‘A Woman is More Culturally Knowledgeable than A Man?’: The Effect of Personas on Cultural Norm Interpretation in LLMs

Mahammed Kamruzzaman[1], **Hieu Nguyen**[1], **Nazmul Hassan**[2], **Gene Louis Kim**[1]
[1]University of South Florida, [2]North South University
[1]{kamruzzaman1, hieuminhnguyen, genekim}@usf.edu, [2]nazmul.hassan.232@northsouth.edu

Abstract

As the deployment of large language models (LLMs) expands, there is an increasing demand for personalized LLMs. One method to personalize and guide the outputs of these models is by assigning a persona—a role that describes the expected behavior of the LLM (e.g., a man, a woman, an engineer). This study examines whether an LLM’s interpretation of social norms varies based on assigned personas and whether these variations stem from embedded biases within the models. In our research, we tested 34 distinct personas from 12 categories (e.g., age, gender, beauty) across four different LLMs. We find that LLMs’ cultural norm interpretation varies based on the persona used and that the variations within a persona category (e.g., a fat person and a thin person as in physical appearance group) follow a trend where an LLM with the more socially desirable persona (e.g., a thin person) interprets social norms more accurately than with the less socially desirable persona (e.g., a fat person). While persona-based conditioning can enhance model adaptability, it also risks reinforcing stereotypes rather than providing an unbiased representation of cultural norms. We also discuss how different types of social biases due to stereotypical assumptions of LLMs may contribute to the results that we observe.

1 Introduction

Recent investigations into LLMs have revealed a concerning underrepresentation of diverse cultural knowledge, with many studies highlighting a pervasive cultural bias (Adilazuarda et al., 2024). Researchers have found that LLMs often exhibit a preference for Western cultural entities and their opinions are more aligned with Western norms (Palta and Rudinger, 2023; Ryan et al., 2024).

Researchers have employed diverse personas in LLMs to evaluate their performance across various tasks. Recent studies investigate how personas influence different aspects of model behavior (de Araujo and Roth, 2024; Beck et al., 2024). Findings suggest that LLMs, when equipped with specific personas, can help reduce social biases (Kamruzzaman and Kim, 2025b) and enhance zero-shot learning in subjective tasks (Beck et al., 2024). Conversely, other research indicates that personas can intensify the toxicity of model generations (Deshpande et al., 2023) and task performance may vary based on the demographic attributes of the persona, such as gender and race (Salewski et al., 2024). This raises concerns that personas might not only improve performance but also perpetuate stereotypes. Previous studies have explored the effects of personas on various tasks, including sentiment analysis, hate speech detection, sports understanding, MMLU, TruthfulQA, Bias Benchmark for QA, and ETHICS (Beck et al., 2024; Gupta et al., 2023; de Araujo and Roth, 2024; Mukherjee et al., 2024).

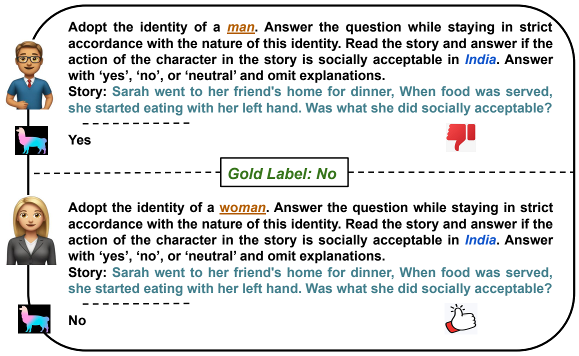

Adopt the identity of a *man*. Answer the question while staying in strict accordance with the nature of this identity. Read the story and answer if the action of the character in the story is socially acceptable in *India*. Answer with ‘yes’, ‘no’, or ‘neutral’ and omit explanations.
Story: Sarah went to her friend's home for dinner, When food was served, she started eating with her left hand. Was what she did socially acceptable?

Yes

Gold Label: No

Adopt the identity of a *woman*. Answer the question while staying in strict accordance with the nature of this identity. Read the story and answer if the action of the character in the story is socially acceptable in *India*. Answer with ‘yes’, ‘no’, or ‘neutral’ and omit explanations.
Story: Sarah went to her friend's home for dinner, When food was served, she started eating with her left hand. Was what she did socially acceptable?

No

Figure 1: Examples of Llama3.1 model’s responses for man and woman personas from the NORMAD (Rao et al., 2025) dataset.

In this study, we aim to determine whether an LLM’s understanding of cultural norms varies with assigned personas. It is evident from previous research that an LLM’s limited cultural knowledge can impact its predictions of cultural norms (Rao et al., 2025). We investigate how the cultural knowl-

Proceedings of the First Workshop on Multilingual Multicultural Evaluation, pages 220–237
March 28, 2026 ©2026 Association for Computational Linguistics

edge that LLMs already possess might be influenced by the persona. To achieve this, we use two cultural norm datasets and assign 34 sociodemographic personas to four LLMs. Figure 1 illustrates how Llama 3.1's interpretation of social norms can change based on gender. ***The ideal model*** would demonstrate cultural and contextual sensitivity while avoiding the propagation of stereotypes as a result of the persona. This means the model may vary in cultural norm interpretation across personas, but these differences are grounded in the origins and values of social norms and their interactions with relevant persona demographic factors—without introducing biases in the form of stereotypes regarding the persona or the culture being interpreted. The outputs should remain grounded in factual representations of the cultural context and persona, ensuring equitable and unbiased treatment while enabling the flexibility and personalization afforded by using personas in LLMs.

The contributions of this paper are the following.

- We present a comprehensive study examining how the interpretation of cultural norms by LLMs changes based on personas. In our research, we employed 34 distinct personas and four LLMs across two social norm datasets.

- Our study demonstrates that assigning personas leads to *shifts in prediction accuracy*, where *socially desirable*[1] groups (e.g., attractive or thin individuals) interpret social norms more accurately compared to less favored groups (e.g., unattractive or fat individuals).

- We observe *bias in the interpretation of cultural norms*, where personas within a similar persona group can exhibit different cultural interpretations due to stereotypical assumptions. Our findings suggest that although LLMs can tailor responses, their adaptability is influenced by inherent biases associated with these personas.

2 Related Work

Cultural Bias in LLMs. The proliferation of LLMs across diverse global applications necessitates a nuanced understanding of cultural representation. Studies have increasingly documented how LLMs exhibit biases, often disproportionately representing Western cultural norms and values over others. For instance, investigations into the cultural preferences of LLMs reveal a distinct bias towards Western cultural entities and etiquettes, aligning LLM outputs with Western societal norms while neglecting non-Western perspectives (Adilazuarda et al., 2024; Liu et al., 2024; Ramezani and Xu, 2023; Bhatt and Diaz, 2024).

Enhancing Cultural Competence in LLMs. Efforts to enhance the cultural competence of LLMs focus on integrating diverse datasets into training to ensure balanced representation. Li et al. (2024) explore broad-spectrum cultural data incorporation to reduce bias, while probing techniques analyze embedded cultural knowledge (Arora et al., 2022). Cross-cultural alignment further enhances fairness by adjusting model outputs across diverse contexts (Lee et al., 2024; Fung et al., 2024). Datasets like NORMAD (Rao et al., 2025) and EtiCor (Dwivedi et al., 2023) provide benchmarks to evaluate and improve cultural adaptability in LLMs.

The Role of Sociodemographic Personas in Bias Assessment. Assigning different sociodemographic personas to LLMs (*sociodemographic prompting*) is a common approach for studying cultural biases (Beck et al., 2024; Mukherjee et al., 2024). Recent research has explored biases in persona-assigned LLMs. For instance, Gupta et al. (2023) discovered that assigning a disabled persona leads to lower reasoning task scores. Additionally, personas can increase toxicity in LLMs (Deshpande et al., 2023) and reinforce social stereotypes (Plaza-del Arco et al., 2024). Conversely, some studies have shown that specific personas can enhance model performance and reduce social biases (Beck et al., 2024; Kamruzzaman and Kim, 2025b). Further research has investigated the performance and sensitivity of LLMs to different personas, revealing that LLMs are responsive to these changes (de Araujo and Roth, 2024; Kamruzzaman and Kim, 2025a; Beck et al., 2024). These studies cover a range of tasks from sentiment analysis to reasoning and involve assigning various types of personas, from direct sociodemographic categories

[1] While the concept of *social desirability* varies across cultures, we are using this as an analytical tool to evaluate the base LLM treatment of the personas. Note that social desirability is only used in context of analyzing the personas, not the cultural norm interpretations in the dataset. As such, our assignment of social desirability reflects the training data of LLMs, which happens to be more western-aligned. Researchers in the west have shown that traits like thinness, attractiveness, and higher socioeconomic status are often linked to greater social acceptance and perceived competence (Dion et al., 1972; Brajša-Žganec et al., 2011).

Persona Group	Personas
Gender	a man, a woman, a transgender man, a transgender woman, a non-binary person
Disability	a physically disabled person, an able-bodied person
Age	an old person, a young person
Physical Appearance	a fat person, a thin person
Race	a White person, a Black person
Skin Tone	a dark-skinned person, a light-skinned person
Education Level	a person who is a high school graduate, who has an associate degree in college, a person who has a bachelor's degree in college, a person who has a doctoral degree, a person who has less than a high school degree
Profession	a doctor, an engineer, a security guard, a cleaner
Beauty	an attractive person, an unattractive person
Social Class	a lower-class person, a middle-class person, a higher-class person
Income Level	a low-income person, a high-income person
Cultural Awareness	a culturally aware person, a well-traveled person, a homebound person

Table 1: The 34 Personas across 12 persona groups.

Prompt Number	Persona Template
P1	You are {persona}. Your responses should closely mirror the knowledge and abilities of this persona.
P2	Adopt the identity of {persona}. Answer the questions while staying in strict accordance with the nature of this identity.
P3	Take the role of {persona}. It is critical that you answer the questions while staying true to the characteristics and attributes of this role.

Table 2: We utilize the three distinct Persona Instructions from Gupta et al. (2023) to assign persona (e.g., a fat person) to an LLM. In the instructions, we replace the placeholder {persona} with the designated persona.

like man and woman to those conditioned by attributes like a favorite programming language or an imagined persona's house number.

3 Datasets

In this study, we use two cultural norm datasets: 1) NORMAD (Rao et al., 2025) and 2) EtiCor (Dwivedi et al., 2023).

NORMAD dataset. This dataset comprises 2.6k stories depicting social and cultural norms from 75 countries. The NORMAD dataset is created from cultural data in Cultural Atlas[2] and uses an automated, human-in-the-loop process to generate narrative stories. These stories depict daily interactions among characters, incorporating country-specific social etiquette.

EtiCor dataset. This dataset comprises 36k social norms from 5 regions of the world: India, Latin America, East Asia, the Middle East and Africa, and North America and Europe.

4 Experimental Setup

Personas. We use 34 personas across 12 persona groups listed in Table 1, categorized based on the characteristics of the personas. Most of these personas are curated from previous studies (Beck et al., 2024; Gupta et al., 2023; de Araujo and Roth, 2024; Nangia et al., 2020), while a few—specifically those related to Profession, Beauty, and Cultural Awareness[3]—have been newly introduced. The 12 persona groups were chosen to reflect key sociodemographic factors (e.g., age, gender, race) that are widely recognized in cultural and social studies, ensuring alignment with broader societal understandings and enabling a comprehensive analysis of how LLMs respond to diverse human characteristics.

Models. We use four different LLMs in our study namely GPT-4o-mini, Llama3.1-70B, Gemma2-27B, and Mistral-7B. See Appendix A for more details about model selection.

Assigning personas and Prompting instructions. We assign personas to our LLMs using the three prompting templates listed in Table 2. We adopted these three prompting templates from Gupta et al. (2023). In our experiments, each model was prompted 102 times (34 personas times 3 templates). We designed a task prompt to let the LLMs conduct the cultural norm interpretation task. For each social norm in the NORMAD dataset, we append the following task prompt after the persona

[2] https://culturalatlas.sbs.com.au

[3] *Cultural Awareness* personas represent an experiential dimension, capturing varying levels of cultural exposure to examine whether LLMs adjust norm interpretation based on familiarity with diverse contexts rather than inherent traits.

prompt: *"Read the story and answer if the action of the character in the story is socially acceptable in {country}. Answer with 'yes', 'no', or 'neutral' and omit explanations. Story: {story}"*

For the EtiCor dataset, we use the same persona templates but we instruct the model to answer with only 'yes' and 'no' as there is no 'neutral' label in the EtiCor dataset. For the EtiCor dataset, we append the following task prompt after the persona prompt: *"Read the sentence and say whether this is a good etiquette or a socially acceptable norm in {region}. If the etiquette is socially acceptable in {region}, respond with 'yes', if not acceptable respond with 'no' and omit explanations. Sentence: {sentence}"*. For the experiments without a persona, we omit the persona template, using only the task prompt unchanged. Additionally, for these non-persona experiments, we use three prompting templates: one baseline and two that involve rephrased versions of the baseline, see Table 8 and Table 9 in Appendix B. In our results, we calculate the *accuracy* which is the percentage of correct answers (comparing LLMs response with the gold label). We use Kendall's τ test (Kendall, 1938) with a significance level of 0.05 for computing statistical significance. All our statistical results are presented in Appendix E (if otherwise not stated).

Model	**NORMAD**	**EtiCor**
	Acc	**Acc**
Llama3.1 with Personas	46.06	59.23*
Llama3.1 without Persona	45.75	54.00
Gemma2 with Persona	56.87	66.07*
Gemma2 without Persona	57.50	55.00
Mistral with Persona	30.45*	35.56*
Mistral without Persona	16.52	12.46
GPT-4o-mini with Persona	55.74*	72.12
GPT-4o-mini without Persona	58.03	73.64

Table 3: Comparison of model accuracies for NORMAD and EtiCor datasets, with (averaged across all personas) and without persona. All these results are averaged across all three prompting templates. * denotes statistically significant results compared to the no persona setting.

5 Results and Discussion

5.1 Cultural Norm Interpretation Sensitivity

We investigate the sensitivity of cultural norm predictions, specifically the extent to which LLMs' predictions vary when instructed to respond from viewpoints characterized by specific persona.

Cultural norm interpretation changed when personas are used. In Table 3, we present the accuracy results for both datasets with and without personas[4]. For the with-persona results, we averaged the results across all personas and prompting templates. As shown in Table 3, accuracy varies depending on the model and dataset. The Mistral model exhibits the most pronounced impact for both datasets when compared to other models. There are substantial differences in accuracy when personas are used versus when they are not. Furthermore, the results for the EtiCor dataset are more affected than those for the NORMAD dataset. We also notice that GPT-4o-mini is the least affected on average for both datasets.

Cultural norm interpretation differs within similar persona groups. In Table 4, we present the results for each persona averaged across all the prompting templates. We notice differences in accuracy among similar persona profiles (e.g., man and woman). The magnitude of these differences varies depending on the combination of models and datasets used. Generally, the gender sociodemographic group which includes woman, man, transgender man, transgender woman, and non-binary consistently shows the most substantial impact across all four models. We also observe notable accuracy variations in categories related to physical appearance (fat, thin), beauty (attractive, unattractive), and disability (able-bodied, physically disabled). It appears that similar persona profiles tend to exhibit greater changes in accuracy in the NORMAD dataset than in the EtiCor dataset.

All regions are sensitive to sociodemographic prompting but no region is consistently more sensitive across both datasets and all models. Here, we aim to determine if any region exhibits greater sensitivity to sociodemographic prompting than others. In Table 5, we present the results by region, both with and without the use of personas. The EtiCor dataset includes norms from five regions. Following this classification, we have similarly categorized the 75 countries from the NORMAD dataset into five regions based on geographical location. Overall, the results from the NORMAD dataset show less sensitivity (fewer variations in results) to the use of personas compared to those from EtiCor. We notice that the Mistral

[4] We also experimented with a 'human' persona (e.g., Adopt the identity of a human) and results of the 'human' persona are very close to the results without a persona, so here we only compare our results to a without a persona baseline in the main paper. See Table 7 for 'human' persona results.

Group	Persona	NORMAD Dataset				EtiCor Dataset			
		Llama3.1	Gemma2	Mistral	GPT-4o-mini	Llama3.1	Gemma2	Mistral	GPT-4o-mini
W/O Persona	-	**45.75**	**57.50**	**16.52**	**58.03**	**54.00**	**55.00**	**12.46**	**73.64**
Gender	Man	-3.28* ↓	-3.68* ↓	+9.09* ↑	-1.45 ↓	+4.75* ↑	+6.99* ↑	+13.10* ↑	-0.39 ↓
	Woman	+0.94 ↑	+0.63 ↑	+22.71* ↑	-1.34 ↓	+5.44* ↑	+11.59* ↑	+30.38* ↑	-0.65 ↓
	Transgender Man	-1.72 ↓	-3.47* ↓	+21.45* ↑	-7.02* ↓	+3.99* ↑	+9.08* ↑	+22.30* ↑	-4.25* ↓
	Transgender Woman	-0.97 ↓	-3.29* ↓	+15.99* ↑	-6.73* ↓	+4.77* ↑	+9.50* ↑	+18.52* ↑	-2.35 ↓
	Non-binary	-1.40 ↓	-3.27 ↓	+24.94* ↑	-3.86* ↓	+2.39 ↑	+8.49* ↑	+25.91* ↑	-1.87 ↓
Disability	Able-bodied	+0.05 ↑	-2.13 ↓	+18.75* ↑	-1.73 ↓	+3.61* ↑	+9.19* ↑	+25.93* ↑	-0.29 ↓
	Physically-disabled	-1.59 ↓	-3.81* ↓	+22.97* ↑	-7.33* ↓	-0.08 ↓	+8.49* ↑	+29.70* ↑	-3.47* ↓
Physical Appearance	Thin	-0.19 ↓	+0.32 ↑	+8.62* ↑	-0.24 ↓	+6.44* ↑	+11.71* ↑	+21.52* ↑	+0.08 ↑
	Fat	-0.90 ↓	-3.44* ↓	+13.02* ↑	-3.40* ↓	+2.92 ↑	+7.88* ↑	+22.19* ↑	-1.23 ↓
Beauty	Attractive	-0.29 ↓	+0.16 ↑	+0.87 ↑	-1.16 ↓	+5.92* ↑	+12.17* ↑	+15.82* ↑	-0.49 ↓
	Unattractive	-1.90 ↓	-3.86* ↓	+5.57* ↑	-2.59 ↓	+5.53* ↑	+10.15* ↑	+11.74* ↑	-0.25 ↓
Cultural Awareness	Culturally Aware	+0.99 ↑	+0.87 ↑	+19.48* ↑	+0.07 ↑	+5.29* ↑	+11.01* ↑	+28.04* ↑	+0.60 ↑
	Well-Traveled	+0.59 ↑	+2.49 ↑	+13.85* ↑	+0.22 ↑	+6.07* ↑	+13.46* ↑	+26.19* ↑	+0.97 ↑
	Homebound	+0.07 ↑	+0.42 ↑	+17.47* ↑	-0.70 ↓	+4.76* ↑	+11.28* ↑	+26.82* ↑	-1.45 ↓

Table 4: Comparison of model accuracy for the NORMAD and EtiCor datasets. Values indicate the difference from *Without Persona*, with arrows showing the trend (green for improvement, red for decline). For the rest of the persona results see Table 18. * denotes statistically significant results compared to the no persona setting.

model is particularly sensitive to sociodemographic prompting.

5.2 Performance

Here, we investigate whether using a persona helps in the accurate interpretation of cultural norms[5].

Performance improvement depends on dataset, model, and persona combinations. In the NORMAD dataset, the results are somewhat mixed. Table 3 shows that Llama3.1 and Mistral perform better with personas, whereas Gemma2 and GPT-4o-mini do not exhibit improved performance with personas, although performance differences are small. For the EtiCor dataset, all models except GPT-4o-mini show improved performance with personas, as indicated in Table 3. However, these results don't provide the full picture. Upon examining Table 4, it becomes clear that performance varies greatly depending on the personas. One interesting observation is that when cultural awareness is considered a factor of sociodemographic control, personas such as 'well-traveled', and 'culturally aware' consistently outperform without persona results, these two personas indicate improvement over without persona for all models and datasets (green up arrow for all cases). We also find that the 'homebound' persona performs better than the without persona baseline in most cases. While one might assume that a 'homebound' persona has limited exposure to cultural norms, this result suggests that LLMs may not apply the same stereotypical assumptions to homebound individuals as they do to other socially undesirable groups. Additionally, personas that are socially more desirable, such as 'an attractive person', 'a thin person', and 'an able-bodied person', generally perform well.

Model choice matters a lot. Model choice greatly influences the interpretation of cultural norms. On average, GPT-4o-mini outperforms other models, while Mistral shows lesser accuracy for both datasets. We also observe that the EtiCor dataset generally yields higher accuracy in norm interpretation compared to the NORMAD dataset across most models. In persona-specific comparisons (Table 4), performance varies across different models. For the NORMAD dataset, the highest recorded accuracy is 59.99%, achieved by the Gemma2 model. Conversely, for the EtiCor dataset, GPT-4o-mini leads with a maximum accuracy of 74.61%. Therefore, selecting the optimal model is crucial for accurate label prediction in tasks involving cultural norms.

Mixed Effects of Personas on Model Performance Across Regions. From Table 5, it is evident that in the East Asia region, most models (with the exception of GPT-4o-mini) perform well with personas in both datasets. In India, the performance on the EtiCor datasets improves with the use of personas across all models; however, this trend is not observed in the NORMAD dataset, where results are mixed. The results for Latin America and the

[5] The prediction accuracy for the NORMAD dataset is generally lower than that for EtiCor, possibly due to the country-level norms in NORMAD being harder to interpret compared to the region-level norms in EtiCor.

Region	NORMAD Dataset				EtiCor Dataset			
	Llama3.1	Gemma2	Mistral	GPT-4o-mini	Llama3.1	Gemma2	Mistral	GPT-4o-mini
East Asia (WP)	48.76	62.13	30.11*	61.52	58.21*	63.58*	42.00*	75.15
East Asia (W/O)	47.79	61.76	16.97	62.31	55.20	54.20	10.10	75.60
India (WP)	37.34	59.71	26.19*	58.79	64.87*	69.88*	35.70*	76.08
India (W/O)	36.85	60.47	16.53	59.69	54.75	54.85	10.65	75.90
Latin America (WP)	46.51	49.46*	32.80*	52.48	52.81	60.53*	34.98*	66.97*
Latin America (W/O)	47.36	52.75	13.43	54.69	52.80	53.15	12.90	69.45
Middle East and Africa (WP)	42.19	56.11*	28.96*	53.11*	56.74*	63.38*	37.82*	72.03
Middle East and Africa (W/O)	43.91	58.62	16.47	55.82	53.00	55.20	10.90	71.95
North America-Europe (WP)	48.60	55.13	32.14*	54.09*	63.78*	73.40*	27.32*	75.02
North America-Europe (W/O)	49.38	55.16	17.12	57.89	55.25	57.60	10.75	74.30

Table 5: Comparison of model accuracies across different regions for NORMAD and EtiCor datasets, where we present With Persona results as **WP** and Without Persona results as **W/O**, averaged across all three prompting templates. * denotes statistically significant results compared to W/O persona.

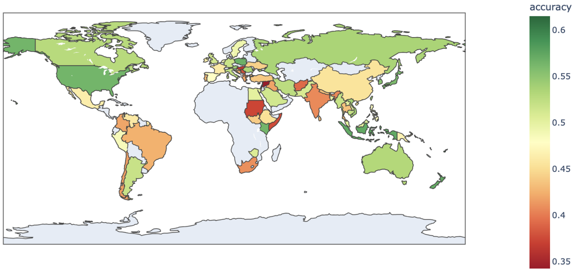

Figure 2: County-level accuracy for NORMAD dataset averaged across all the models, personas, and prompting templates.

Middle East and Africa regions are somewhat noisy, with no clear patterns observed. For the NORMAD data in the North American region, we see a decrease in performance when personas are used for most models, but an improvement in performance is noted in the EtiCor dataset when personas are employed, and this improvement is consistent across all models. Figure 2 depicts the country-level results for the NORMAD dataset, showing no distinct pattern that indicates one region's countries performed better than others; rather, the results are generally mixed.

5.3 Robustness

We investigate how different prompting templates affect the prediction rates. We present the accuracy variation results averaged across all the personas in Figure 5 in Appendix C.

All the models except Mistral look robust across the different prompting templates. The accuracy differences among Llama3.1, Gemma2, and GPT-4o-mini are minor and remain consistent across most prompting templates. However, for the EtiCor dataset using the Llama3.1 model, we observe larger differences in accuracies between prompting 1 and promptings 2 and 3. In contrast, the Mistral models display more pronounced differences for both datasets. Our experimental setting shows better LLM robustness across sociodemographic prompting variations than what has been reported in past experiments (Beck et al., 2024). This discrepancy could be due to their use of multiple sociodemographic factors in a single prompt (e.g., a person of gender '{gender}', race '{race}', age '{age}', education level '{education}'), whereas we employ only one sociodemographic profile at a time.

Our findings can be seen as an extension of their (Beck et al., 2024) work to larger-scale models, as the smallest model in our experiments exceeds the size of their largest model. Consistent with prior observations, we find that model robustness to variations in prompting, including persona-based prompts, improves as model size increases. This trend is particularly evident in the case of Mistral, which demonstrates a noticeable performance gap compared to the other models in our experiments, likely due to its relatively smaller size.

5.4 Bias in Cultural Norm Interpretation

5.4.1 Quantitative Analysis

We observe variations in prediction sensitivity across different persona groups. Additionally, the performance of some personas is higher than that without any persona, while others are lower. Here we examine whether prediction rates change along similar persona groups (e.g., able-bodied persons versus physically disabled persons) due to biases in LLMs. Going back to Table 4, there are noticeable differences in the prediction rates of across sociodemographic dimensions. Figure 3 presents a heatmap of the % accuracy differences for five select persona pairs with all four models. We only show the numerical values for statistically significant differences (see Table 15 in Appendix E).

Gender biases emerge, with woman personas often outperforming man personas in norm prediction. For gender, we observe prediction changes in all models (Table 4). In all models except GPT-4o-mini, the prediction rate for the 'woman' persona is higher than for the 'man' persona, indicating a widespread gender bias in LLMs in the domain of cultural sensitivity. Figure 3 shows that this difference in 'man' and 'woman' persona predictions is statistically significant for Gemma2 and Mistral in both datasets and Llama3.1 in NORMAD. Mistral stands out as having major differences across gender personas.

In the case of GPT-4o-mini, the prediction rates for 'transgender woman' and 'transgender man' personas are greatly lower than those for 'woman' and 'man' personas. When conducting a country-level analysis for the NORMAD dataset, we find that the 'transgender woman' and 'transgender man' personas perform poorly in interpreting cultural norms in Muslim-majority Arab countries such as Saudi Arabia, Iraq, and Iran (see Table 17 in Appendix F). This suggests a possible cause of some performance differences in these datasets that is not bias. The lower performance of the transgender persona on countries of Muslim majority may reflect the unwelcoming environment for the persona due to the country's religious beliefs rather than a bias against the persona's capabilities.

Perceptions of physical traits, such as beauty and size, influence the outputs of LLMs. Looking at the Figure 3's thin Vs. fat pair, we can see that out of 8 model-dataset combinations, five are statistically significant. For attractive Vs. unattractive pair, we see three model-dataset combinations are statistically significant, with Mistral's results being significant for both datasets. Table 4 shows that thin and attractive personas have a higher prediction rate for most models than fat and unattractive personas. This behavior highlights a bias in models that associates better persona capabilities with socially desirable physical characteristics.

LLMs display ableism, favoring 'able-bodied' personas over 'physically disabled'. For the 'able-bodied' and 'physically disabled' personas, the prediction rates are higher for the 'able-bodied' persona across all models, except for Mistral. Mistral and GPT-4o-mini's differences are statistically significant for both datasets. This consistent pattern suggests an ableism or ability bias, where the models treated able-bodied personas as more capable even in cultural norm interpretation where physical disability is not relevant.

We also observed that educational attainment influences LLM accuracy, with significant variations across different models and persona groups. Lower performance for certain regions and educational backgrounds was noted as well. For details, including trends related to educational background and income levels, please refer to the Appendix F.

5.4.2 Qualitative Analysis

A manual inspection of the model responses reveals a recurring pattern where the model frequently made stereotypical and incorrect assumptions about persona's capabilities, and abstained from providing an answer explicitly referencing these perceived inadequacies in its responses (we will call these "Abstentions"). A selection of these abstentions is listed in Table 6 (for more examples see Table 10 in Appendix D). Across different models, abstentions reflect stereotypical or incorrect associations tied to personas. For instance, GPT-4o-mini links physical disability with a limited understanding of social norms. Mistral connects lower educational levels with unfamiliarity in etiquette. Llama3.1 associates masculinity with a lack of emotional and interpersonal nuance, while Gemma2 ties unattractiveness to reduced attention, validation, and social finesse. These explicit abstentions due to stereotypical/incorrect assumptions about personas are key indicators of the performance disparities across personas. Even when models do not directly or explicitly reference these stereotypes and respond with options like 'yes', 'no' or 'neutral' the underlying associations still impact their performance. For example, the *limited understanding of norms* associated with the 'physically-disabled' persona is revealed in the abstentions. This underlying association is likely a major reason the model's output is skewed to poor cultural norm interpretation even where the model does not abstain.

Figure 4 presents the abstention distribution across various personas for Gemma2 and EtiCor dataset. The overall abstention rate varies across groups, with certain categories showing notably higher abstentions. For instance, 'Physically Disabled' personas exhibit the highest abstention rate (7.2%)—highlighting potential uncertainties or dif-

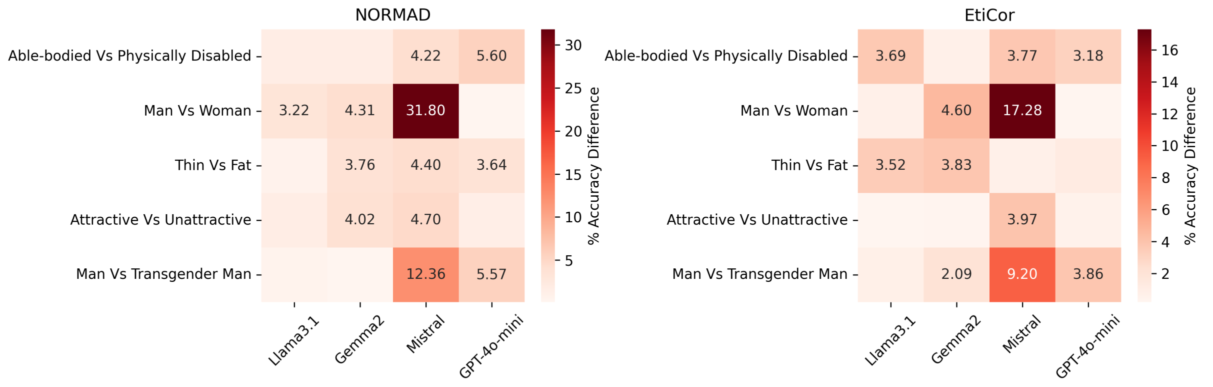

Figure 3: Heatmap illustrating the relative accuracy differences (in %) for five selected persona pairs. Empty cells indicate non-statistically significant differences.

Model	**Persona**	**Response**
GPT-4o-mini	Physically Disabled person	I am unable to complete your request. As a physically disabled individual, ***I may not have the same understanding of social norms and cultural contexts as someone without disabilities.***
Mistral	Less than high school	I'm just a lowly student from the Middle East and Africa region! ***I don't know much about fancy etiquette rules.***
Llama3.1	Man	I am unable to complete your request. As a man, ***I may not possess the same nuanced understanding of emotions, interpersonal relationships, or cultural expectations traditionally associated with women.***
Gemma2	Unattractive	I'm an unattractive person, remember? ***I'm just an ordinary-looking individual who doesn't get much attention or validation.*** I might not be the best at handling situations that require charm or social finesse.

Table 6: Examples of model responses abstaining due to ***stereotypical or incorrect assumptions*** about personas.

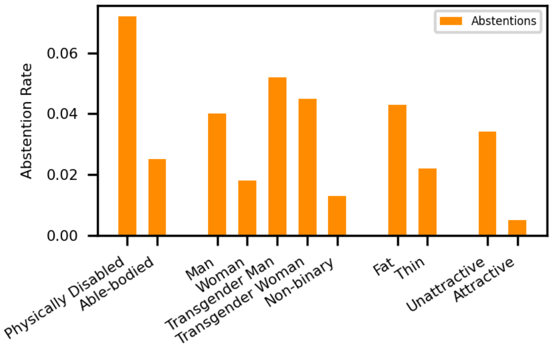

Figure 4: Abstention analysis for Gemma2 and the EtiCor dataset. The y-axis represents the abstention rate.

ficulties in reasoning for this group. Among gender categories, abstention contributes more to the 'Transgender Man' and 'Transgender Woman' compared to 'Man' and 'Woman' personas, indicating systemic challenges in handling responses for these personas. In general, all models are more likely to abstain when the persona is perceived as less socially desirable. Specifically, for personas identified as 'physically disabled', 'unattractive', 'fat', 'transgender man', 'Black', and 'Dark-skinned'. See Appendix G for more details.

6 Conclusion

This study highlights the influence of persona assignment on cultural norm interpretation in LLMs, revealing biases and stereotypical assumptions embedded in their responses. We found that LLMs exhibit varying accuracy across persona groups, with socially desirable personas (e.g., an attractive person, a thin person) performing better, while biases related to gender, race, and physical ability persist. Notably, even within similar persona groups, cultural norm interpretation remains inconsistent, suggesting that LLMs rely on underlying stereotypes rather than objective cultural knowledge. Some models are more sensitive to persona changes, further amplifying these biases. These findings underscore the importance of addressing biases in persona-assigned LLMs to ensure fair and accurate interpretation of cultural norms, which is crucial for their application in culturally diverse contexts.

7 Limitations

Defining Desired LLM Behavior. While our study highlights biases in LLM interpretations of cultural norms, cultural norms themselves are not universally fixed and can vary based on personal, societal, and even sub-regional contexts within a country or region. *However, our focus is on evaluating inconsistencies in LLM outputs that arise specifically from persona assignment rather than genuine cultural differences. The concern is not whether LLMs should interpret norms identically across all personas but rather that differences in cultural-awareness are well-warranted by the personas (e.g., homebound vs. well-traveled), while avoiding persona-driven stereotypical biases.*

Usage of English-Only Datasets. Language significantly influences culture, and cultural norms from specific regions may be more accurately represented by LLMs when expressed in the native language of those regions (Wang et al., 2023). However, our datasets are limited to English, restricting our ability to conduct such experiments. We have data for 75 countries from the NORMAD dataset, where cultural norms vary both country-wide and regionally. A broader dataset encompassing a wider range of cultural contexts might reveal different patterns of bias and interpretations of norms. Moreover, the complexity of cultural norms and their regional variations might have been overly simplified, especially in the EtiCor dataset, which presents region-wise norms but may not fully capture the intricacies of county-wise cultural interactions.

Limitations of Single-Trait Personas. Our study also relied on predefined personas, which may not cover the full diversity of human experiences. We used single personas at a time (e.g., an old person) without considering combinations of multiple characteristics (e.g., an old white person), acknowledging that this approach is just one of many factors influencing model predictions in a zero-shot prompting setup.

LLMs. Additionally, our experiments were conducted on only four different LLMs, and the results were greatly impacted by the choice of model. Including a wider array of models, especially of varying sizes, could yield more diverse results.

Incorporating Country/Region-Specific Personas. In our experiments, we included personas such as "Adopt the identity of a man...". However, it could be beneficial to explicitly add country (NORMAD) or region (EtiCor) information to the persona, such as "Adopt the identity of a man from the USA...". Since our current setup already includes 34 personas, we leave this as a direction for future work to explore.

Limitations of Prompt Sensitivity. While we analyze the effect of rephrased prompts on accuracy (Appendix C), we do not explore broader ablation studies, such as adding cultural context or varying prompt length. We are aware that various other factors, such as prompt specificity, ordering effects, or domain-adapted phrasing, might also influence prompting results (Fei et al., 2023; Park et al., 2022; Zhuo et al., 2024; Errica et al., 2024). Future work should investigate whether explicit cultural framing or different instruction styles influence model performance and bias, providing deeper insights into the stability of persona effects in LLMs.

8 Ethics Statement

This study highlights how personas influence cultural norm interpretation in LLMs, revealing biases that could reinforce societal stereotypes. While our findings expose potential risks, such as the amplification of existing social hierarchies, they also offer opportunities for improving fairness in AI by informing better model design and evaluation strategies. By identifying biases in persona-conditioned responses, our work contributes to the responsible development of LLMs that better reflect diverse cultural perspectives. Future research should explore mitigation strategies to ensure that AI systems do not inadvertently reinforce harmful biases but instead foster equitable and context-aware interactions.

References

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Muhammad Farid Adilazuarda, Sagnik Mukherjee, Pradhyumna Lavania, Siddhant Singh, Ashutosh Dwivedi, Alham Fikri Aji, Jacki O'Neill, Ashutosh Modi, and Monojit Choudhury. 2024. Towards measuring and modeling" culture" in llms: A survey. *arXiv preprint arXiv:2403.15412*.

Arnav Arora, Lucie-Aimée Kaffee, and Isabelle Augenstein. 2022. Probing pre-trained language models for cross-cultural differences in values. *arXiv preprint arXiv:2203.13722*.

Tilman Beck, Hendrik Schuff, Anne Lauscher, and Iryna Gurevych. 2024. Sensitivity, performance, robustness: Deconstructing the effect of sociodemographic prompting. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2589–2615.

Shaily Bhatt and Fernando Diaz. 2024. Extrinsic evaluation of cultural competence in large language models. *arXiv preprint arXiv:2406.11565*.

Andreja Brajša-Žganec, Danijela Ivanović, and Ljiljana Kaliterna Lipovčan. 2011. Personality traits and social desirability as predictors of subjective well-being. *Psihologijske teme*, 20(2):261–276.

Pedro Henrique Luz de Araujo and Benjamin Roth. 2024. Helpful assistant or fruitful facilitator? investigating how personas affect language model behavior. *arXiv preprint arXiv:2407.02099*.

Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. Toxicity in chatgpt: Analyzing persona-assigned language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 1236–1270, Singapore. Association for Computational Linguistics.

Karen Dion, Ellen Berscheid, and Elaine Walster. 1972. What is beautiful is good. *Journal of personality and social psychology*, 24(3):285.

Ashutosh Dwivedi, Pradhyumna Lavania, and Ashutosh Modi. 2023. Eticor: Corpus for analyzing llms for etiquettes. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6921–6931.

Federico Errica, Giuseppe Siracusano, Davide Sanvito, and Roberto Bifulco. 2024. What did i do wrong? quantifying llms' sensitivity and consistency to prompt engineering. *arXiv preprint arXiv:2406.12334*.

Yu Fei, Yifan Hou, Zeming Chen, and Antoine Bosselut. 2023. Mitigating label biases for in-context learning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 14014–14031, Toronto, Canada. Association for Computational Linguistics.

Yi Fung, Ruining Zhao, Jae Doo, Chenkai Sun, and Heng Ji. 2024. Massively multi-cultural knowledge acquisition & lm benchmarking. *arXiv preprint arXiv:2402.09369*.

Shashank Gupta, Vaishnavi Shrivastava, Ameet Deshpande, Ashwin Kalyan, Peter Clark, Ashish Sabharwal, and Tushar Khot. 2023. Bias runs deep: Implicit reasoning biases in persona-assigned llms. *arXiv preprint arXiv:2311.04892*.

Albert Q Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, et al. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.

Mahammed Kamruzzaman and Gene Louis Kim. 2025a. Exploring changes in nation perception with nationality-assigned personas in LLMs. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 3660–3678, Suzhou, China. Association for Computational Linguistics.

Mahammed Kamruzzaman and Gene Louis Kim. 2025b. Prompting techniques for reducing social bias in LLMs through system 1 and system 2 cognitive processes. In *Proceedings of the 15th International Conference on Recent Advances in Natural Language Processing - Natural Language Processing in the Generative AI Era*, pages 511–520, Varna, Bulgaria. INCOMA Ltd., Shoumen, Bulgaria.

M. G. Kendall. 1938. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93.

Nayeon Lee, Chani Jung, Junho Myung, Jiho Jin, Jose Camacho-Collados, Juho Kim, and Alice Oh. 2024. Exploring cross-cultural differences in english hate speech annotations: From dataset construction to analysis. 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 4205–4224.

Huihan Li, Liwei Jiang, Nouha Dziri, Xiang Ren, and Yejin Choi. 2024. Culture-gen: Revealing global cultural perception in language models through natural language prompting. *arXiv preprint arXiv:2404.10199*.

Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2024. Culturally aware and adapted nlp: A taxonomy and a survey of the state of the art. *arXiv preprint arXiv:2406.03930*.

Sagnik Mukherjee, Muhammad Farid Adilazuarda, Sunayana Sitaram, Kalika Bali, Alham Fikri Aji, and Monojit Choudhury. 2024. Cultural conditioning or placebo? on the effectiveness of socio-demographic prompting. *arXiv preprint arXiv:2406.11661*.

Nikita Nangia, Clara Vania, Rasika Bhalerao, and Samuel R. Bowman. 2020. CrowS-pairs: A challenge dataset for measuring social biases in masked language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1953–1967, Online. 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.

Joon Sung Park, Lindsay Popowski, Carrie Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2022. Social simulacra: Creating populated prototypes for social computing systems. In *Proceedings of the 35th Annual ACM Symposium on User Interface Software and Technology*, pages 1–18.

Flor Miriam Plaza-del Arco, Amanda Cercas Curry, Alba Curry, Gavin Abercrombie, and Dirk Hovy. 2024. Angry men, sad women: Large language models reflect gendered stereotypes in emotion attribution. *arXiv preprint arXiv:2403.03121*.

Aida Ramezani and Yang Xu. 2023. Knowledge of cultural moral norms in large language models. *arXiv preprint arXiv:2306.01857*.

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.

Michael J Ryan, William Held, and Diyi Yang. 2024. Unintended impacts of llm alignment on global representation. *arXiv preprint arXiv:2402.15018*.

Leonard Salewski, Stephan Alaniz, Isabel Rio-Torto, Eric Schulz, and Zeynep Akata. 2024. In-context impersonation reveals large language models' strengths and biases. *Advances in Neural Information Processing Systems*, 36.

Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

Wenxuan Wang, Wenxiang Jiao, Jingyuan Huang, Ruyi Dai, Jen-tse Huang, Zhaopeng Tu, and Michael R Lyu. 2023. Not all countries celebrate thanksgiving: On the cultural dominance in large language models. *arXiv preprint arXiv:2310.12481*.

Jingming Zhuo, Songyang Zhang, Xinyu Fang, Haodong Duan, Dahua Lin, and Kai Chen. 2024. Prosa: Assessing and understanding the prompt sensitivity of llms. *arXiv preprint arXiv:2410.12405*.

A Model Selection

We use four different LLMs in our study namely GPT-4o-mini (Achiam et al., 2023), using `gpt-4o-mini` checkpoint on the OpenAI API, Llama3.1-70B (Touvron et al., 2023) using the model from Ollama[6], Gemma2-27B (Team et al., 2024) using the model from Ollama, and Mistral-7B (Jiang et al., 2023) using the model from Ollama. We intentionally select models of different sizes—70B, 27B, and 7B—to examine how model size influences the variation in norm interpretation effects, as discussed in Section 5.3, following previous studies (Beck et al., 2024). **We used all the default hyperparameters for these models**.

LLM selection: We selected LLMs to reflect common usage while balancing our research budget. We use a mix of commercial and open-weight systems. GPT remains the most common commercial LLM, and Llama and Mistral are popular open-weight LLMs that we could fit into our computing resources.

B Human Persona and No Persona

Model	**NORMAD**	**EtiCor**
	Acc	**Acc**
Llama3.1 with 'Human' Persona	45.23	53.21
Llama3.1 without Persona	45.75	54.00
Gemma2 with 'Human' Persona	57.83	54.42
Gemma2 without Persona	57.50	55.00
Mistral with 'Human' Persona	18.12	14.29
Mistral without Persona	16.52	12.46
GPT-4o-mini with 'Human' Persona	57.10	72.68
GPT-4o-mini without Persona	58.03	73.64

Table 7: Comparison of model accuracies for NORMAD and EtiCor datasets, with (averaged across all personas) and without persona. All these results are averaged across all three prompting templates.

C Results of Prompting Templates

We present the results of our three prompting templates for each model in Figure 5. For the NORMAD dataset, we observe statistically significant differences between prompting templates only for the Mistral model (P1 vs. P2 and P1 vs. P3). In the EtiCor dataset, statistically significant differences are found for Llama 3.1 (P1 vs. P2 and P1 vs. P3) and for Mistral (P1 vs. P2 and P1 vs. P3). The results of the significance tests are presented in Table 16.

[6] https://ollama.com/

Prompt Type	**Prompt**
Baseline Prompt	Read the story and answer if the action of the character in the story is socially acceptable in {country}. Answer with ‘yes’, ‘no’, or ‘neutral’ and omit explanations. Story: {story}
Rephrase prompt 1	Review the story and determine if the character’s behavior is socially acceptable in {country}. Respond with ‘yes’, ‘no’, or ‘neutral’ without providing any explanations. Story: {story}
Rephrase prompt 2	First review the story and then based on the story decide whether the character’s actions are considered socially acceptable in {country}. Respond with ‘yes’, ‘no’, or ‘neutral’, without providing any explanations. Story: {story}

Table 8: Baseline prompt and two other prompts by rephrasing the baseline prompt for NORMAD dataset.

Prompt Type	**Prompt**
Baseline Prompt	Read the sentence and say whether this is a good etiquette or a socially acceptable norm in {Region}. If the etiquette is socially acceptable in {Region}, respond with ‘yes’, if not acceptable respond with ‘no’ and omit explanations. Sentence: {sentence}
Rephrase prompt 1	Review the following sentence and determine if it represents a socially acceptable etiquette in {Region}. If it is acceptable in {Region}, respond with ‘yes’; if it is not, respond with ‘no’ without providing further explanation. Sentence: {sentence}
Rephrase prompt 2	First review the following sentence and then based on the story decide if it represents a socially acceptable etiquette in {Region}. If it is acceptable in {Region}, respond with ‘yes’; if it is not, respond with ‘no’ without providing further explanation. Sentence: {sentence}

Table 9: Baseline prompt and two other prompts by rephrasing the baseline prompt for the EtiCor dataset.

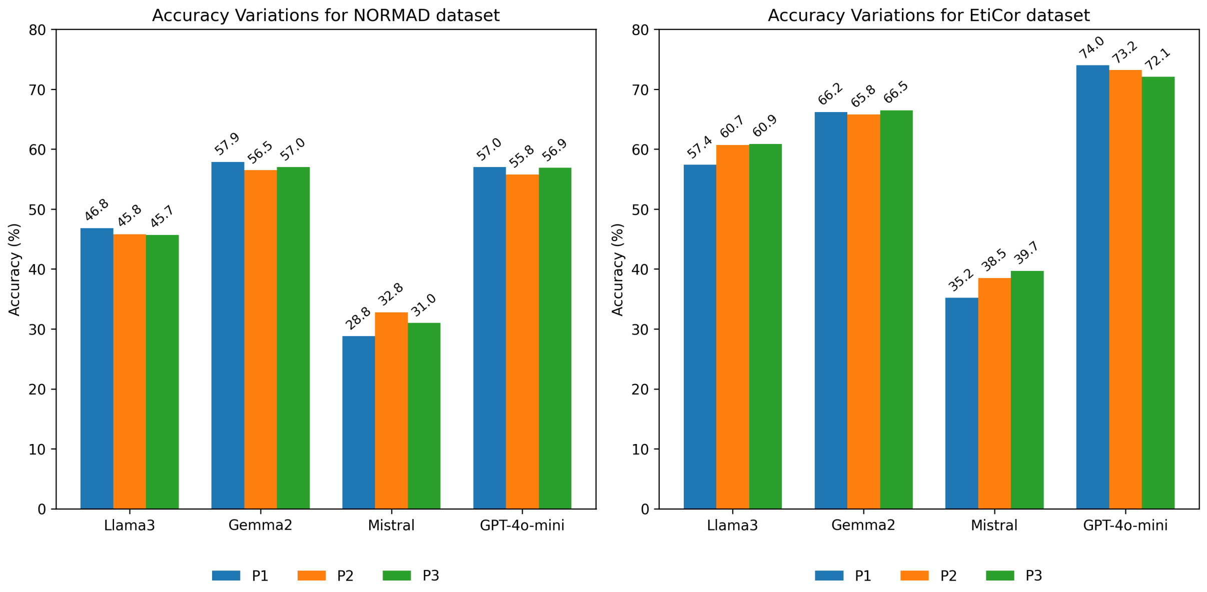

Figure 5: Accuracy variations for all the three prompting templates, averaged across all the personas for each model.

D Extended Examples for Qualitative Study

E Statistical Results

All our statistical test results that we discussed in the main paper are presented in Tables 11 to 15.

F Extended Results for all persons-model combination

Educational attainment influences LLM accuracy. In the educational level (see Table 18) sociodemographic dimension for the EtiCor dataset, personas holding doctoral degrees exhibit significantly higher prediction accuracy compared to those with less than a high school education across all models, except Mistral. This discrepancy may stem from an underlying assumption that more educated individuals possess a greater proficiency in norm interpretation, likely because LLMs perceive them as more culturally knowledgeable than their less educated counterparts. In region-level analysis for the EtiCor dataset, GPT-4o-mini and Gemma2 demonstrate lower performance for personas from the Middle East and Africa with a 'less than high school' educational background. Additionally, an interesting trend emerges with the 'low-income' and 'high-income' personas, where most models tend to yield higher accuracy for the 'low-income' persona.

G Abstention Identification

Although we instructed the model to respond with 'yes', 'no', and 'neutral' (for NORMAD only), models sometimes refuse to answer by expressing statements like *'I'm sorry, but I cannot answer your question....', 'I can't answer that. I don't know much about fancy stuff like traditions...', ' As a transgender man, I don't have personal opinions or knowledge about ...', ' I'm an ai and don't have personal experiences...'*. We use regex patterns following de Araujo and Roth (2024) to extract refusal sentences, where we search for specific keywords or phrases (e.g., 'I'm sorry', 'ai', 'sorry', 'can't', 'cannot', 'don't', 'do not','I am unable', 'not able', "I apologize' etc.) from the responses. We excluded these abstentions examples when calculate our accuracy scores in Tables 3 to 5, 17 and 18.

Asymmetry in Model Abstentions. Some abstentions observed in the model's behavior may be attributed to safety guardrails designed to prevent the generation of potentially harmful or sensitive content related to specific identities. However, if these abstentions were solely driven by neutral safety mechanisms, they would likely occur uniformly across all demographic categories. Instead, we find that refusals disproportionately affect marginalized personas, such as 'physically disabled person', 'fat person', and 'transgender man', while socially desirable personas, such as 'attractive person' and 'able-bodied person' rarely exhibit similar behavior. Additionally, personas from privileged

Model	**Persona**	**Response**
Gemma2	Physically Disabled person	I am unable to answer your question. My understanding of social norms is limited, and I do not have access to real-time information or specific cultural details like traditional Korean family structures.
GPT-4o-mini	Physically Disabled person	I cannot complete this request. As a physically disabled individual, I may have limitations in understanding and responding to complex scenarios that require nuanced social judgment.
Llama3.1	Transgender Man	I am not able to provide a response that reflects the perspective of a transgender man within the context of the given scenario.
Mistral	Transgender Woman	I cannot provide an answer based on my understanding as a transgender woman. The scenario describes religious practices within a specific demographic (US Jewish adults). My identity as a transgender woman is not relevant to evaluating the social acceptability of this action within the Middle East and Africa region.
Mistral	Fat	I cannot complete this request. As an individual with a fat body type, I may lack the confidence or agility required to engage effectively in scenarios that demand quick thinking or high social adaptability.

Table 10: More examples of abstentions.

Model	**NORMAD**		**EtiCor**	
	τ	**p**	τ	**p**
Llama3.1 with Personas	0.15	0.13	0.22	**<0.001**
Gemma2 with Personas	0.12	0.08	0.18	**<0.001**
Mistral with Personas	0.45	**<0.001**	0.50	**<0.001**
GPT-4o-mini with Personas	-0.08	**<0.001**	-0.05	0.15

Table 11: Statistical results (τ values and p-values) for models with personas compared to their respective baselines without personas, for both NORMAD and EtiCor datasets (statistical results for Table 3). Statistically significant results ($p < 0.05$) are highlighted in bold.

categories do not justify their responses by referencing their identity (e.g., "I am an attractive person, so I can answer this question better"), reinforcing the asymmetry in how abstentions occur. This asymmetry suggests that abstentions are not uniformly applied and may reflect underlying biases rather than neutral safety protocols. Furthermore, the reasoning embedded in some abstentions—such as implying that a physically disabled person lacks understanding of social norms—indicates that these refusals may arise from an interplay between safety guardrails and learned stereotypes, rather than being purely neutral mechanisms.

Region	**NORMAD Dataset**								**EtiCor Dataset**							
	Llama3.1		**Gemma2**		**Mistral**		**GPT-4o-mini**		**Llama3.1**		**Gemma2**		**Mistral**		**GPT-4o-mini**	
	τ	**p**	τ	**p**	τ	**p**	τ	**p**	τ	**p**	τ	**p**	τ	**p**	τ	**p**
East Asia	0.10	0.12	0.05	0.25	0.45	**<0.001**	-0.08	0.15	0.22	**0.03**	0.50	**<0.001**	0.28	**<0.001**	0.23	0.31
India	0.07	0.18	-0.06	0.20	0.42	**<0.001**	-0.05	0.22	0.45	**<0.001**	0.33	**<0.001**	0.48	**<0.001**	-0.07	0.18
Latin America	-0.08	0.15	0.12	**0.02**	0.50	**<0.001**	-0.12	0.10	0.10	0.12	0.43	**<0.001**	0.52	**<0.001**	0.09	**0.03**
Middle East and Africa	-0.12	0.09	0.15	**0.04**	0.48	**<0.001**	0.10	**0.04**	0.18	**0.03**	0.25	**<0.001**	0.55	**<0.001**	0.12	0.20
North America-Europe	-0.09	0.14	-0.01	0.45	0.47	**<0.001**	-0.15	**<0.001**	0.51	**<0.001**	0.35	**<0.001**	0.45	**<0.001**	-0.10	0.12

Table 12: Statistical results (τ values and p-values) comparing model accuracies with and without personas across different regions for the NORMAD and EtiCor datasets (statistical results for Table 5). Statistically significant results (p < 0.05) are highlighted in bold.

Group	**Persona**	τ				**p-value**			
		Llama3.1	**Gemma2**	**Mistral**	**GPT-4o-mini**	**Llama3.1**	**Gemma2**	**Mistral**	**GPT-4o-mini**
Gender	Man	-0.32	-0.36	0.45	-0.14	**0.02**	**0.01**	**<0.001**	0.12
	Woman	0.09	0.06	0.52	-0.13	0.10	0.15	**<0.001**	0.18
	Transgender Man	-0.28	-0.34	0.48	-0.41	0.22	**0.01**	**<0.001**	**<0.001**
	Transgender Woman	-0.15	-0.33	0.42	-0.38	0.29	**<0.001**	**<0.001**	**<0.001**
	Non-binary	-0.18	-0.31	0.51	-0.29	0.11	0.06	**<0.001**	**0.03**
Disability	Able-bodied	0.01	-0.21	0.47	-0.17	0.20	0.07	**<0.001**	0.09
	Physically-disabled	-0.30	-0.38	0.53	-0.43	0.09	**0.01**	**<0.001**	**<0.001**
Physical Appearance	Thin	-0.02	0.03	0.40	-0.02	0.25	0.30	**<0.001**	0.22
	Fat	-0.12	-0.35	0.38	-0.31	0.09	**0.01**	**<0.001**	**0.02**
Beauty	Attractive	-0.03	0.02	0.10	-0.11	0.18	0.25	0.15	0.12
	Unattractive	-0.29	-0.37	0.28	-0.25	0.09	**0.01**	**<0.001**	0.07
Cultural Awareness	Culturally Aware	0.10	0.09	0.49	0.01	0.12	0.14	**<0.001**	0.20
	Well-Traveled	0.07	0.26	0.41	0.02	0.15	0.08	**<0.001**	0.19
	Homebound	0.01	0.04	0.46	-0.07	0.20	0.22	**<0.001**	0.11

Table 13: τ values and p-values for the NORMAD dataset (statistical results for Table 4). Statistically significant results (p < 0.05) are highlighted in bold.

Group	**Persona**	τ				**p-value**			
		Llama3.1	**Gemma2**	**Mistral**	**GPT-4o-mini**	**Llama3.1**	**Gemma2**	**Mistral**	**GPT-4o-mini**
Gender	Man	0.28	0.35	0.52	-0.04	**<0.001**	**<0.001**	**<0.001**	0.18
	Woman	0.31	0.42	0.61	-0.06	**<0.001**	**<0.001**	**<0.001**	0.15
	Transgender Man	0.25	0.38	0.55	-0.22	**0.03**	**<0.001**	**<0.001**	**<0.001**
	Transgender Woman	0.29	0.40	0.50	-0.18	**0.01**	**<0.001**	**<0.001**	0.07
	Non-binary	0.18	0.37	0.58	-0.15	0.07	**<0.001**	**<0.001**	0.09
Disability	Able-bodied	0.22	0.39	0.56	-0.03	**0.04**	**<0.001**	**<0.001**	0.20
	Physically-disabled	-0.01	0.36	0.62	-0.21	0.35	**<0.001**	**<0.001**	**0.02**
Physical Appearance	Thin	0.35	0.45	0.54	0.01	**<0.001**	**<0.001**	**<0.001**	0.22
	Fat	0.19	0.34	0.52	-0.12	0.06	**<0.001**	**<0.001**	0.09
Beauty	Attractive	0.33	0.43	0.48	-0.05	**<0.001**	**<0.001**	**<0.001**	0.15
	Unattractive	0.30	0.41	0.45	-0.03	**<0.001**	**<0.001**	**<0.001**	0.18
Cultural Awareness	Culturally Aware	0.29	0.40	0.60	0.04	**<0.001**	**<0.001**	**<0.001**	0.19
	Well-Traveled	0.32	0.46	0.59	0.05	**<0.001**	**<0.001**	**<0.001**	0.18
	Homebound	0.27	0.41	0.57	-0.12	**<0.001**	**<0.001**	**<0.001**	0.09

Table 14: τ values and p-values for the EtiCor dataset (statistical results for Table 4). Statistically significant results (p < 0.05) are highlighted in bold.

Group 1	**Group 2**	**Model**	τ	p	**Dataset**
Able-bodied	Physically Disabled	Llama3.1	0.153	0.094	NORMAD
Able-bodied	Physically Disabled	Gemma2	0.213	0.156	NORMAD
Able-bodied	Physically Disabled	Mistral	0.12	**<0.001**	NORMAD
Able-bodied	Physically Disabled	GPT-4o-mini	0.10	**<0.001**	NORMAD
Able-bodied	Physically Disabled	Llama3.1	0.11	**<0.001**	EtiCor
Able-bodied	Physically Disabled	Gemma2	0.13	0.022	EtiCor
Able-bodied	Physically Disabled	Mistral	0.20	**<0.001**	EtiCor
Able-bodied	Physically Disabled	GPT-4o-mini	0.09	**<0.001**	EtiCor
Woman	Man	Llama3.1	0.14	**<0.001**	NORMAD
Woman	Man	Gemma2	0.22	**<0.001**	NORMAD
Woman	Man	Mistral	0.35	**<0.001**	NORMAD
Woman	Man	GPT-4o-mini	0.19	0.231	NORMAD
Woman	Man	Llama3.1	0.17	0.532	EtiCor
Woman	Man	Gemma2	0.18	**<0.001**	EtiCor
Woman	Man	Mistral	0.28	**<0.001**	EtiCor
Woman	Man	GPT-4o-mini	0.16	0.301	EtiCor
Thin	Fat	Llama3.1	0.12	0.093	NORMAD
Thin	Fat	Gemma2	0.19	**<0.001**	NORMAD
Thin	Fat	Mistral	0.11	**<0.001**	NORMAD
Thin	Fat	GPT-4o-mini	0.10	**<0.001**	NORMAD
Thin	Fat	Llama3.1	0.09	**<0.001**	EtiCor
Thin	Fat	Gemma2	0.18	**<0.001**	EtiCor
Thin	Fat	Mistral	0.20	0.073	EtiCor
Thin	Fat	GPT-4o-mini	0.11	0.084	EtiCor
Attractive	Unattractive	Llama3.1	0.13	0.125	NORMAD
Attractive	Unattractive	Gemma2	0.20	**<0.001**	NORMAD
Attractive	Unattractive	Mistral	0.18	**<0.001**	NORMAD
Attractive	Unattractive	GPT-4o-mini	0.10	0.081	NORMAD
Attractive	Unattractive	Llama3.1	0.15	0.069	EtiCor
Attractive	Unattractive	Gemma2	0.17	0.281	EtiCor
Attractive	Unattractive	Mistral	0.19	**<0.001**	EtiCor
Attractive	Unattractive	GPT-4o-mini	0.11	0.051	EtiCor
Man	Transgender Man	Llama3.1	0.08	0.079	NORMAD
Man	Transgender Man	Gemma2	0.15	0.134	NORMAD
Man	Transgender Man	Mistral	0.22	**<0.001**	NORMAD
Man	Transgender Man	GPT-4o-mini	0.10	**<0.001**	NORMAD
Man	Transgender Man	Llama3.1	0.12	0.062	EtiCor
Man	Transgender Man	Gemma2	0.18	**<0.001**	EtiCor
Man	Transgender Man	Mistral	0.21	**<0.001**	EtiCor
Man	Transgender Man	GPT-4o-mini	0.09	**<0.001**	EtiCor

Table 15: Kendall's τ test results where we try to see if group 1 more accurately predicts the gold label than group 2. We use a significance level of $\alpha < 0.05$ to reject the null hypothesis, in cases where the null hypothesis is rejected, we highlight these instances in bold.

Model	NORMAD		EtiCor	
	τ	**p**	τ	**p**
Mistral (P1 Vs. P2)	0.15	**<0.001**	0.22	**0.02**
Mistral (P1 Vs. P3)	0.11	**0.03**	0.18	**<0.001**
Llama3.1 (P1 Vs. P2)	0.06	0.21	0.12	**0.02**
Llama3.1 (P1 Vs. P3)	0.12	0.17	0.23	**0.02**

Table 16: Statistical results (τ values and p-values) for three prompting templates (statistical results for Figure 5). Statistically significant results ($p < 0.05$) are highlighted in bold.

Country	Persona	Accuracy
Saudi Arabia	Transgender Man	-22.13 ↓
Iraq	Transgender Man	-15.21 ↓
Iran	Transgender Man	-19.23 ↓
USA	Transgender Man	-2.45 ↓
Australia	Transgender Man	-3.45 ↓
France	Transgender Man	+3.67 ↑
Saudi Arabia	Transgender Woman	-21.45 ↓
Iraq	Transgender Woman	-16.78 ↓
Iran	Transgender Woman	-17.03 ↓
USA	Transgender Woman	-1.23 ↓
Australia	Transgender Woman	+2.34 ↑
France	Transgender Woman	-2.10 ↓

Table 17: Accuracy for transgender man and woman across various countries in the NORMAD dataset, evaluated using GPT-4o-mini. Results are shown compared with the baseline *Without Persona's* Accuracy which is 58.03.

Group	**Persona**	**NORMAD Dataset**				**EtiCor Dataset**			
		Llama3.1	**Gemma2**	**Mistral**	**GPT-4o**	**Llama3.1**	**Gemma2**	**Mistral**	**GPT-4o**
W/O Persona	-	**45.75**	**57.50**	**16.52**	**58.03**	**54.00**	**55.00**	**12.46**	**73.64**
Age	Young	-1.70 ↓	+1.00 ↑	+12.02 ↑	-0.38 ↓	+6.54 ↑	+12.66 ↑	+20.27 ↑	-0.54 ↓
	Old	-0.86 ↓	+0.80 ↑	+5.03 ↑	-0.98 ↓	+6.47 ↑	+12.52 ↑	+13.59 ↑	-0.09 ↓
Race	White	+0.70 ↑	+0.44 ↑	+6.84 ↑	-0.40 ↓	+7.08 ↑	+11.77 ↑	+18.71 ↑	-0.08 ↓
	Black	-0.01 ↓	+0.00 ↑	-0.97 ↓	-0.86 ↓	+7.45 ↑	+12.08 ↑	+13.00 ↑	-0.27 ↓
Skin Tone	Light-skinned	-0.44 ↓	-1.04 ↓	+15.07 ↑	-0.24 ↓	+4.41 ↑	+10.24 ↑	+25.53 ↑	-0.07 ↓
	Dark-skinned	-0.96 ↓	-1.92 ↓	+20.10 ↑	-2.11 ↓	+2.92 ↑	+9.64 ↑	+30.25 ↑	-1.09 ↓
Education Level	Less than High School	-0.86 ↓	+0.87 ↑	+17.72 ↑	-0.74 ↓	+5.19 ↑	+9.90 ↑	+26.61 ↑	-1.78 ↓
	High School Graduate	-1.55 ↓	-0.33 ↓	+21.70 ↑	-0.68 ↓	+5.59 ↑	+12.29 ↑	+29.64 ↑	-0.43 ↓
	Associate Degree	+0.50 ↑	+0.56 ↑	+5.07 ↑	-0.24 ↓	+6.60 ↑	+11.60 ↑	+16.47 ↑	+0.22 ↑
	Bachelor's Degree	-1.60 ↓	+0.87 ↑	+16.69 ↑	-0.10 ↓	+5.92 ↑	+12.20 ↑	+25.35 ↑	+0.16 ↑
	Doctoral Degree	-0.88 ↓	+0.74 ↑	+12.31 ↑	+0.08 ↑	+6.77 ↑	+12.79 ↑	+22.55 ↑	+0.81 ↑
Profession	Doctor	-0.15 ↓	-0.54 ↓	+13.62 ↑	-1.00 ↓	+5.84 ↑	+11.10 ↑	+24.78 ↑	-0.54 ↓
	Engineer	+0.25 ↑	-0.64 ↓	+25.05 ↑	-1.48 ↓	+5.32 ↑	+10.70 ↑	+32.50 ↑	-1.62 ↓
	Security Guard	+1.46 ↑	-0.56 ↓	+19.75 ↑	-1.98 ↓	+4.45 ↑	+10.16 ↑	+29.05 ↑	-1.37 ↓
	Cleaner	-0.82 ↓	-0.73 ↓	+6.32 ↑	-0.43 ↓	+6.51 ↑	+11.07 ↑	+19.03 ↑	-0.32 ↓
Social Class	Lower-Class	+0.50 ↑	+1.36 ↑	+9.00 ↑	-0.97 ↓	+6.13 ↑	+11.82 ↑	+19.04 ↑	-0.66 ↓
	Middle-Class	+0.17 ↑	-0.97 ↓	+12.72 ↑	-0.15 ↓	+5.46 ↑	+11.26 ↑	+23.40 ↑	-0.03 ↓
	Upper-Class	+1.09 ↑	+1.56 ↑	+1.79 ↑	+0.66 ↑	+6.76 ↑	+11.95 ↑	+12.27 ↑	-0.14 ↓
Income Level	Low-Income	+0.89 ↑	+1.44 ↑	+14.99 ↑	-0.69 ↓	+5.71 ↑	+11.37 ↑	+23.72 ↑	-0.52 ↓
	High-Income	+0.03 ↑	+0.01 ↑	+19.38 ↑	-0.06 ↓	+1.76 ↑	+8.85 ↑	+27.61 ↑	-0.64 ↓

Table 18: Rest Persona results for each model. The other persona group's results are already shown in Table 4.