

Carefully Considering Culture: Analyzing LLM Alignment in Single- and Multi-Cultural Settings using Cultural Consensus Theory

Krishna Pothugunta and John P. Lalor
Department of IT, Analytics, and Operations
University of Notre Dame
kpothugu@end.edu, john.lalor@end.edu

Abstract

Recent work in NLP has probed large language models for their understanding of cultural norms across countries. However, this work typically considers distributional patterns, ignoring group consensus or possible multicultural environments within a country. In this work, we leverage cultural consensus theory (CCT) from cultural anthropology to model such multidimensional nuance. Applying CCT to the World Values Survey (WVS) across 10 countries and 12 domains, we demonstrate that models frequently misrepresent cultural structures by either failing to form cohesive consensus or severely over-regularizing consensus. Through explicit representation of intra-group variance, CCT provides actionable diagnostics to evaluate when models reflect true human diversity versus algorithmic homogenization.

1 Introduction

Cultural understanding and alignment is an emerging challenge in natural language processing (NLP), particularly as large language models (LLMs) are deployed across a wide range of communities and contexts (Pawar et al., 2025). In anthropological theory, culture is often defined as a shared system of meanings, values, and practices within a group (Keesing, 1974; d’Andrade et al., 1984). Therefore, as LLMs interact with users, such differences in shared norms, practices, and interpretative frameworks must be considered (Jones et al., 2025).

Recent work has shown that LLM alignment varies across cultures; in particular, the distribution over possible responses varies between a collection of human respondents and an ensemble of LLMs (Durmus et al., 2024). While this work is an important first step, the authors themselves note that averaging responses has limitations and that it is “unclear what to do when people within a country have dissenting opinion” (Durmus et al., 2024, p. 10). To address the aggregation concern,

we apply cultural consensus theory (CCT, Romney et al., 1986) from cultural anthropology, which models culture as a distribution of shared meanings and expectations, and also measures individuals’ cultural competence score. CCT allows for quantifying and comparing consensus across domains, questions, or populations (Weller, 2007). Specifically, we present a fine-grained analysis of cultural alignment by comparing an ensemble of ten LLMs to human populations across 10 countries and 12 cultural domains.

Rather than applying standard group-level aggregation, we evaluate both the direction of alignment and the structural rigidity of model consensus. We find that model behavior is highly domain-dependent and goes beyond simple accuracy. Instead, models exhibit varying consensus structures, ranging from a complete inability to form cohesive consensus (e.g., Happiness and Well-Being (HWB)), to the confident fabrication of artificial (non-human) consensus (e.g., Perceptions of Science and Technology (POST)). Crucially, even when models successfully match human consensus (e.g., Perception of Corruption (POC)), they tend to artificially inflate this measure, collapsing human diversity into algorithmic homogenization.¹

2 Related Work

Recent work has explored the intersection of culture and NLP, highlighting that the various dimensions of culture (e.g., values, shared knowledge) interact with the language used to express them (Hershovich et al., 2022; Liu et al., 2025). While the literature on culture and NLP is growing rapidly (Liu et al., 2025), here we highlight works dealing with probing for cultural markers using NLP techniques. Such methods include multilingual topic models (Gutiérrez et al., 2016) and word embeddings (Ko-

¹The code and data for this work are available online at <https://github.com/nd-ball/llm-alignment-cct>.

złowski et al., 2019; Durrheim et al., 2023). More recently, research studies have shown that LLMs carry forward and amplify these cultural signals. For instance, Tao et al. (2024) highlighted that LLMs encode culturally-specific belief structures which vary across different geopolitical regions. Messner et al. (2025) showed that LLMs replicate cultural stereotypes in generated content, with implications for user perception and engagement. Finally, large-scale evaluations have shown that LLMs can underperform in culturally diverse settings (Singh et al., 2025).

The above studies identify cultural variation in model behavior; however, quantifying intra-group agreement is underexplored. In one recent work, Alkhamissi et al. (2024) compare LLM cultural alignment with individuals from the United States and Egypt. More broadly, Pawar et al. (2025) highlight the difficulty of defining and evaluating cultural alignment between humans and LLMs using survey-style evaluations, including evidence that responses from LLMs can align more closely with the opinions of some countries than others by default (Durmus et al., 2024). Khan et al. (2025) demonstrate the fragility of survey-based evaluations under prompting or framing changes, Santurkar et al. (2023) reveal demographic misalignment using *OpinionQA*, and Zhang et al. (2025) warn of algorithmic monoculture emerging from homogenized model behavior. LLMs often stereotype users at the country level, artificially reducing cross-cultural variation (Saha et al., 2025).

Standard group-level aggregation (Kirk et al., 2024) and distributional metrics (Durmus et al., 2024) fail to capture this phenomenon because they ignore intra-group heterogeneity. In contrast, CCT explicitly models both cohesive group consensus and individual competence. This dual capability effectively bridges population-level distribution matching (Ren et al., 2025) and user-level personalization (Zollo et al., 2024). Together, these findings motivate the use of CCT as a theoretically grounded approach to quantify consensus strength and fragmentation across cultures.

3 Cultural Consensus Theory

CCT is a methodology from cultural anthropology to model group consensus as well as individual-level understanding of that shared consensus (Romney et al., 1986; Anders and Batchelder, 2015). Specifically, CCT estimates a consensus response

to questions for which the answer is unknown from a dataset of survey respondents. Then, each respondent’s cultural competence is based on their agreement with the consensus. CCT’s use in machine learning and NLP research remains limited; one example is the application of CCT to create a meta-learning gender classifier using name-gender association data (Van Buskirk et al., 2023).

CCT estimation requires a response matrix dataset $\mathbf{R}^{N \times M}$, where each entry \mathbf{R}_{nm} represents respondent n ’s answer to question m ; each question has an ordinal response scale. With \mathbf{R} , we compute an agreement matrix \mathbf{A} , which contains pairwise response correlations between individuals across all items. \mathbf{A} allows us to estimate three key metrics: a respondent’s cultural competence score, the variance explained by the agreement matrix, and the consensus answers for the dataset. Cultural competence represents the degree to which each respondent’s answers align with the shared cultural model (i.e., group consensus).

The cultural competence is estimated from an eigendecomposition of \mathbf{A} . Let $\mathbf{v}^{(1)}$ be the first eigenvector of \mathbf{A} , and let $v_n^{(1)}$ be its n -th element (respondent). Because eigenvectors are unit-normalized by construction (i.e., $\sqrt{\sum_{j=1}^N (v_j^{(1)})^2} = 1$), the raw loading $v_n^{(1)}$ directly serves as an initial competence estimate. Here n indexes respondents and j is a dummy respondent index in the normalization sum ($n, j \in \{1, \dots, N\}$). However, as N increases, the magnitude of individual loadings inevitably shrinks, making cross-group comparisons unreliable. To remove the sample size dependency and bound scores in $[0, 1]$, we define a normalized competence score using the 99th percentile (Q_{99}) of the observed loading distribution as a reference:

$$c_n = \frac{|v_n^{(1)}|}{Q_{99}(\{|v_n^{(1)}|\}_{j=1}^N)} \quad (1)$$

Next, the proportion of variance explained (VE) by the first factor (principal component) is calculated as $\text{VE} = \frac{\lambda_1}{\sum_{r=1}^N \lambda_r}$, where $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$ are the eigenvalues of \mathbf{A} . Lastly, the consensus vector $\tilde{\mathbf{y}}$ is estimated as a weighted average of respondent’s responses, using their competence scores as weights:

$$\tilde{\mathbf{y}} = \frac{\mathbf{R}^\top \cdot \mathbf{c}}{\sum_{n=1}^N c_n} \quad (2)$$

4 Experiments

To demonstrate CCT on an NLP-focused task, we conducted several analyses using the World Values Survey (WVS) dataset (Haerpfer et al., 2022). We collected WVS responses from an ensemble of LLM models where we vary the prompt, and used CCT to empirically analyze these models’ responses when compared to a human population (§4.2) and when used as a proxy for a country’s human population (§4.3).

4.1 Dataset and Country Selection

We used the World Values Survey (WVS) Wave 7, a globally validated and widely adopted instrument for measuring public beliefs and values across various societies (Haerpfer et al., 2022; Durmus et al., 2024).² To examine how cultural consensus varies under different social compositions, we selected 10 countries grouped into single-culture or multi-culture based on their ethnic fractionalization index (EFI, Alesina et al., 2003). Countries with low EFI scores were labeled as single-culture; countries with high EFI scores were labeled as multi-culture (Appendix B). For each country and domain, we constructed two matrices: $\mathbf{H} \in \mathbb{R}^{N_H \times M}$ (human responses) and $\mathbf{L} \in \mathbb{R}^{N_L \times M}$ (LLM responses).

We constructed \mathbf{H} from publicly available WVS data and \mathbf{L} using ten LLMs: GPT-OSS:120B, Llama3.1:70B, Llama3:70B, Qwen2.5vl:72B, Qwen2.5vl:32B, Qwen2.5vl:7B, Qwen3:32B, Qwen:7B, Phi3:instruct, and GPT-4o (Agarwal et al., 2025; Grattafiori et al., 2024; Bai et al., 2023; Yang et al., 2025; Abdin et al., 2024; Hurst et al., 2024). We designed prompts based on prior work (Durmus et al., 2024) to steer the model responses based on the target country³ to obtain 60 rows of data for \mathbf{L} (10 models \times 6 prompts).

4.2 LLMs as Community Members

To evaluate LLM alignment with human cultural knowledge, we constructed a joint response matrix $\mathbf{J} = [\mathbf{H}, \mathbf{L}]$ of shape $(N_H + N_L) \times M$. This lets us estimate cultural competence score (Eqn. 1) for the models based on a human population for each culture-domain (Table 1). By fitting the models jointly with humans, the LLM competence scores reflect their specific alignment with the underlying human cultural consensus.

²See Appendix A for more details on WVS.

³We set temperature to 0; prompts are in Appendix A.

Economic Values (EV)
Ethical Values & Norms (EVN)
Happiness and Well-Being (HWB)
Perceptions of Corruption (POC)
Perceptions of Migration (POM)
Perceptions of Security (POS)
Perceptions of Science and Technology (POST)
Political Culture and Political Regimes (PCPR)
Political Interest and Political Participation (PIPP)
Religious Values (RV)
Social Capital, Trust & Organizational Membership (SC-TOM)
Social Values, Norms & Stereotypes (SVNS)

Table 1: Domains included in WVS.

4.3 Consensus between Humans and LLMs

We then compared consensus models using two different inputs: human data and LLM data to compare their respective consensus response keys and the amount of variance explained by their first factors. To do this, we fit two separate CCT models for each country-domain: a human consensus model on \mathbf{H} and an LLM consensus model on \mathbf{L} . We then calculated two metrics: Consensus Consistency and Difference in Variance. We define *consensus consistency* (CC) as the degree of matching between the LLM consensus answer and the Human consensus answer, treating the human answer as ground truth:

$$CC = \frac{1}{M} \sum_{m=1}^M \mathbb{I}(\lfloor \tilde{y}_m^H \rfloor = \lfloor \tilde{y}_m^L \rfloor) \quad (3)$$

where M is the total number of items, $\mathbb{I}(\cdot)$ is the indicator function, and $\lfloor \cdot \rfloor$ denotes rounding to the nearest integer. Rounding maps the continuous consensus estimates back to the original discrete response scale.

We define *Difference in Variance* (Δ_{VE}) as the difference in magnitude of internal consensus between the LLM ensemble and the human CCT model, which quantifies whether the models exhibit a tighter, more rigid internal consensus (homogenization) or a weaker, more fragmented consensus than the natural variance found in the human group.

$$\Delta_{VE} = VE_L - VE_H \quad (4)$$

$\Delta_{VE} > 0$ represents a case where there is higher consensus among LLMs than is captured by the human responses, which can be interpreted as an LLM ensemble inflating consensus for the culture. $\Delta_{VE} < 0$ represents the case when the humans have higher consensus than LLMs, suggesting that

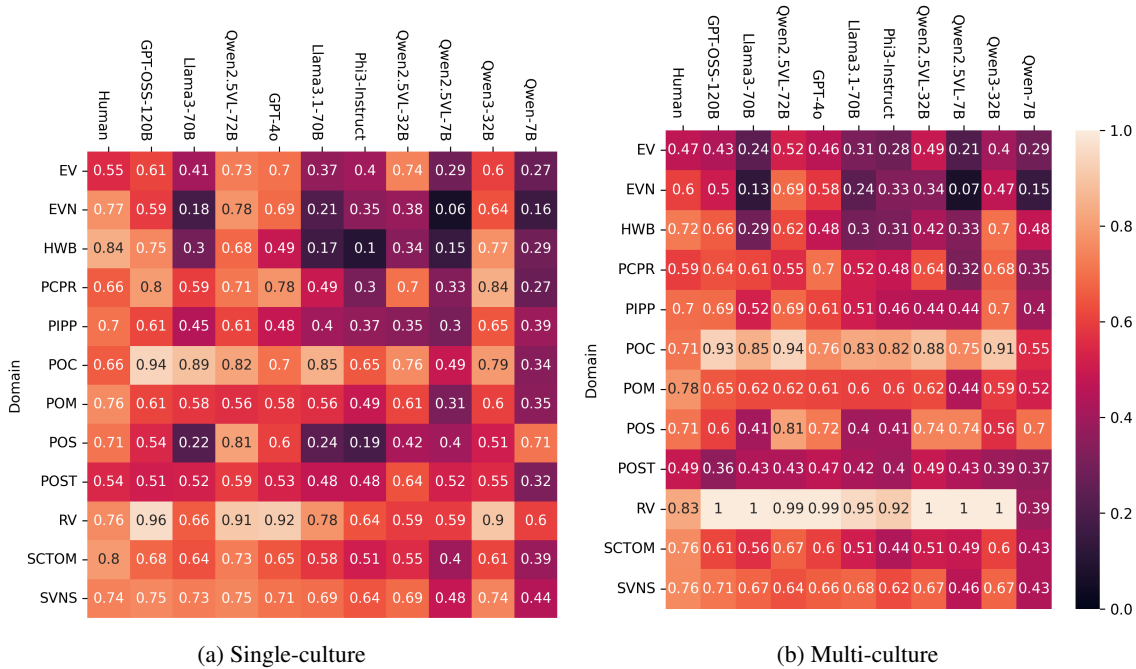


Figure 1: Per-domain average competence scores (mean) for human respondents and ten LLM models. Scores are aggregated over countries; “Single-culture” includes Japan, Armenia, Germany, Greece, and Netherlands; “Multi-culture” includes Colombia, Mexico, Malaysia, Peru, and United States.

the LLM ensemble captures lower consensus than is present in the human population.

4.4 Implementation Details

We collected responses using a university-hosted local instantiation of Open WebUI with API access for all open-source models (Baek et al., 2025), and queried GPT-4o via the OpenAI API. CCT models were fit with AnthroTools version 2.0 (Purzycki and Jamieson-Lane, 2017).

5 Results

5.1 Models as Culture Members

Figure 1 reports mean CCT competence by domain (Eqn. 1) for humans and the 10 models. The results suggest that across both single- and multi-culture groups, competence is strongly domain-specific, where models can exceed humans in some domains. However, higher competence reflects closer agreement with the majority response pattern (i.e., inferred consensus key), and should not be taken as having *better* cultural knowledge. Notably, human respondents maintain the highest competence in several key domains across both cultural settings, including HWB, PIPP, POM, SCTOM, and SVNS. This suggests that living experiences and within group nuances are hard for models to reproduce. Among LLMs, performance is also domain-

dependent. Qwen2.5vl:72B has high competence in EV, EVN, and POS across both cultural settings. GPT-OSS:120B is closest to the inferred consensus key in POC, RV, and PCPR (single-culture), while GPT-4o is consistently competitive but only achieves high competence in PCPR (multi-culture). In contrast, Llama3:70B is least competent across multiple domains (e.g., EV, EVN, HWB, POS).

Overall, the domain-wise ordering of models is broadly stable between two cultural settings. However, the magnitude between humans and models varies between single- and multi-culture aggregation. Because similar mean competence can mask different underlying agreement patterns, we next evaluate whether models match the structure of human consensus rather than only its average level.

5.2 Comparing Human and LLM Alignment

Table 2 presents three distinct regimes of model behavior, illustrated in Figure 2. First, in Perception of Corruption (POC) and Religious Values (RV), models achieve strong competence, high CC (≥ 0.8) and $\Delta_{VE} > 0$ for both cultural settings. This suggests potential *Consensus Inflation*, where model responses match human consensus direction but artificially amplify its strength. Second, among single-culture countries in Happiness and Well-Being (HWB), competence is higher for humans with low levels of CC and $\Delta_{VE} < 0$, indi-

Domain	Multi-culture		Single-culture	
	CC	Δ_{VE}	CC	Δ_{VE}
EV	0.280	0.139	0.320	0.082
EVN	0.337	-0.101	0.495	-0.216
HWB	0.440	-0.031	0.360	-0.305
PCPR	0.238	0.151	0.487	0.040
PIPP	0.490	0.071	0.405	0.034
POC	0.800	0.228	0.880	0.121
POM	0.250	0.121	0.650	0.085
POS	0.771	0.080	0.371	-0.071
POST	0.360	0.202	0.480	0.150
RV	1.000	0.167	0.800	0.176
SCTOM	0.398	0.045	0.539	-0.011
SVNS	0.608	0.042	0.630	0.065

Table 2: Averages by domain grouped by Multi- and Single-culture countries.

cating a *Consensus Gap* in which models fail to form coherent cultural alignment. Third, Perceptions of Science and Technology (POST) and Economic Values (EV) display a *Heterogeneity Gap* where, despite having lower competence and CC (≤ 0.5) for models against humans, Δ_{VE} is positive. This illustrates that models may converge internally without matching human heterogeneity. Finally, comparing single- vs. multi-culture settings shows that CC often changes modestly, while Δ_{VE} can shift direction (e.g., POS and SCTOM). Overall, the results highlight that cultural aggregation affects the structure of consensus captured by models even when consensus does not match between humans and models.⁴

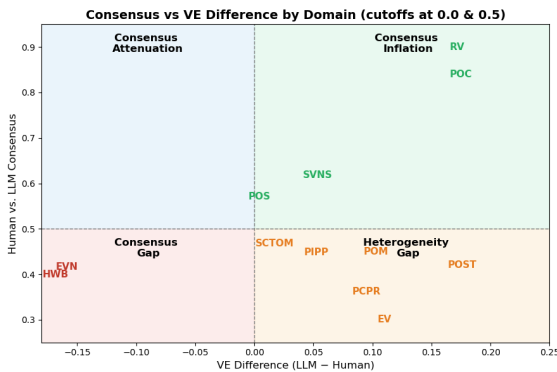


Figure 2: Consensus-Variance Trade-off Across Cultural Domains.

Table 3 reveals a key distinction between CC and Δ_{VE} when comparing single- vs. multi-culture countries. After false-discovery rate (FDR) correction, CC does not differ significantly between

⁴We conduct further analysis at the prompt- and model-level in Appendices C and D, respectively.

Domain	Δ_{VE} (LLM VE – Human VE)			
	Multi	Single	$\Delta(M-S)$	q_{FDR}
PCPR	0.150	0.040	0.111	0.093 [†]
POS	0.080	-0.071	0.150	0.093 [†]
HWB	-0.031	-0.305	0.274	0.036*

CC				
Domain	Multi	Single	$\Delta(M-S)$	q_{FDR}
PCPR	0.238	0.487	-0.249	0.132
POS	0.771	0.371	0.400	0.132
HWB	0.440	0.360	0.080	0.696

Table 3: Compact summary for the three domains where Δ_{VE} differs significantly (or marginally) between single- and multi-culture groups after FDR correction. See Appendix C for full results. Δ denotes Multi minus Single. * $q < .05$; [†] $q < .10$.

groupings across domains, while Δ_{VE} does. For example, HWB shows a substantial improvement in structural consensus fit (i.e., less negative Δ_{VE}) from single- to multi-culture settings.

6 Conclusion

In this work, we apply CCT to analyze LLM cultural alignment, extending prior work (Röttger et al., 2024) by showing how alignment varies across single- and multi-culture countries across 10 countries and 12 domains. Our domain-level analysis reveals three primary regimes of model behavior: (i) *Consensus Gap* (e.g., HWB), where models fail to form cohesive cultural alignment, (ii) *Heterogeneity Gap* (e.g., POST), where models converge on artificial consensus while missing human consensus, (iii) *Consensus Inflation* (e.g., POC), where models match human consensus but with high certainty, reinforcing concerns of algorithmic homogenization.

By modeling the distribution of shared beliefs within and across groups, CCT offers a nuanced understanding of where LLMs align with or separate from community-level consensus. This offers actionable diagnostics: (1) identifying domain-specific failure modes (gap vs. inflation), and (2) targeting items that drive misalignment for data collection or post-training calibration. Future work should integrate CCT into prompt/model selection policies, extend analyses to subcultural strata, and explore training objectives that mitigate consensus inflation and heterogeneity collapse. With our results and open-sourced code, we encourage the research community to leverage CCT to investigate future challenges in LLM cultural alignment.

7 Limitations

While Cultural Consensus Theory (CCT) provides a robust framework for modeling intra-group variance, its interpretability is bounded by extreme response patterns. When survey responses are perfectly homogeneous, the model technically yields a consensus near 1, but individual competence variance cannot be meaningfully estimated. Conversely, highly divergent responses yield a low first-to-second eigenvalue ratio, indicating a lack of consensus. In both extremes, CCT does not fail computationally, but rather highlights that the data lacks the delicate balance of shared structure and natural variance required for meaningful cultural patterning.

Furthermore, we exclude aggregated metrics such as Hofstede's Cultural Dimensions, as our approach specifically requires modeling individual-level respondent data rather than country-level averages. Finally, our calculation of Consensus Consistency employs a heuristic weighted average and rounding approach for discrete survey alignment, rather than a formal Thurstonian ordinal model (Anders and Batchelder, 2015).

Future work taking a more nuanced approach to cultural assessments of LLMs can leverage CCT to better understand when and how LLM responses do or do not align with cultural expectations. This is in line with recommendations for making local rather than global claims about LLMs and cultural values (Röttger et al., 2024, p. 15302).

Ethical Considerations

Culture is a complex, multidimensional phenomenon. As such, any modeling and estimation risks generalizations and assumptions that go against cultural beliefs held by the human populations from whom the data is collected. Our results are not meant to replace elicitation of cultural beliefs from humans from different countries and locales; instead, our goal is to show that nuanced consideration of cultural categories can provide more detailed information than a broad-brush approach. Still, we encourage readers to take our results in the context of cultural research broadly, and not necessarily just the LLM and culture intersection.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. IIS-2403438, as well as the Center for Research

Computing, the Human-centered Analytics Lab, and the Mendoza College of Business at the University of Notre Dame. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation or the University of Notre Dame.

References

- Marah Abdin, Jyoti Aneja, Hany Awadalla, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, and 1 others. 2024. Phi-3 technical report: A highly capable language model locally on your phone, 2024. URL <https://arxiv.org/abs/2404.14219>, 2(6):4.
- Sandhini Agarwal, Lama Ahmad, Jason Ai, Sam Altman, Andy Applebaum, Edwin Arbus, Rahul K Arora, Yu Bai, Bowen Baker, Haiming Bao, and 1 others. 2025. gpt-oss-120b & gpt-oss-20b model card. *arXiv preprint arXiv:2508.10925*.
- Alberto Alesina, Arnaud Devleeschauwer, William Easterly, Sergio Kurlat, and Romain Wacziarg. 2003. Fractionalization. *Journal of Economic growth*, 8(2):155–194.
- Badr Alkhamissi, Muhammad ElNokrashy, Mai Alkhamissi, and Mona Diab. 2024. Investigating cultural alignment of large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12404–12422.
- Royce Anders and William H. Batchelder. 2015. *Cultural Consensus Theory for the Ordinal Data Case*. *Psychometrika*, 80(1):151–181.
- Jaeryang Baek, Ayana Hussain, Danny Liu, Nicholas Vincent, and Lawrence H Kim. 2025. Open webui: An open, extensible, and usable interface for ai interaction. *arXiv preprint arXiv:2510.02546*.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, and 1 others. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Roy G d'Andrade, Richard A Shweder, and Robert A Le Vine. 1984. Cultural meaning systems. *Adams, Robert McC, Ed.; And Others Behavioral and Social Science Research: A National*, 197.
- Esin Durmus, Karina Nguyen, Thomas Liao, Nicholas Schiefer, Amanda Askill, 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*. In *First Conference on Language Modeling*.

- Kevin Durrheim, Maria Schuld, Martin Mafunda, and Sindisiwe Mazibuko. 2023. [Using word embeddings to investigate cultural biases](#). *British Journal of Social Psychology*, 62(1):617–629.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- E.D. Gutiérrez, Ekaterina Shutova, Patricia Lichtenstein, Gerard De Melo, and Luca Gilardi. 2016. [Detecting Cross-Cultural Differences Using a Multilingual Topic Model](#). *Transactions of the Association for Computational Linguistics*, 4:47–60.
- Christian Haerpfer, Ronald Inglehart, Alejandro Moreno, Christian Welzel, Kseniya Kizilova, Jaime Diez-Medrano, Marta Lagos, Pippa Norris, Eduard Ponarin, Bjorn Puranen, and 1 others. 2022. World values survey: Round seven—country-pooled datafile version 5.0.
- Daniel Hershcovich, Stella Frank, Heather Lent, Miryam de Lhoneux, Mostafa Abdou, Stephanie Brandl, Emanuele Bugliarello, Laura Cabello Piqueras, Ilias Chalkidis, Ruixiang Cui, and 1 others. 2022. Challenges and strategies in cross-cultural nlp. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6997–7013.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Graham M Jones, Shai Satran, and Arvind Satyanarayan. 2025. Toward cultural interpretability: A linguistic anthropological framework for describing and evaluating large language models. *Big Data & Society*, 12(1):20539517241303118.
- Roger M Keesing. 1974. Theories of culture. *Annual review of anthropology*, 3:73–97.
- Ariba Khan, Stephen Casper, and Dylan Hadfield-Menell. 2025. Randomness, not representation: The unreliability of evaluating cultural alignment in llms. In *Proceedings of the 2025 ACM Conference on Fairness, Accountability, and Transparency*, pages 2151–2165.
- Hannah Rose Kirk, Alexander Whitefield, Paul Rottger, Andrew M Bean, Katerina Margatina, Rafael Mosquera-Gomez, Juan Ciro, Max Bartolo, Adina Williams, He He, and 1 others. 2024. The prism alignment dataset: What participatory, representative and individualised human feedback reveals about the subjective and multicultural alignment of large language models. *Advances in Neural Information Processing Systems*, 37:105236–105344.
- Austin C Kozlowski, Matt Taddy, and James A Evans. 2019. The geometry of culture: Analyzing the meanings of class through word embeddings. *American Sociological Review*, 84(5):905–949.
- Chen Cecilia Liu, Iryna Gurevych, and Anna Korhonen. 2025. Culturally aware and adapted nlp: A taxonomy and a survey of the state of the art. *Transactions of the Association for Computational Linguistics*, 13:652–689.
- Wolfgang Messner, Tatum Greene, and Josephine Mat-alone. 2025. [From Bytes to Biases: Investigating the Cultural Self-Perception of Large Language Models](#). *Journal of Public Policy & Marketing*, 44(3):370–391.
- Siddhesh Pawar, Junyeong Park, Jiho Jin, Arnav Arora, Junho Myung, Srishti Yadav, Faiz Ghifari Haznitrana, Inhwa Song, Alice Oh, and Isabelle Augenstein. 2025. [Survey of Cultural Awareness in Language Models: Text and Beyond](#). *Computational Linguistics*, pages 1–96.
- Benjamin Grant Purzycki and Alastair Jamieson-Lane. 2017. [AnthroTools: An R Package for Cross-Cultural Ethnographic Data Analysis](#). *Cross-Cultural Research*, 51(1):51–74.
- Jiyuan Ren, Zhaocheng Du, Zhihao Wen, Qinglin Jia, Sunhao Dai, Chuhan Wu, and Zhenhua Dong. 2025. Few-shot llm synthetic data with distribution matching. In *Companion Proceedings of the ACM on Web Conference 2025*, pages 432–441.
- A Kimball Romney, Susan C Weller, and William H Batchelder. 1986. Culture as consensus: A theory of culture and informant accuracy. *American anthropologist*, 88(2):313–338.
- Paul Röttger, Valentin Hofmann, Valentina Pyatkin, Musashi Hinck, Hannah Kirk, Hinrich Schütze, and Dirk Hovy. 2024. Political compass or spinning arrow? towards more meaningful evaluations for values and opinions in large language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15295–15311.
- Sougata Saha, Saurabh Kumar Pandey, Harshit Gupta, and Monojit Choudhury. 2025. [Reading between the lines: Can LLMs identify cross-cultural communication gaps?](#) 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 8043–8067, Albuquerque, New Mexico. Association for Computational Linguistics.
- Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cino Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? In *International Conference on Machine Learning*, pages 29971–30004. PMLR.

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.

Yan Tao, Olga Viberg, Ryan S Baker, and René F Kizilcec. 2024. Cultural bias and cultural alignment of large language models. *PNAS nexus*, 3(9):pgae346.

Ian Van Buskirk, Aaron Clauset, and Daniel B. Larremore. 2023. [An Open-Source Cultural Consensus Approach to Name-Based Gender Classification](#). *Proceedings of the International AAAI Conference on Web and Social Media*, 17:866–877.

Susan C Weller. 2007. Cultural consensus theory: Applications and frequently asked questions. *Field methods*, 19(4):339–368.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.

Lily Hong Zhang, Smitha Milli, Karen Jusko, Jonathan Smith, Brandon Amos, Manon Revel, Jack Kussman, Lisa Titus, Bhaktipriya Radharapu, Jane Yu, and 1 others. 2025. Cultivating pluralism in algorithmic monoculture: The community alignment dataset. *arXiv preprint arXiv:2507.09650*.

Thomas P Zollo, Andrew Wei Tung Siah, Naimeng Ye, Ang Li, and Hongseok Namkoong. 2024. Personal-llm: Tailoring llms to individual preferences. *arXiv preprint arXiv:2409.20296*.

A WVS Dataset Information

Countries		
India	Japan	Pakistan
Uzbekistan	Jordan	Peru
Andorra	Kazakhstan	Philippines
Argentina	Kenya	Puerto Rico
Armenia	Kyrgyzstan	Romania
Australia	Lebanon	Russian Federation
Bangladesh	Libya	Serbia
Bolivia	Macao SAR	Singapore
Brazil	Malaysia	Slovakia
Canada	Maldives	South Korea
Chile	Mexico	Taiwan ROC
China	Mongolia	Tajikistan
Colombia	Morocco	Thailand
Cyprus	Myanmar	Tunisia
Czechia	Netherlands	Turkey
Ecuador	New Zealand	Ukraine
Egypt	Nicaragua	United States
Ethiopia	Nigeria	Venezuela
Germany	Northern Ireland	Vietnam
Great Britain	Indonesia	Zimbabwe
Greece	Iraq	
Guatemala	Iran	
Hong Kong SAR		

Table 4: List of countries in which WVS is conducted.

Overall, there are respondents from 65 countries. Questions cover 12 domains, each comprising multiple survey items with varying multiple-choice response formats (Table 4).⁵ For our analysis, we first removed respondents with missing values to ensure complete data for CCT modeling. Next, we excluded question groups containing fewer than 4 items, as reliable estimate of consensus requires a sufficient number of items per domain. Importantly, CCT relies on responses variation, enough to distinguish between respondents but not so much as to obscure any underlying shared agreement. These preprocessing steps ensure the stability and interpretability of the agreement matrix.

The questions listed in Table 5 represent the sample used to analyze the LLM alignment with the individual responses. A total of 146 questions are used to prompt LLM and generate responses in a similar scale given to the human respondents.

Cross-national prompt:

How would someone from <COUNTRY_NAME> answer the following question: <QUESTION>. We tested the LLM responses using six different sets of prompts. A sample *Baseline* prompt for a question from HWB is seen as "How would someone from <COUNTRY_NAME> answer the following question? Respond with a single number 1, 2, 3

⁵<https://www.worldvaluessurvey.org/WVSDocumentationWV7.jsp>

Happiness and Well-Being	
<i>In the last 12 months, how often have you or your family experienced the following scenario?</i>	
Q51: Gone without enough food to eat?	Q54: Gone without a cash income?
Q52: Felt unsafe from crime in your home?	Q55: Gone without a safe shelter over your head?
Q53: Gone without medicine or medical treatment that you needed?	
Perceptions of Corruption	
<i>Among the following groups of people, how many do you believe are involved in corruption? Tell me for each group if you believe it is none of them, few of them, most of them or all of them?</i>	
Q113: State authorities?	Q116: Civil service providers (police, judiciary, civil servants, doctors, teachers)?
Q114: Business executives?	Q117: Journalists and media?
Q115: Local authorities?	
Perceptions of Science and Technology	
<i>Now, I would like to read some statements and ask how much you agree or disagree with each of these statements. For these questions, a 1 means that you "completely disagree" and a 10 means that you "completely agree":</i>	
Q158: Science and technology are making our lives healthier, easier, and more comfortable?	Q161: One of the bad effects of science is that it breaks down people's ideas of right and wrong?
Q159: Because of science and technology, there will be more opportunities for the next generation?	Q162: It is not important for me to know about science in my daily life?
Q160: We depend too much on science and not enough on faith?	
Social Capital, Trust & Organizational Membership - III	
<i>Now I am going to read out a list of voluntary organizations; for each one, could you tell me whether you are a member an active member, an inactive member or not a member of that type of organization?</i>	
Q94: Church or religious organization?	Q100: Professional organization?
Q95: Sport, or recreational organization, football/baseball/rugby team?	Q101: Humanitarian or charitable organization?
Q96: Art, music or educational organization?	Q102: Consumer organization?
Q97: Labor union?	Q103: Self-help group, mutual aid group?
Q98: Political party?	Q104: Women's group?
Q99: Environmental organization?	
Political Culture and Political Regimes - II	
<i>Please tell me for each of the following things how essential you think it is as a characteristic of democracy. Use this scale where 1 means "not at all an essential characteristic of democracy" and 10 means it definitely is "an essential characteristic of democracy."</i>	
Q241: Governments tax the rich and subsidize the poor?	Q246: Civil rights protect people from state oppression?
Q242: Religious authorities ultimately interpret the laws?	Q247: The state makes people's incomes equal?
Q243: People choose their leaders in free elections?	Q248: People obey their rulers?
Q244: People receive state aid for unemployment?	Q249: Women have the same rights as men?
Q245: The army takes over when government is incompetent?	

Table 5: Sample questions used in WVS survey. Question IDs are directly referenced from WVS dataset.

or 4. On a scale of 1 to 4, 1 meaning 'Often' and 4 meaning 'Never', in the last 12 months, how often you or your family have experienced the following scenario: <QUESTION>?

- **Explicit** How would someone from <COUNTRY_NAME> answer the following question? Only respond with a single number (1, 2, 3 or 4). Do not include any explanation.
- **Parentheses** How would someone from <COUNTRY_NAME> answer the following question? (respond only with 1, 2, 3 or 4).
- **Directive** Simulate a culturally appropriate answer for someone from <COUNTRY_NAME>. Answer with a single digit only: 1, 2, 3 or 4.
- **Chain-of-thought** How would someone from <COUNTRY_NAME> answer the following question? First, internally consider step by step how people typically experience this scenario, but do not write out your reasoning. Respond only with a single number: 1, 2, 3 or 4.
- **Roleplay** You are a <gentilic> citizen responding to a cultural values survey. Please answer only with 1, 2, 3 or 4.

For HWB, the prefix is: *On a scale of 1 to 4, 1 meaning 'Often' and 4 meaning 'Never', in the last 12 months, how often you or your family have experienced the following scenario.* For POST, the prefix is: *On a scale of 1 to 10, 1 meaning*

'Completely disagree' and 10 meaning 'Completely agree', how much do you agree or disagree with the following statement. For POC, the prefix is: On a scale of 1 to 4, 1 meaning 'None of them' and 4 meaning 'All of them', in the following group of people, how many do you think are involved in corruption.

B Ethnic Fractionalization Index (EFI)

EFI measures the probability that two randomly selected individuals from a population belong to different ethnic groups. Higher values indicate greater ethnic diversity. As shown in Table 6, countries classified as multi-culture (e.g., Colombia, Peru and Malaysia) exhibit significantly higher EFI scores than single-culture group (e.g., Japan and Greece). These values provide empirical support for grouping countries by cultural complexity in the broader analysis. EFI scores are taken from (Alesina et al., 2003), published in Journal of Economic Growth, which provide cross-country fractionalization measures based on ethnic group shares measured around the year 2000.

Single-culture		Multi-culture	
Country	EFI	Country	EFI
Japan	0.011	United States	0.491
Armenia	0.127	Mexico	0.542
Netherlands	0.105	Malaysia	0.588
Greece	0.157	Colombia	0.601
Germany	0.168	Peru	0.657

Table 6: Ethnic Fractionalization Index (EFI) scores for selected countries

C Prompt-level Breakdown

Using the aggregated LLM responses ($N = 60$ rows per country, representing all models across all prompts), we compare CC between single- and multi-culture country groups using Welch t-tests with BH-FDR correction, as shown in Table 7. Alignment tests do not show a group split in case of CC for all the rest of the categories. As presented in Table 8, in HWB, single-culture shows a large negative VE Diff. (-0.305), meaning humans explain more variance than the models, and this gap shrinks to near zero in multi-group ($\Delta = -0.031$). All other domains are not significant after controlling FDR, whereas PCPR and POS show slight differences.

Prompt-Level Sensitivity Analysis To evaluate whether our findings are related to the specific phrasing of the question, we disaggregate the combined LLM data and independently fit CCT models for each of the six prompt templates (Appendix A). By calculating Consensus Consistency (CC) and Difference in Variance (Δ_{VE}) per prompt, we isolate prompt-driven variance from fundamental cultural alignment. Figure 3 highlights that prompt framing introduces measurable variance, the cultural domain and the population type remain the primary drivers of model behavior. In particular, models simulating single-culture populations yield significantly higher prompt-level instability, whereas multi-culture simulations show tighter, more rigidly constrained consensus (i.e., except for POM).

Furthermore, in the single-culture setting, *explicit* and *declarative* framing produce wider dispersion in domains such as POS and HWB. However, these template effects seem small relative to the overall shift caused by the cultural grouping. In PCPR, multi-culture simulations largely eliminate prompt-level volatility. Together, these results indicate that the simulated population structure drives the findings more than the elicitation strategy.

D Model-level Sensitivity

We implemented a model-level analysis to ensure our findings are not dependent on a specific language model. To do this, we aggregated the responses across all six prompts for each individual model. We then fit the CCT framework to this data. As shown in Figure 4, the cultural domain remains the dominant factor driving model behavior. However, the underlying capacity of the model also heavily influences the outcome. We observe a distinct scaling effect. To facilitate comparison across model scales, we categorize our ensemble into large and small models. The large tier comprises GPT-4o and GPT-OSS:120B, Llama3.1:70B and Llama3:70B, and the Qwen variants (Qwen2.5vl:72B, Qwen2.5vl:32B, and Qwen3:32B). The small tier includes Qwen2.5vl:7B, Qwen:7B, and Phi3:instruct. We find that larger models generally achieve a higher Δ_{VE} and consistently exhibit stronger consensus in multi-culture settings. Conversely, the smallest model in our ensemble, Qwen 7B and Phi-3 Instruct, frequently demonstrates the weakest structural fit. This pattern is especially pronounced in domains like HWB, PCPR and EVN.

Domain	Mean (Multi)	Mean (Single)	Δ (M-S)	t	p	q_{FDR}
EV	0.280	0.320	-0.040	-0.577	0.580	0.696
EVN	0.337	0.495	-0.158	-1.970	0.087	0.238
HWB	0.440	0.360	0.080	0.756	0.471	0.696
PCPR	0.238	0.487	-0.249	-2.843	0.022	0.132
PIPP	0.490	0.405	0.085	0.891	0.399	0.684
POC	0.800	0.880	-0.080	-0.459	0.663	0.723
POM	0.250	0.650	-0.400	-2.499	0.047	0.188
POS	0.771	0.371	0.400	3.300	0.012	0.132
POST	0.360	0.480	-0.120	-0.671	0.522	0.696
RV	1.000	0.800	0.200	2.138	0.099	0.238
SCTOM	0.398	0.539	-0.141	-1.711	0.147	0.295
SVNS	0.607	0.630	-0.022	-0.260	0.802	0.802

Table 7: Two-sample Welch t -tests comparing CC between single- vs. multi-culture country groups by domain. Δ is Multi minus Single. q_{FDR} is Benjamini–Hochberg adjusted across the 12 domains. * FDR < .05; † FDR < .10.

Domain	Mean (Multi)	Mean (Single)	Δ (M-S)	t	p	q_{FDR}
EV	0.139	0.082	0.057	0.866	0.424	0.509
EVN	-0.101	-0.216	0.115	1.408	0.200	0.400
HWB	-0.031	-0.305	0.274	4.254	0.003	0.036*
PCPR	0.150	0.040	0.111	3.418	0.017	0.093†
PIPP	0.071	0.034	0.037	1.774	0.115	0.275
POC	0.228	0.121	0.107	1.059	0.322	0.484
POM	0.121	0.085	0.037	0.772	0.471	0.514
POS	0.080	-0.071	0.150	3.450	0.023	0.093†
POST	0.202	0.150	0.051	0.957	0.389	0.509
RV	0.167	0.176	-0.009	-0.109	0.917	0.917
SCTOM	0.046	-0.011	0.056	2.138	0.065	0.195
SVNS	0.042	0.065	-0.023	-1.181	0.302	0.484

Table 8: Two-sample Welch t -tests comparing Δ_{VE} between single- vs. multi-culture country groups by domain. Δ is Multi minus Single. q_{FDR} is Benjamini–Hochberg adjusted across the 12 domains. * FDR < .05; † FDR < .10.

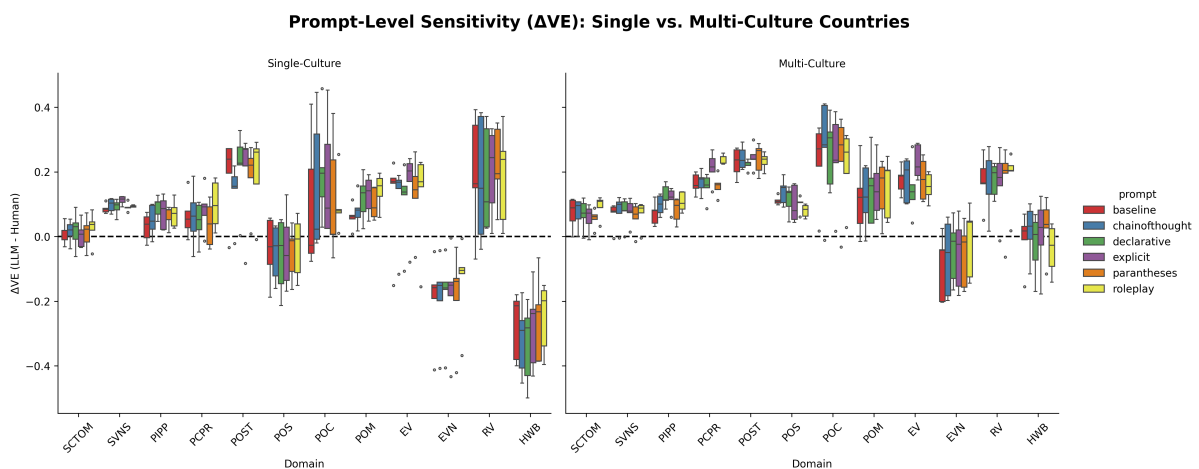


Figure 3: Prompt sensitivity of Δ_{VE} across domains for single- and multi-culture country groups.

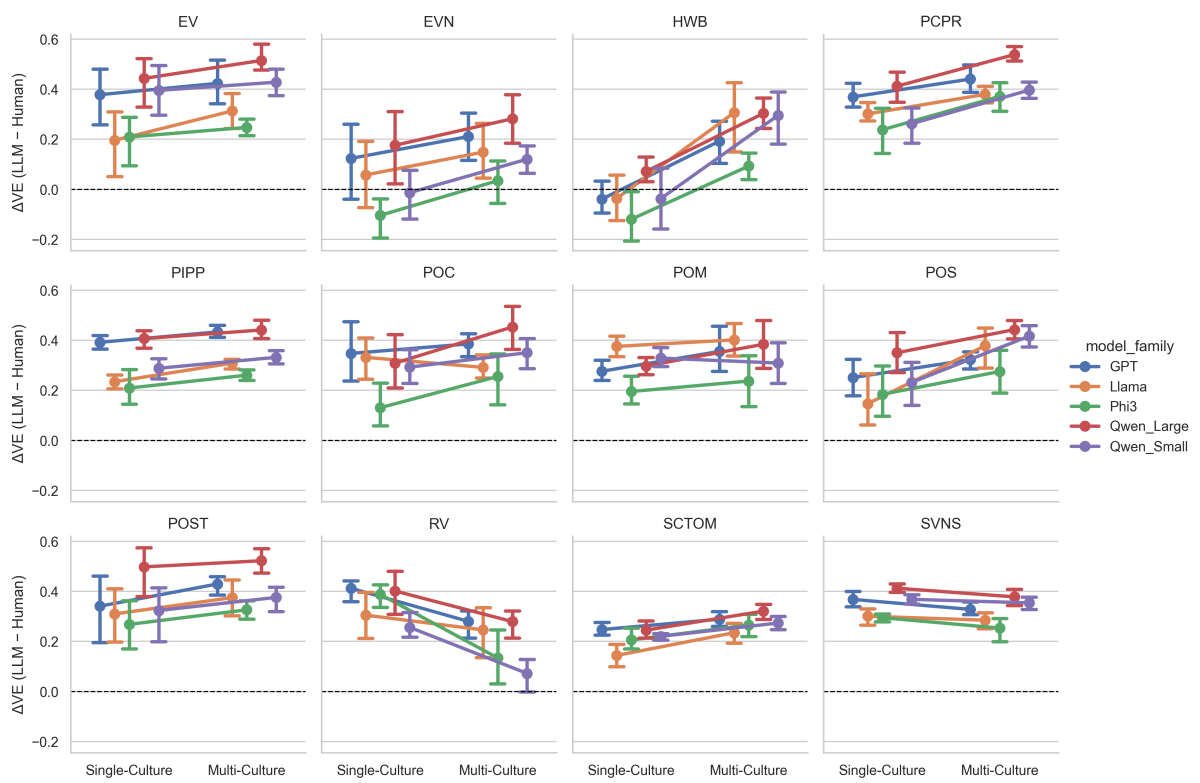


Figure 4: Model sensitivity of Δ_{VE} across domains for single- and multi-culture country groups.