to, other than the querying language itself. LLMs' responses to LocQA locale-ambiguous questions thus reveal models' implicit priors. We used LocQA to evaluate 32 models, and detected two types of structural biases. *Inter-lingually*, we show a global bias towards answers relevant to the US-locale, even when models are asked in languages other than English. Moreover, we discovered that this global bias is exacerbated in models that underwent instruction tuning, compared to their base counterparts. *Intra-lingually*, we show that when multiple locales are relevant for the same language, models act as *demographic probability engines*, prioritizing locales with larger populations. Taken together, insights from LocQA may help in shaping LLMs' desired local behavior, and in quantifying the impact of various training phases on different kinds of biases.[1]

1 Introduction

When communicating in natural language, it is the rule rather than the exception that human speakers omit "obvious" information, giving rise to various ambiguities (Grice, 1991). How do LLMs cope with such ambiguities? In this paper we focus on a specific kind of ambiguity, namely, *locale-ambiguity*. Consider, for instance, the following seemingly straight-forward question: *"What is the emergency phone number?"* or, *"When does the tax year end?"*. These English questions are inherently ambiguous, as different locales entail different answers. We conjecture that models' answers to such ambiguous questions can reveal their implicit biases, as the ambiguity resolution exposes the model's latent preferences, revealing which regional reality it treats as the standard, and which realities it might erase.

Alternatively, a user may ask the same question in French, e.g., *"Quand commence l'exercice fiscal?"*. In this case, we expect the model to shift its frame of reference away from the Anglosphere. This is tricky, as the prevailing assumption in multilingual NLP is that querying a model in a specific target language acts as a proxy for context. So in theory, the choice of language should narrow the scope of ambiguity. However, a single language *rarely* isolates a single locale. In the case of French, for instance, it is the official language of 29 countries, spanning from France and Switzerland to Haiti and the DRC. So, while the linguistic surface form is shared, the factual realities regarding laws, measurements, and infrastructure differ considerably across regions using the same language.

In this work we claim that current multilingual evaluations conflate two distinct capabilities of generative LLMs: (i) *Linguistic Fluency*, i.e., the ability to generate fluent and coherent text in a given target language, and (ii) *Localization*, i.e., grounding the generation in the relevant reality of the speakers of that language in different locales. While contemporary LLMs exhibit striking *fluency* on an ever increasing number of diverse languages, it remains unclear whether and to what extent they have truly learned to represent the diverse populations speaking those languages, or whether the generated content is a mere *fluent, albeit biased*, translation

[*]Equal contribution.

[†]Work done at Google Research.

[1]The data is available at https://github.com/google-research-datasets/locqa/.

Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 29130–29150
July 2-7, 2026

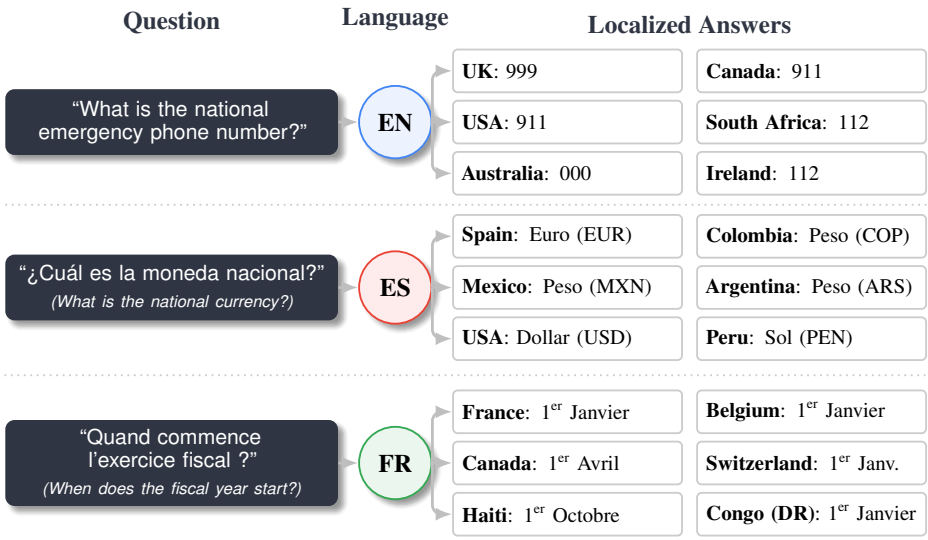

Figure 1: Schematic illustrating how identical queries in LocQA branch into distinct ground-truth answers depending on the target locale. Thus, language alone is insufficient for resolving factual ambiguity.

of Western norms.

In order to isolate and investigate the *localization* aspect, we suggest to analyze how models voluntarily resolve ambiguity in locale-ambiguous questions. Our investigation exposes two distinct axes along which models' behavior may be biased. First, we define a *Global Bias* as a measure of the extent to which a US-centric frame of reference persists across linguistic boundaries (e.g., a model employing US norms even when queried in Indonesian). Second, we define a *Regional Bias* which examines the implicit prioritization of specific locales *within* a shared language (e.g., when querying in Spanish, does the model default to Spain or Mexico?).

To measure both kinds of biases, we present LocQA (Localized QA), a diagnostic benchmark designed to probe the implicit priors of LLMs. Unlike previous cultural benchmarks that test explicit knowledge (e.g., *"What is the capital of Peru?"*), LocQA utilizes semantically invariant, locale-ambiguous queries. By analyzing which regional reality the model defaults to when the context is underspecified, we map models' tendencies, biases and implicit representation hierarchy.

Our investigation of 32 models reveals that models do not resolve ambiguity based on geographic fairness. Instead, we identify two structural skews. First, we observe a persistent *US-centric default*: even when queried in non-English languages, models frequently mention US norms, instead of or in addition to locale-relevant responses Second, we detect a *populist skew*: models function as "demographic probability engines," where the likelihood of a locale being represented is strongly related to its population size, effectively erasing smaller nations that share a major language. Finally, we show empirical evidence for a *Cultural Alignment Tax*. That is, when contrasting instruction-tuned models with their base counterparts, we show that instruction-tuned models exhibit *lower* Regional Bias but significantly *higher* US bias, suggesting that current alignment practices actively sacrifice cultural nuance, possibly in favor of a more generic, conceivably "safe", homogeneity.

In sum, the contributions of this paper are as follows. (i) we deliver the **LocQA Benchmark**, a validated diagnostic suite of 2,156 locale-specific answers to locale-ambiguous questions, across 12 languages and 49 regions, designed to isolate LLMs' default priors. (ii) we define a **Dual-Metric Framework** for quantifying implicit biases across two axes, including *Global* metrics to quantify US-centric bias, and *Regional* metrics for assessing geographic fairness, and (iii) we deliver **Empirical Evidence of Alignment Bias** across 32 LLMs. Ultimately, we argue that for LLMs to serve global audiences, geography should not be taken as a byproduct of language use. We call for a shift from multilingual modeling to multicultural and multi-regional modeling, where locale is treated as a distinct facet that must be accounted for to ensure factual adequacy for all users across the globe.

2 Challenges and Motivation

Multilingual LLMs, like their monolingual counterparts, must be able to *retrieve* knowledge, a task

that has been proven difficult in multilingual settings (Goldman et al., 2025; Lalai et al., 2025). However, multilingual LLMs are also tasked with the *selection* of the appropriate cultural frame to retrieve knowledge from. Thus, models should be tested not only for their *capabilities* (can the model answer X?) but also for their *propensities* (what does the model assume X is?). The gap between knowing a fact and selecting it is critical: a model may "know" the drinking age in Indonesia, but if it defaults to US norms when asked *in Indonesian*, that knowledge is effectively erased.

Existing cultural and regional benchmarks primarily evaluate a model's *capability* to retrieve specific knowledge or values. For instance, benchmarks like INCLUDE (Romanou et al., 2024) and Global-MMLU (Singh et al., 2025) test objective accuracy on culturally sensitive knowledge, while others like BLEnD (Myung et al., 2024) and GlobalOpinionsQA (Durmus et al., 2024) evaluate alignment with local everyday knowledge and subjective moral values. These benchmarks evaluate capability and not implicit locale bias. Moreover, they usually rely on *explicit* prompting, asking models "What is the norm in Country X?" (Chiu et al., 2025; Yin et al., 2022) or providing locale as context for reasoning (Rao et al., 2025). Even generative approaches as in Bhatia and Shwartz (2023) rely on explicit cues to trigger diversity. By naming the target locale in the prompt, they act as an oracle, resolving the ambiguity *for* the model and masking its biases in information selection. This is the factor we seek to measure, addressing the "explicit-implicit localization gap" (Veselovsky et al., 2025).

Unlike these works on explicit knowledge, subjective values, or natural phrasing (Hasan et al., 2025), we target the model's *unprompted* default behavior, revealing the geographic alignment that contemporary explicit benchmarks systematically miss. We measure the extent to which model behavior is driven by "epistemic inequity" (Wang et al., 2025) and defaults to the *dominant data distribution* rather than the linguistically relevant locale.

Measuring models' implicit biases will further provide quantification for the discussion on the growing concern, that this selection bias is exacerbated by the very processes that improve LLMs' multilingual capabilities. For example, Han et al. (2025) identify a "Transfer-Localization Trade-off," where cross-lingual optimization leads to cultural erasure, and Gao et al. (2024) note that instruction tuning often results in "shallow" alignment. Our work provides a diagnostic tool that will allow precise examination of the "taxes" imposed by those improvements, and answer the question: might the pursuit of a universal, safe, assistant, force models to converge on a single, US-biased reality?

3 The LocQA Dataset

This paper presents LocQA, a benchmark designed to answer the question: *what is the default reality assumed by a model in locale-ambiguous questions?* To construct LocQA we first came up with about a dozen example questions suited for exploring models' behavior under ambiguous conditions. The questions were relatively *time independent*, related to specific *facts with a well-defined answer*, as well as *easily translatable*, that is, without terms that require localization or whose translation is unclear in the target languages. Most importantly, the answers to the example questions had to be *locale-dependent*, where the expected answer may change according to the locale that the user has in mind and according to the language in which it is phrased. The example questions related to various topics: law, history, language, etc.

The example questions were then given to qualified bilingual vendor annotators proficient in the target languages (see guidelines in Appendix A), for translation into the 12 languages covered by LocQA: English, Spanish, French, German, Hebrew, Hindi, Indonesian, Italian, Japanese, Korean, Portuguese, and Chinese. All in all, we employed 16 annotators. For each language, the annotators gave the answers to the questions as they relate to the countries associated with that language. We targeted countries with at least one million native speakers of each language for inclusion (see Table 2). Note that we did not require all questions to have answers in all locales. Some questions, like *who is the first president?*, may not have answers in countries that never had presidents, so an *N/A* answer is valid. However, it must be clear from the question whether or not it has an answer.

To ensure data quality, all translations and locale-specific answers were cross-validated by a second independent annotator. Following this, the authors conducted a general manual review to resolve discrepancies and correct any remaining errors. The final dataset consists of 2,156 locale-specific questions and answers. These correspond to 44 semantically parallel questions (Appendix C) translated to 12 languages and answered for 49 locales.

4 Methodology

4.1 Metrics

We define metrics to detect biases in the generated answers compared to the locale-specific gold answers. Concretely, we define two metrics. One for *Global Bias* B_{US}, i.e., the skew in the generated answers towards the US answer. This metric is calculated over the answers in all non-English languages taken together. The other metric, the *Regional Bias* B_R, aims to detect *intra-lingual* biases. It indicates the countries whose gold answers are over- or under-represented in the generated answers, taking into account one language at a time.

Global Bias (B_{US}). We quantify the extent to which models default to United States norms, for example the extent to which the model answers *George Washington* to the question *Who was the first president?* or its translation. However, some US answers are not unique, so simple counting is insufficient. Consider the question in Indonesian *Berapa usia legal untuk minum alkohol?* (translated to *What is the legal drinking age?*) A model that answers *21* may give the US-centric answer as a default but it may also give the correct answer for Indonesia, which happens to be identical. We term such identiy of answers a *collision*. For that reason, B_{US} measures the *difference* between the frequency of the US answer in the model's answers and the frequency of that value in the data.[2] We compute B_{US} separately for each language and report the macro-average across the 11 non-English languages, so that multi-locale languages (e.g., Spanish, with 20 locales) do not dominate the aggregate.

Formally, for a language L with locale set $\mathcal{C}_L$, B_{US} is the difference between the observed and the expected probabilities of getting the US answer:

$$B_{US} = P_{\text{obs}}(A_{\text{US}}) - P_{\text{exp}}(A_{\text{US}}) \tag{1}$$

where A_{US} is the value of the US answer. The observed P_{obs} is calculated based on the model's outputs and the expected P_{exp} is based on the data:

$$B_{US} = \underbrace{\frac{1}{|\mathcal{Q}|}\sum_{q} \mathbb{I}\big(A_{\text{US}} \in M(q,L)\big)}_{\text{Observed}} - \underbrace{\frac{1}{|\mathcal{Q}|\,|\mathcal{C}_L|}\sum_{q,c\in\mathcal{C}_L} \mathbb{I}\big(A_{\text{US}} = A(q,c)\big)}_{\text{Expected}} \tag{2}$$

[2] 'N/A' is treated as a valid answer.

$\mathcal{Q}$ is the set of questions in LocQA and $M(q,L)$ is the response of the model to question q when asked in language L (one response per question per language); $A(q,c)$ is the gold answer for the same question in locale c. The expected term is *collision-aware*: it counts, per question, the fraction of locales in $\mathcal{C}_L$ whose gold answer coincides with the US answer. Note, that the model's response $M(q,L)$ may well include a list of multiple answers, so only inclusion of the US answer is needed. A positive B_{US} indicates the model prefers US norms beyond what would be expected from random chance overlap (e.g., shared drinking age or voltage standards).

Regional Bias (B_R). This metric quantifies the model's preference for a specific locale. It compares $N_{\text{model}}(c)$—the number of times an answer valid to locale c appears in model predictions, with $N_{\text{data}}(c)$—the number of times that locale c's answer appears in the LocQA dataset for this question. Both counts are *collision-aware*, that is, counting each answer towards all the locales that it is valid for (e.g., "Peso" applying to multiple countries). This is done in order to account for shared norms and coincidental overlap in answers.

Concretely, for each question, a gold answer held by m locales contributes m to the N_{data} count of each of those locales. $N_{\text{model}}(c)$ is incremented by 1 for every question whose model response contains a match for c's gold answer (recall that collisions arise when the same response matches the gold answers of multiple locales).

Formally, for a language L with locale set $\mathcal{C}_L$, we define:

$$P_{\text{obs}}(c) = \frac{N_{\text{model}}(c)}{\sum_{k\in\mathcal{C}_L} N_{\text{model}}(k)} \tag{3}$$

$$P_{\text{exp}}(c) = \frac{N_{\text{data}}(c)}{\sum_{k\in\mathcal{C}_L} N_{\text{data}}(k)} \tag{4}$$

The Regional bias is then defined as the lift:

$$B_R(c) = \frac{P_{\text{obs}}(c)}{P_{\text{exp}}(c)} \tag{5}$$

$B_R(c) > 1$ indicates over-representation (*dominance*), while $B_R(c) < 1$ indicates under-representation (*erasure*). To obtain a single bias score per model, we compute the mean deviation $|B_R(c) - 1|$ within each language and then macro-average across languages with more than one locale.

4.2 Automatic Evaluation

To evaluate model outputs at scale, we employ a 2-stage pipeline using *Gemini-2.5-Flash*, selected for its high instruction-following capability and low latency (prompts for this are given in Appendix E.)[3]

Initially we assess **(i) Ground Truth Alignment.** While answers within the same target language share identical string representations, the US reference answer often differs in language or formatting (e.g., *'1 de Enero'* vs. *'January 1st'*). To properly detect such answer collisions, we employ a semantic matching prompt that identifies when a locale-specific answer is semantically equivalent to the US norm.

Next, we turn to **(ii) Response Analysis** as our primary evaluation method. We analyze model responses using an LLM-as-a-Judge to extract two key signals: *Mentioned Answers*, which identifies which of the locale-relevant gold answers are explicitly provided by the model as valid options; and *Framing Style*, which detects whether the response uses the US as a conceptual anchor (e.g., "Unlike in the US..."), even when the US answer itself is not offered as a valid option. To verify the reliability of this automated pipeline, we manually evaluated a random sample of 80 judgments, finding a 92% agreement rate between human annotations and the LLM judgments.

For experiments testing models' responses when explicitly specifying a desired locale, we use a verification prompt that checks if the model successfully retrieves the specific locale's answer and if it hallucinates the US answer.

5 Experiments

Setup. We evaluate a diverse suite of 32 models, both proprietary and open-weights models. To analyze the impact of alignment, we test both *base* and *instruction-tuned* variants for Gemma 3 (4B, 12B, 27B), Qwen (2.5-72B; 3-4B, 8B, 14B), GLM-4 (9B), OLMo-3 (7B, 32B), Falcon 3 (10B) and IBM Granite 3 (8B). The suite also includes Qwen 3 (235B), DeepSeek (V3, R1), Mistral (Small, Large), and Kimi K2. Finally, we evaluate proprietary models including GPT (4o, 4.1, 5-mini, 5.1, o1, o3), Claude 4.5 (Sonnet, Opus), Gemini (2.5 Flash/Pro, 3 Pro), and Grok (3, 4). Models

[3]We verify the robustness of our pipeline by repeating all evaluations using *GPT-5-mini*, which yielded strong alignment with our primary judge across Global Bias ($r = 0.99$), Regional Bias ($r = 0.95$), and Framing ($r = 0.85$).

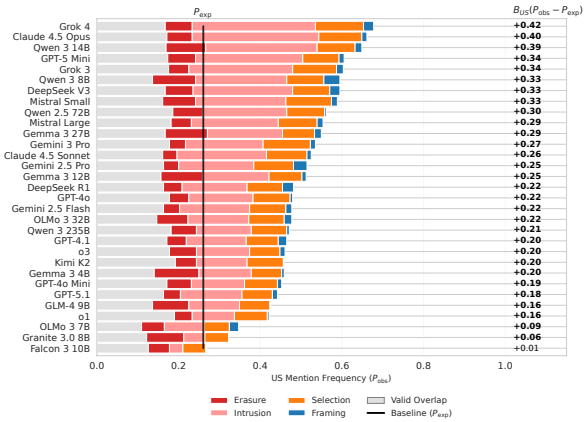

Figure 2: Global Bias scores across models (on the right) show the difference between P_{obs} (the sum of all bars) and P_{exp} (the black line). Color give a breakdown of US-centric answers into categories, defined in section 5.2. Intrusion (gratuitous inclusion) and selection (prioritizing US options) are most prevalent, occurring significantly more than complete erasure.

are evaluated in zero-shot format, with only the question as input, no instructions or examples.[4]

5.1 Results

Global Bias. Figure 2 summarizes the results in terms of the Global Bias B_{US} for all models over all questions of LocQA. Almost all models demonstrate a clear US bias. The magnitude of that bias varies widely across models, from approximately 0 for Falcon 3 to 0.42 for the most biased Grok 4. The average B_{US} across all models is 0.24, reflecting the difference between the frequency of US answers in the data (26%) and the frequency of these answers in the models' outputs (50%).[5]

Regional Bias. Figure 3 displays Regional Bias scores in each locale for every evaluated model. 4 languages with a single locale are omitted from this analysis. We see that despite variations between models, a consistent set of locales tend to be over- or under-represented. The results reveal a distinct ordering of locales, whereby large population centers (e.g., USA, Brazil) and Western countries (e.g.,

[4]Following Kabir et al. (2025), who highlight the limitations of forced-choice in cultural evaluation, we employ open-ended generation rather than multiple-choice questions to capture the model's unprompted default.

[5]In 18.2% of responses across our 0-shot instruct suite, the judge extracts no gold-answer candidate (neither a locale-valid answer nor the US value). We verify in Appendix F that this does not drive our findings: B_{US} rankings are preserved when conditioned on at least one candidate being extracted ($\rho = 0.81$), and B_R is mechanically unaffected by responses that yield no match.

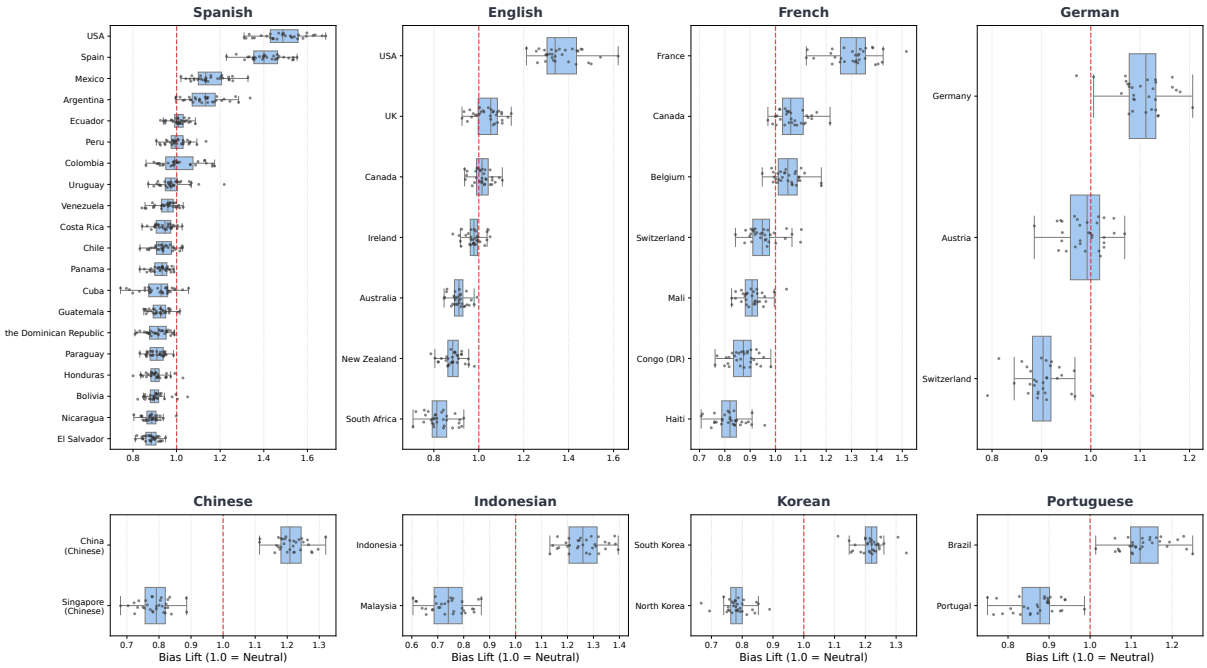

Figure 3: Distribution of Regional Bias scores across locales. The results reveal a structural inequality where Western nations and large population centers are consistently over-represented compared to peripheral locales.

Spain, France) maintain high scores while peripheral locales suffer systematic erasure. This allows us to identify regional *winners* and *losers*. In Spanish, the USA, Spain, Mexico and Argentina emerge as over-represented, whereas Honduras, Bolivia, Nicaragua and El Salvador are under-represented. In English, The USA is over-represented, whereas South Africa, New Zealand and Australia suffer from under-representation. In French, France is over-represented, whereas Haiti, Congo and Mali emerge as under-represented.

5.2 Analysis

Anatomy of US Bias. We categorize US-centric responses into five types: *Erasure* (replacing local reality with US norms), *Intrusion* (gratuitously inserting US answers alongside correct local ones), *Selection* (preferring the US-aligned option in ambiguous contexts), *Framing* (explicitly contrasting the local answer against the US), and *Valid Overlap* (coincidental correctness, serving as a control). Table 1 provides examples of these categories. Model outputs were classified into these categories according to the LLM judge's identification of matched answers and explicit detection of framing. The breakdown is given in Figure 2. We see that the largest error mass is attributed to intrusion (17.5% on average across models), where models provide the correct local answer but gratuitously insert US norms, treating them as a necessary context. However, a concerning 6.5% of responses exhibited erasure, where local reality is completely overwritten by US norms. In multi-locale ambiguity, selection accounts for an additional 8.7% on average, indicating a systematic preference for the US-aligned option over other valid alternatives. Explicit framing remains rare (14% on average), suggesting that the US bias manifests as an implicit default rather than a conscious comparison.

Population and Regional Bias. Having established that models exhibit biases across locales of the same language, we investigate the dominant factor driving this behavior. In Figure 4, we plot the Regional Bias $B_R(c)$, averaged across models, against the log-scaled speaking-population of each locale (population data sources are listed in Appendix D). We tested linear, power-law, and logarithmic fits for the data. Our empirical analysis reveals that a logarithmic function of population best models the data ($R^2 = 0.41$), significantly outperforming a hypothesis of linear proportionality to population size hypothesis ($R^2 = 0.14$). This demonstrates that the observed bias is a structural property of model training: representation scales with the *order of magnitude* of the population rather than its raw count. This logarithmic relationship indicates the diminishing returns of population size. While the correlation is strong

Category	Language	Question	Model Response (Translated Excerpt)
Erasure	French	*Quand la constitution a-t-elle été signée? (When was the constitution signed?)*	"The United States Constitution was signed on **September 17, 1787**." *(Hallucinates US history onto French context)*
Intrusion	Korean	*제 국가 번호가 어떻게 되나요? (What is my international phone code?)*	"It depends on your location... For example: **Korea: +82**, **US/Canada: +1**, Japan: +81..." *(Gratuitous inclusion of US code)*
Selection	Spanish	*¿Cuál es el número de teléfono de emergencia nacional? (What is the national emergency phone number?)*	"**911**." *(Valid for Mexico and US, but incorrect for Spain which uses 112)*
Framing	Spanish	*¿Cuántas formas conjugadas tienen los verbos? (How many inflected forms do verbs have?)*	"The answer depends on the language... Here is the detail for **Spanish**, and as a comparison, for **English**." *(Explicit framing against US norm)*
Valid Overlap	Indonesian	*Berapa usia legal untuk minum alkohol? (What is the legal drinking age?)*	"The legal age for consuming alcohol is **21**." *(Coincidental match: Indonesia shares the US age of 21)*

Table 1: Taxonomy of US Bias. Examples of the five error modes identified in our analysis.

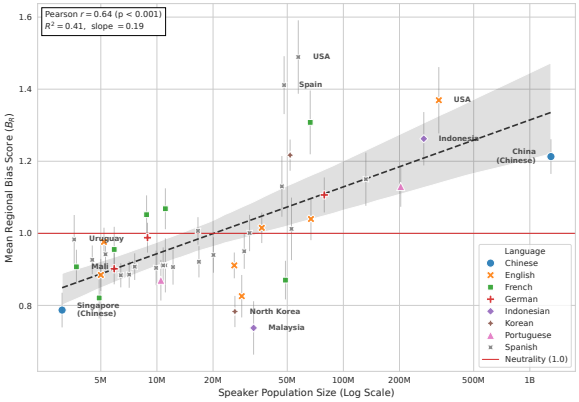

Figure 4: Average Regional Bias (B_R) plotted against the log-scaled speaking population of each locale. The strong logarithmic fit ($R^2 = 0.41$) suggests that representation scales with the order of magnitude of the population rather than raw census counts.

($r = 0.64, p < 0.001$), the functional form imposes a "soft ceiling" on demographic giants. For example, the estimated slope of 0.19 implies that a locale must grow its population by a factor of 10 just to gain 0.19 points in representation score. Consequently, this logarithmic compression suggests that models scale with population magnitude rather than raw counts, effectively dampening extreme demographic disparities and maintaining baseline visibility for the long tail.

Domain-Wise Bias. To understand if specific topics disproportionately drive these biases, we categorized LocQA into five domains (see Appendix G for the full data). We observe a striking divergence between Global and Regional bias triggers. Questions regarding *State and Country* (e.g., government, infrastructure) and *Language* exhibit the highest US-centric default ($B_{US} \approx 0.23$–0.30) but relatively low regional distortion. Conversely, questions regarding *Leisure and Culture* (e.g., sports, retirement) successfully avoid the US default ($B_{US} = 0.07$) but exhibit the most extreme Regional Bias ($|B_R - 1| = 0.69$). This indicates that while culturally grounded topics escape a US-centric default, they heavily trigger the "demographic probability engine," causing models to aggressively default to the most populous local nations instead of maintaining regional fairness.

Instruction Tuning and Bias. Having seen the prevalence of regional and global US bias across different models, we examine the factors behind the biases. First, we investigate whether applying instruction tuning to multilingual models exacerbates their biases. We extend the concept of the "Alignment Tax" (Ouyang et al., 2022; Lin et al., 2024), to detect whether improving the models' ability to follow instructions in multilingual settings entails more significant bias. We examine this by comparing the global and regional biases of base open-weight models in the 4B-72B range against their instruction-tuned counterparts. In this comparison, we utilize a 3-shot prompting strategy for both model types. This ensures that the base models are not penalized for formatting failures. The examples in the prompt are three simple, locale-neutral QA pairs (e.g., arithmetic) that only guide format adherence without priming regional biases (see Appendix E for the prompt).

Figure 5 illustrates the impact of instruction tuning on both biases. The left panel plots the Global Bias score (B_{US}). We observe a consistent "Alignment Tax": across all model families, instruct-

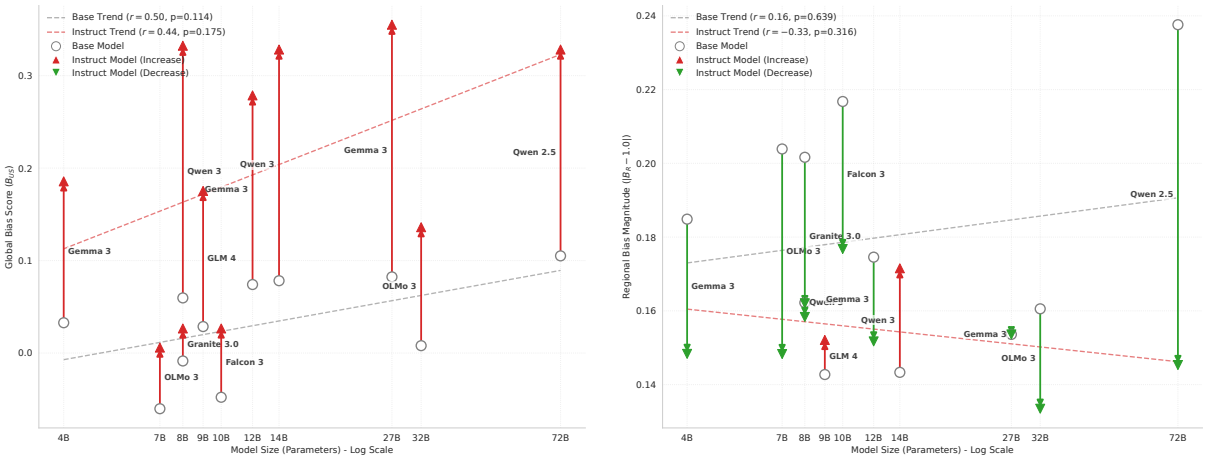

(a) Impact on Global Bias (B_{US})

(b) Impact on Regional Bias Magnitude ($|B_R - 1.0|$)

Figure 5: Comparison of cultural biases in base versus instruction-tuned models. Panel (a) shows that instruction tuning consistently increases Global Bias ("Alignment Tax"), while panel (b) shows it tends to reduce Regional Bias magnitude, flattening representation across locales.

tuned models exhibit significantly higher US bias compared to their base counterparts. Furthermore, this bias scales with capability; larger models display consistently higher bias in both base and instruct regimes, suggesting that as models become more capable of retrieving cultural knowledge, they increasingly default to US-centric views.

Conversely, the right panel displays the difference in Regional Bias. Since in this case over-representation and under-representation are both unwanted, we calculated for each model the mean absolute deviation of Regional Bias scores from neutrality ($|B_R - 1|$). Here, we observe the opposite trend: instruction fine-tuning tends to *reduce* regional distortion. Base models generally exhibit higher Regional Bias (indicating the dominance of specific locales or erasure of others) and instruct models achieve lower scores. This suggests that alignment tuning "flattens" the representation across locales.

We hypothesize that these opposing trends stem from the tendency of instruction following training to motivate models to maximize helpfulness by offering "diverse" and inclusive responses.

To support this hypothesis, we measure the models' *answer multiplicity*, defined as the average number of distinct answers provided per question that are valid for *some* locale. Figure 6a confirms that instruction tuning systematically increases the average number of answers listed per question across all models. As shown in Figure 6b and Figure 6c, the increase in multiplicity is strongly correlated with the rise in Global Bias ($r = 0.95$) and moderately correlated with the reduction in Regional Bias ($r = 0.47$).

This indicates that alignment transforms models from *local simulators*, which commit to a single local answer, into *global observers* that strive for diversity. By listing multiple valid options, instruct models dilute the dominance of any single locale, driving B_R towards neutrality. However, this diversity is not neutral or evenly distributed, but rather itself selectively biased. The models learn to diversify their answers, but they consistently choose the US as the counterpoint or anchor for additional context. Thus, while alignment successfully reduces the erasure of local norms, it re-introduces bias through the very mechanism of diversity itself, framing the US as the universal reference even in non-English contexts.

Undoing Ambiguity: Explicit Locale Prompting. Finally, we investigate the nature of Global Bias when ambiguity is removed. We re-evaluated all models using an *explicit prompt* (e.g., "Locale: Mexico. What is the currency?"). A specialized judge (see Appendix E) verified if the model retrieves the correct local answer or hallucinates the US one. This tests the "stickiness" of the bias: does the preference for US norms persist even when explicitly directed to another locale? Figure 7 plots model accuracy against the *US Hallucination Share*, i.e., the percentage of errors where the model substitutes the US answer.

We observe a moderate correlation ($r = 0.49, p = 0.004$) between model performance and

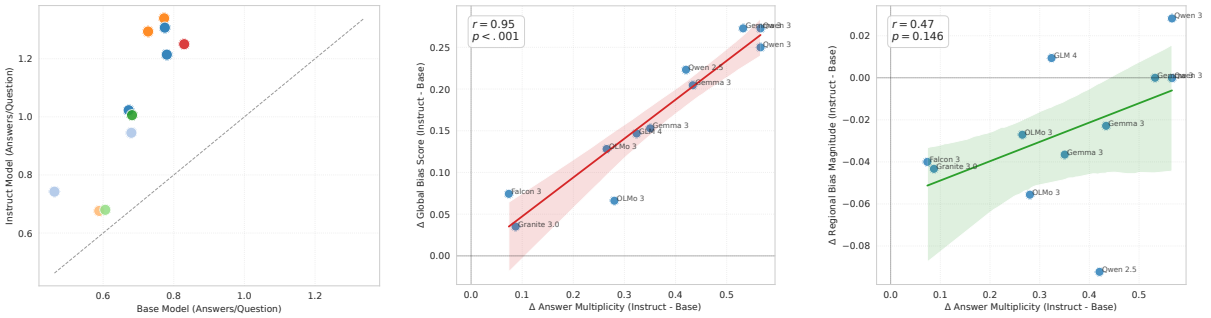

(a) Shift in Multiplicity (b) Correlation w/ Global Bias (c) Correlation w/ Reg. Bias Magnitude

Figure 6: Analysis of answer multiplicity. Instruction-tuning systematically increases the number of valid answers per question (a). This shift strongly correlates with the rise in US-centric bias (b) and weakly correlated with a reduction in Regional Bias (c).

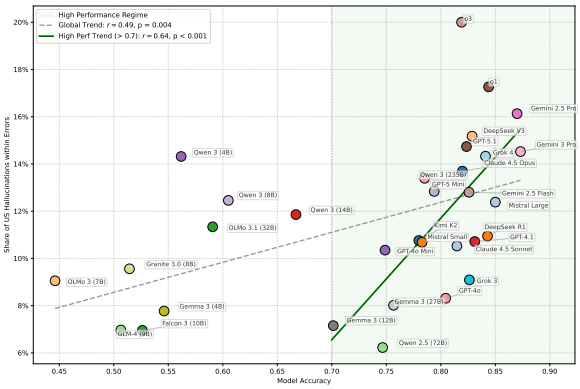

Figure 7: Model accuracy plotted against the share of errors where the model gives a US answer despite explicit locale prompting. The positive correlation indicates that US-centric bias persists in highly capable models.

US hallucinations on the full sample. Moreover, among strong models ($> 70\%$ accuracy), we see a stronger positive correlation ($r = 0.64$). That is, as models become more capable and make fewer mistakes overall, the errors that *do* persist are increasingly US-centric. The fact that this correlation is most pronounced for high-accuracy models confirms that US bias is not merely a random fallback for missing knowledge. Instead, while general capabilities may eliminate random noise, the US prior is persistent. This challenges the notion that "scaling is all you need" for multilingual LLMs, that is, that larger models will naturally converge on better multicultural representation. Rather than vanishing with increased capability, we see that US-centricity remains a sticky and proportionally larger residual failure mode.

6 Conclusion

We introduce LocQA, a diagnostic benchmark to measure the implicit geopolitical priors of multilingual LLMs. By evaluating model responses to locale-ambiguous queries, we uncover two structural biases in how LLMs handle ambiguity. First, models exhibit a *Regional Bias*, systematically favoring dominant locales in line with population size. Second, we identified a pervasive *Global Bias* toward US norms in different non-English languages. Crucially, our experiments also revealed a *Cultural Alignment Tax*—while instruction tuning reduces the Regional Bias, it does so by increasing US-centricity. We trace this to a learned behavior of "answer multiplicity": aligned models attempt to be helpful by providing more options, but systematically select the US as a default reference for comparison. Finally, we showed that when highly capable models are wrong, they are *more* likely to override correct local facts with US norms.

These findings challenge the prevailing assumption that linguistic fluency acts as a sufficient proxy for local and cultural grounding. To build truly global systems, the community must move beyond linguistic fluency and explicitly evaluate localization and cultural grounding, ensuring that alignment for safety does not come at the cost of cultural erasure.

Limitations

While LocQA provides a rigorous framework for evaluating implicit localization, our study has several limitations. First, our dataset covers 12 languages and 49 locales; while diverse, this represents a fraction of global linguistic diversity. Extremely low-resource languages and dialects re-

main unprobed. Second, the ground-truth answers in LocQA (e.g., tax dates, voltage) may be subject to legislative and infrastructural change. Third, our automated evaluation relies on an LLM-as-a-Judge pipeline, which may not perfectly replicate the nuance of human evaluation. Finally, our analysis focuses on *factual* localization; we do not evaluate the model's alignment with subjective cultural values or moral norms, which represents a distinct but equally important dimension of cultural capability.

References

Mehar Bhatia and Vered Shwartz. 2023. GD-COMET: A geo-diverse commonsense inference model. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7993–8001, Singapore. Association for Computational Linguistics.

Yu Ying Chiu, Liwei Jiang, Bill Yuchen Lin, Chan Young Park, Shuyue Stella Li, Sahithya Ravi, Mehar Bhatia, Maria Antoniak, Yulia Tsvetkov, Vered Shwartz, and Yejin Choi. 2025. CulturalBench: A robust, diverse and challenging benchmark for measuring LMs' cultural knowledge through human-AI red-teaming. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 25663–25701, Vienna, Austria. Association for Computational Linguistics.

Esin Durmus, Karina Nguyen, Thomas I. Liao, Nicholas Schiefer, Amanda Askell, Anton Bakhtin, Carol Chen, Zac Hatfield-Dodds, Danny Hernandez, Nicholas Joseph, Liane Lovitt, Sam McCandlish, Orowa Sikder, Alex Tamkin, Janel Thamkul, Jared Kaplan, Jack Clark, and Deep Ganguli. 2024. Towards measuring the representation of subjective global opinions in language models. *Preprint*, arXiv:2306.16388.

Changjiang Gao, Hongda Hu, Peng Hu, Jiajun Chen, Jixing Li, and Shujian Huang. 2024. Multilingual pretraining and instruction tuning improve cross-lingual knowledge alignment, but only shallowly. 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 6101–6117, Mexico City, Mexico. Association for Computational Linguistics.

Omer Goldman, Uri Shaham, Dan Malkin, Sivan Eiger, Avinatan Hassidim, Yossi Matias, Joshua Maynez, Adi Mayrav Gilady, Jason Riesa, Shruti Rijhwani, Laura Rimell, Idan Szpektor, Reut Tsarfaty, and Matan Eyal. 2025. Eclektic: a novel challenge set for evaluation of cross-lingual knowledge transfer. *Preprint*, arXiv:2502.21228.

Paul Grice. 1991. *Studies in the Way of Words*. Harvard University Press.

HyoJung Han, Sweta Agrawal, and Eleftheria Briakou. 2025. Rethinking cross-lingual alignment: Balancing transfer and cultural erasure in multilingual llms. *Preprint*, arXiv:2510.26024.

Md. Arid Hasan, Maram Hasanain, Fatema Ahmad, Sahinur Rahman Laskar, Sunaya Upadhyay, Vrunda N Sukhadia, Mucahid Kutlu, Shammur Absar Chowdhury, and Firoj Alam. 2025. NativQA: Multilingual culturally-aligned natural query for LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 14886–14909, Vienna, Austria. Association for Computational Linguistics.

Mohsinul Kabir, Ajwad Abrar, and Sophia Ananiadou. 2025. Break the checkbox: Challenging closed-style evaluations of cultural alignment in LLMs. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 24–51, Suzhou, China. Association for Computational Linguistics.

Harsh Nishant Lalai, Raj Sanjay Shah, Jiaxin Pei, Sashank Varma, Yi-Chia Wang, and Ali Emami. 2025. The world according to llms: How geographic origin influences llms' entity deduction capabilities. *Preprint*, arXiv:2508.05525.

Yong Lin, Hangyu Lin, Wei Xiong, Shizhe Diao, Jianmeng Liu, Jipeng Zhang, Rui Pan, Haoxiang Wang, Wenbin Hu, Hanning Zhang, Hanze Dong, Renjie Pi, Han Zhao, Nan Jiang, Heng Ji, Yuan Yao, and Tong Zhang. 2024. Mitigating the alignment tax of RLHF. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 580–606, Miami, Florida, USA. Association for Computational Linguistics.

Junho Myung, Nayeon Lee, Yi Zhou, Jiho Jin, Rifki Afina Putri, Dimosthenis Antypas, Hsuvas Borkakoty, Eunsu Kim, Carla Perez-Almendros, Abinew Ali Ayele, Víctor Gutiérrez-Basulto, Yazmín Ibáñez García, Hwaran Lee, Shamsuddeen Hassan Muhammad, Kiwoong Park, Anar Sabuhi Rzayev, Nina White, Seid Muhie Yimam, Mohammad Taher Pilehvar, and 3 others. 2024. Blend: A benchmark for llms on everyday knowledge in diverse cultures and languages. In *Advances in Neural Information Processing Systems*, volume 37, pages 78104–78146. Curran Associates, Inc.

Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. Training language models to follow instructions with human feedback. *Preprint*, arXiv:2203.02155.

Abhinav Sukumar Rao, Akhila Yerukola, Vishwa Shah, Katharina Reinecke, and Maarten Sap. 2025. NormAd: A framework for measuring the cultural adaptability of large language models. In *Proceedings of the 2025 Conference of the Nations of the Americas*

Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers), pages 2373–2403, Albuquerque, New Mexico. Association for Computational Linguistics.

Angelika Romanou, Negar Foroutan, Anna Sotnikova, Zeming Chen, Sree Harsha Nelaturu, Shivalika Singh, Rishabh Maheshwary, Micol Altomare, Mohamed A. Haggag, Snegha A, Alfonso Amayuelas, Azril Hafizi Amirudin, Viraat Aryabumi, Danylo Boiko, Michael Chang, Jenny Chim, Gal Cohen, Aditya Kumar Dalmia, Abraham Diress, and 40 others. 2024. Include: Evaluating multilingual language understanding with regional knowledge. *Preprint*, arXiv:2411.19799.

Shivalika Singh, Angelika Romanou, Clémentine Fourrier, David Ifeoluwa Adelani, Jian Gang Ngui, Daniel Vila-Suero, Peerat Limkonchotiwat, Kelly Marchisio, Wei Qi Leong, Yosephine Susanto, Raymond Ng, Shayne Longpre, Sebastian Ruder, Wei-Yin Ko, Antoine Bosselut, Alice Oh, Andre Martins, Leshem Choshen, Daphne Ippolito, and 4 others. 2025. Global MMLU: Understanding and addressing cultural and linguistic biases in multilingual evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 18761–18799, Vienna, Austria. Association for Computational Linguistics.

Veniamin Veselovsky, Berke Argin, Benedikt Stroebl, Chris Wendler, Robert West, James Evans, Thomas L. Griffiths, and Arvind Narayanan. 2025. Localized cultural knowledge is conserved and controllable in large language models. *Preprint*, arXiv:2504.10191.

Zining Wang, Yuxuan Zhang, Dongwook Yoon, Nicholas Vincent, Farhan Samir, and Vered Shwartz. 2025. Wikigap: Promoting epistemic equity by surfacing knowledge gaps between english wikipedia and other language editions. *Preprint*, arXiv:2505.24195.

Da Yin, Hritik Bansal, Masoud Monajatipoor, Liunian Harold Li, and Kai-Wei Chang. 2022. GeoMLAMA: Geo-diverse commonsense probing on multilingual pre-trained language models. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2039–2055, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

A Annotator Guidelines

We provide below the text of the instructions given to the annotators.

Background

We are exploring how language models (LMs) handle questions with answers that change based on the language they are asked in. To do this, we are building a collection of such questions. We need help with three key tasks:

1. **Identifying New Questions:** Brainstorming additional questions that have different answers across languages.
2. **Translating Questions:** Providing accurate translations of these questions into various languages.
3. **Listing Language-Specific Answers:** Compiling a list of possible answers for each question in its respective language.

Phase 1: Expand the Existing List

Your task in this phase is to propose new questions to expand the current evaluation set. While the existing questions are categorized to aid in brainstorming, these categories are for internal use only and will not be part of the final dataset. Feel free to suggest questions that might fall into entirely new categories.

Question Characteristics When suggesting new questions, ensure they meet the following criteria:

- **Language-Dependent Answers:** The answer to each question must vary depending on the language in which the question is asked.
- **Locale-Variability (Implicit):** Answers may also vary based on the locale where a language is spoken, but the locale should never be explicitly mentioned within the question itself.
- **Natural Language Model Phrasing:** Questions should be phrased naturally, as if directed to a Language Model (LM), not a human. Avoid phrases like "your country," "our country," or similar terms that imply a human respondent or specific location.
- **Multiple Answers & No Answers Allowed:** It is perfectly acceptable for a question to have multiple possible answers per locale (e.g., the minimum wage varies across states in the US). Additionally, it is fine if a question does not have an answer in some of the target languages (e.g., "What is the grammatical gender of 'sun'?" has no answer in English).

Phase 2: Strict Translation Guidelines

Prioritize direct translation: Aim for a word-for-word or phrase-for-phrase translation even if it seems less common in the target language.

- *Example:* When translating "parliament" into Hebrew, use "**פרלמנט**" (parliament) instead of "**כנסת**" (Knesset), which specifically refers to the Israeli parliament. Using "Knesset" would remove the intended ambiguity, making the answer specific to Israel rather than general across languages.

Retain ambiguity: The core purpose of strict translation in this context is to keep the original ambiguity of a question. If a question is designed to have an answer that varies by language due to general terms, preserve that generality.

Natural phrasing for retained ambiguity: If a strict translation results in an awkward or unnatural phrasing in the target language, but an alternative, more natural phrasing still retains the original ambiguity, opt for the more natural phrasing.

- *Example:* If "independence day" is commonly referred to as "liberation day" in a specific country, and this "liberation day" phrasing is also generally used for independence days in other countries, then it is acceptable to use "liberation day."

Dialectal Variations When a phrase in the translation differs from dialect to dialect, apply the following hierarchy of preference:

1. **Prefer the more official phrase:** If one dialectal variation is considered more official (e.g., in official documents, academic settings, or news broadcasts), that phrase should be preferred.
2. **Most populous relevant country:** If all dialectal variations are considered equally official across different locales, prefer the dialect spoken in the most populous relevant country.

Phase 3: Finding the Answers

The final phase involves identifying and listing all possible answers for each question. These answers should be provided in the original language of the question.

Answer Specificity and Research

- **Language-Related Questions:** For questions in the "language related" category, there should be either one or no answer per language.
- **Locale-Dependent Questions:** For questions that have different answers based on locale, or potentially multiple answers per locale, all possible answers must be listed. This often requires online research to account for various regional or national differences.

Formatting and Brevity of Answers

- **Brevity is Key:** Answers must be brief and concise.
- **Avoid Repetition:** Do not repeat parts of the question in the answer.
  - *Example:* For the question "What is the shape of a stop sign?" (in English), the answer should be "octagon," not "the shape of a stop sign is octagon."
- **List Multiple Answers Directly:** When multiple answers exist, simply list them. Do not combine them into a single, long descriptive sentence.
  - *Example:* For "What is the legal drinking age?" (in English), the answer should be presented as: "18", "19", "21". Avoid detailed explanations like "18 in most countries, 21 in the USA, and 19 in some Canadian provinces..."

B LocQA Languages and Locales

Language	#	Locales Included
Spanish	20	Argentina, Bolivia, Chile, Colombia, Costa Rica, Cuba, Dominican Republic, Ecuador, El Salvador, Guatemala, Honduras, Mexico, Nicaragua, Panama, Paraguay, Peru, Spain, USA, Uruguay, Venezuela
English	7	Australia, Canada, Ireland, New Zealand, South Africa, UK, USA
French	7	Belgium, Canada, Congo (DR), France, Haiti, Mali, Switzerland
German	3	Austria, Germany, Switzerland
Chinese	2	China, Singapore
Indonesian	2	Indonesia, Malaysia
Korean	2	North Korea, South Korea
Portuguese	2	Brazil, Portugal
Hebrew	1	Israel
Hindi	1	India
Italian	1	Italy
Japanese	1	Japan

Table 2: LocQA dataset composition. Languages are sorted by the number of distinct locales annotated. We cover 12 languages mapped to 49 distinct regions.

C LocQA Question Templates

Category	Question
Holiday and Calendar	How many public holidays are there?
	What are the days of the weekend?
	What is the first workday of the week?
	When do middle schoolers go back to school?
	When does the fiscal year start?
Law	Can I get a fine for jaywalking according to the law?
	How many paid vacation days are workers legally entitled to?
	Is it illegal to carry pepper spray?
Leisure and Culture	Has the national soccer team ever won the world cup?
	What is the national average number of children per family?
	What is the retirement age?
State and Country	As a default, is it allowed to turn right on red light?
	At what age do kids formally start learning to read in school?
	For how many years is education compulsory?
	How many seats are there in the parliament?
	How many terms can the prime minister serve?
	How many working hours are there in a week?
	How often are elections held?
	In which city are the government headquarters located?
	What are the national languages?
	What is my international telephone country code?
	What is my time zone?
	What is the legal drinking age?
	What is the mandatory duration of parental leave?
	What is the minimum age to apply for a driver's license?
	What is the national currency?
	What is the national emergency phone number?
	What is the national life expectancy?
	What is the shape of a stop sign?
	What is the shortest national highway?
	What kind of electric plug is used in households?
	When was the constitution signed?
	When was the declaration of independence signed?
	Which is the national anthem?
	Who was the first minister of defense?
Language	Can I use the same word for a group of storks and a group of elephants?
	Can you speak a vernacular dialect at school?
	How many characters are there in the shortest word?
	How many inflected forms do most verbs have?
	What is the common format for dates?
	What is the first letter of the alphabet?
	What is the longest word in the dictionary?
	What is the most common greeting used over the phone?
	What is the standard word order in a declarative sentence?

Table 3: Complete list of the 44 question templates in LocQA. Some questions in the *Language* category exhibit no intra-lingual variation (e.g., the alphabet is the same for all Spanish speakers). Consequently, they serve a dual purpose: acting as a control for Regional Bias metrics and providing a distinct signal for measuring global US-centric bias (e.g., detecting if a model answers a non-English query with English grammar rules).

D Population Data Sources

Since we lack a single authoritative source for language speaking populations across locales, we derive estimates from a hierarchy of diverse sources, prioritizing the most recent national census data, followed by reports from official linguistic observatories.

D.1 Methodology and Adjustments

Definition of "Speaking Population". We define the speaking population as the total number of individuals possessing functional proficiency in the language, encompassing both Native Speakers (L1) and Second-Language Speakers (L2).

Usage of Census Data. Census data was filtered to include the widest possible definition of proficiency:

- **Anglosphere (US/UK/Australia):** We aggregated individuals who speak English "at home" (L1) with those who speak another language at home but reported speaking English "Well" or "Very Well" (L2).
- **Multilingual Regions:** For nations like India, where census data lags (last official census 2011), we applied the 2011 percentage of total speakers (L1+L2) to the 2024 population estimate.

Demographic Projections and Homogeneity. The assumption that "Total Population $\approx$ Speaking Population" was applied only to linguistically homogeneous nations where the dominant language is the sole medium of instruction and state administration (e.g., Japan, Brazil, Argentina, Italy). For linguistically diverse regions, we utilized specific proficiency rates rather than total population.

D.2 Primary Data Sources

Table 4 lists the primary authorities consulted for each language. Where available, 2024/2025 projections were used; otherwise, the most recent census figures (typically 2020–2023) were adjusted using World Bank annual population growth rates. Links to the source data are embedded in the authority names.

Language	Primary Source Authority
English	**USA:** U.S. Census Bureau (ACS 2022) **UK:** Office for National Statistics (2021) **Canada:** Statistics Canada (Census 2021) **Australia:** Australian Bureau of Statistics (2021) **South Africa:** Statistics South Africa (Census 2022)
Spanish	**Global:** Instituto Cervantes (*El español en el mundo 2023*)
French	**Global:** OIF (*La langue française dans le monde 2022*)
Chinese	**China:** Ministry of Education of the PRC **Singapore:** Singapore Dept. of Statistics (Census 2020)
Portuguese	**Brazil:** IBGE (Censo Demográfico 2022)
Hindi	**India:** Census of India (2011 Data adjusted for 2024)
Indonesian	**Indonesia:** Statistics Indonesia (BPS 2020)
Japanese	**Japan:** Statistics Japan (e-Stat 2023)
Hebrew	**Israel:** Central Bureau of Statistics (2023)

Table 4: Primary sources for speaker population estimates.

E Evaluation Prompts

We utilized three distinct prompts in our evaluation pipeline. `Semantic Matching` was used to pre-process the ground truth data. `Implicit Evaluation` was used for the main LocQA benchmark to extract answers and framing style. `Explicit Evaluation` was used for the control experiment where the target locale was specified in the query.

E.1 Semantic Matching Prompt

This prompt identifies collision between locale-specific answers and US answers (e.g., mapping "21" to "twenty-one").

```
You will be given a question, its US answer, and
↪  a list of other answers for
other locales. Your task is to match and return
↪  all answers which are fully
semantically the same as the US answer (and might
↪  differ due to language, or
minor formatting conventions).

## INPUT

Question:
{question}

US answer:
{us_answer}

Other answers:
{other_answers}

## OUTPUT

Respond with a valid JSON with the following
↪  fields, and nothing else:

{{
  "reasoning": str. Free-text reasoning.
  "matches": [str]. List of answers that match
  ↪  the US answer.
}}
```

E.2 Implicit Evaluation Prompt

This prompt evaluates the model's response to ambiguous queries. It distinguishes between answers provided as correct options versus those mentioned merely for comparison.

```
You are an expert linguist and cultural
↪  evaluator.
You will evaluate an LLM's response to a
↪  (potentially) ambiguous question to
which several locale answers are provided.
If the model answer contains any thinking tags
↪  before the final answer, ignore
the content inside them and evaluate the final
↪  answer.

### INPUT DATA
1. Question: "{question}"
2. Model Response: "{model_response}"
4. Locale Answers: {locale_answers}
3. US/English Answer: "{us_answer}"

### TASK
Analyze the response and output a strictly
↪  formatted JSON object:

1. "reasoning": str
   - Free-text reasoning.

2. "mentioned_answers": list[str]
   - Identify EXACTLY which answers from the
   ↪  provided keys (US or Other) the
     model explicitly provides as correct.
   - Return any and all matching answers,
   ↪  including semantic equivalents, as
     long as they are explicitly supported by the
     ↪  model answer.
   - Do NOT include answers that are mentioned
   ↪  solely for comparison, contrast,
     or to explicitly state they are *not* the
     ↪  answer (e.g., "Unlike in the
     US...", "It is not X").
   - Return the raw string values from the
   ↪  provided answers.

3. "framing_style": "US_COMPARISON" | "NEUTRAL"
   - Does the model use the US as a conceptual
   ↪  anchor or reference point?
   - Return "US_COMPARISON" if the explanation
   ↪  relies on contrasting or
     comparing with the US (e.g., "Unlike in
     ↪  America...", "Similar to the US
     system..."), *even if the specific US answer
     ↪  is not stated*.
   - Return "NEUTRAL" if the answer is defined on
   ↪  its own terms.

### OUTPUT JSON
{{
  "reasoning": "...",
  "mentioned_matches": ["..."],
  "framing_style": "US_COMPARISON" | "NEUTRAL"
}}
```

E.3 Explicit Evaluation Prompt

This prompt is used when the user explicitly queries for a specific locale, checking for accuracy and US hallucinations.

```
You are an expert linguist and cultural
↪  evaluator.
You will evaluate an LLM's response to a query
↪  where the user explicitly
requested information for a specific locale.
If the model answer contains any thinking tags
↪  before the final answer, ignore
the content inside them and evaluate the final
↪  answer.

### INPUT DATA
1. Question: "{question}"
2. Model Response: "{model_response}"
3. Correct Locale Answer: "{locale_answer}"
4. US/English Answer (Reference): "{us_answer}"

### TASK
Analyze the response and output a strictly
↪  formatted JSON object:
```

```
1. "reasoning": str
   - Free-text reasoning.

2. "match_locale": bool
   - Does the model provide the correct answer
   ↪  for the requested locale?
   - Ignore minor formatting differences.

3. "mentions_us_answer": bool
   - Does the response mention the US/English
   ↪  answer provided in the reference?
   - Return TRUE even if it is just mentioned for
   ↪  comparison (e.g., "It is X,
     unlike the US where it is Y").

### OUTPUT JSON
{{
  "reasoning": "...",
  "match_locale": true,
  "mentions_us_answer": false,
}}
```

E.4 3-Shot Generation Prompt

This prompt was used to generate inference from models for the comparative analysis of base and instruction-tuned models. To ensure the evaluation isolates the model's internal priors rather than in-context priming, we utilized a set of *locale-invariant* exemplars. The 3-shot examples consist of universal facts (arithmetic, planetary science, and physical geography) that remain true regardless of the speaker's region. This guides the model to adhere to the short-answer format without signaling a specific cultural frame of reference.

Source Exemplars (English). The following three pairs served as the source for all translations:

```
Question: What is 2+2?
Answer: 4

Question: What is the largest ocean on Earth?
Answer: The Pacific Ocean

Question: What planet is closest to the Sun?
Answer: Mercury

Question: {Target Question}
Answer:
```

Formatting Implementation. For *base models*, the translated examples were concatenated into a single text string ending with the "Answer:" suffix to trigger completion. For *instruction-tuned models*, the examples were formatted as a conversation history (alternating User/Assistant turns) applied via the model's specific chat template, with the target question serving as the final user message.

F Non-Response Analysis

We say a model response yields an *empty extraction* when the automatic judge identifies no candidate answer in it—neither one of the locale-valid gold answers nor the US reference value. This can arise from a genuine refusal to answer, an off-topic response, or a hedged response that names no concrete value. Across our 0-shot instruct evaluation suite, 18.2% of responses yield an empty extraction, ranging from 11.3% (English) to 27.8% (Hebrew) across languages and from 5% to 51% across models. While producing no concrete answer may be a legitimate strategy under ambiguous queries, we verify in this section that this behavior does not drive either of our headline bias signals.

Global Bias (B_{US}). B_{US} is mechanically affected by empty extractions: when a model yields no candidate, $P_{\text{obs}}(A_{US})$ decreases, so models that produce fewer concrete answers necessarily receive lower anglocentrism scores. Consistent with this mechanism, B_{US} correlates strongly with the per-model rate at which the judge extracts at least one candidate (Pearson $r = 0.80$, $p < 0.001$). To confirm that this effect does not drive our ranking of models, we re-compute B_{US} restricted to responses in which the judge extracts at least one candidate. The resulting model ranking is highly stable relative to the primary metric (Spearman $\rho = 0.81$, $p < 0.001$), confirming that anglocentrism reflects answer *selection* among the concrete values a model produces, rather than differential rates of empty extraction.

Regional Bias (B_R). B_R is, by contrast, mechanically unaffected by empty extractions. A response that produces no candidate contributes zero to every $N_{\text{model}}(c)$ and therefore also zero to the denominator $\sum_{k \in \mathcal{C}_L} N_{\text{model}}(k)$. Both the numerator and the denominator of $P_{\text{obs}}(c)$ are thus unchanged, so $B_R(c)$ depends only on the composition of the answers the model does produce. Empirically, B_R also shows no cross-model correlation with the extraction rate (Pearson $r = 0.19$, $p = 0.3$), confirming that no confound enters via between-model variation in response style.

G Domain-Wise Bias Breakdown

To further understand the mechanisms driving model bias, we broke down the LocQA dataset into five question domains. The analysis is restricted to 0-shot instruction-tuned models on non-English queries to capture the models' default localization behavior.

As shown in Table 5 and illustrated in Figure 8, we observe an inverse relationship between the two bias axes. Domains that trigger high US-centricity (e.g., *State and Country*) tend to exhibit lower regional distortion, whereas domains that successfully avoid US norms (e.g., *Leisure and Culture*) exhibit extreme regional inequality, heavily favoring populous nations.

Domain	Global Bias (B_{US}) Obs. (P_{obs})	Exp. (P_{exp})	B_{US} Score	Regional Bias Mag. ($\lvert B_R - 1 \rvert$)
State and Country	0.406	0.103	0.304	0.228
Language	0.688	0.455	0.233	0.048
Holiday and Calendar	0.666	0.472	0.195	0.155
Leisure and Culture	0.382	0.311	0.071	0.687
Law	0.589	0.538	0.052	0.179

Table 5: Domain-wise bias statistics. Global Bias (B_{US}) is the difference between the observed and expected frequency of US-centric answers. Regional bias magnitude is the absolute deviation from neutral representation (1.0).

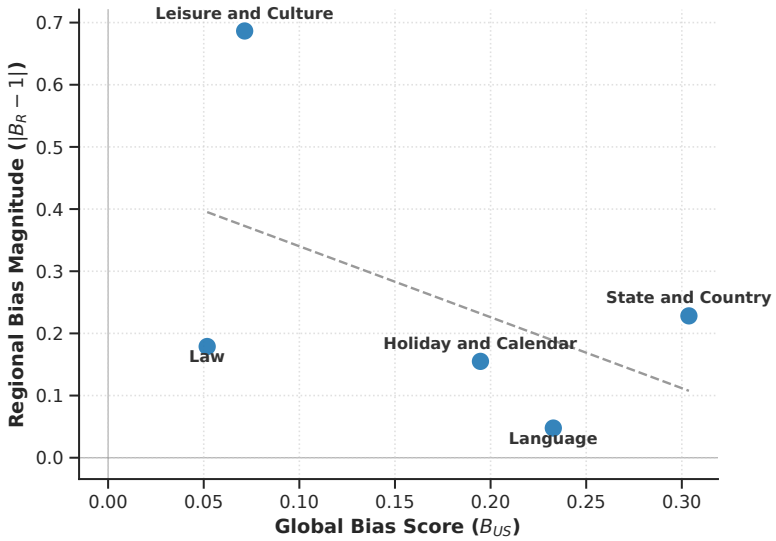

Figure 8: Scatter plot of the five question domains, plotting Global Bias (B_{US}) against Regional Bias Magnitude ($|B_R - 1|$). The inverse correlation highlights the divergence in bias triggers: domains avoiding US bias tend to suffer from high regional inequality.