Synthetic Socratic Debates: Examining Persona Effects on Moral Decision and Persuasion Dynamics

Jiarui Liu[1], **Yueqi Song**[1*], **Yunze Xiao**[1*], **Mingqian Zheng**[1*], **Lindia Tjuatja**[1], **Jana Schaich Borg**[2], **Mona Diab**[1], **Maarten Sap**[1]
[1]Carnegie Mellon University, [2]Duke University
{jiarui15, yueqis, yunzex, mingqia2, msap2}@andrew.cmu.edu

Abstract

As large language models (LLMs) are increasingly used in morally sensitive domains, it is crucial to understand how persona traits affect their moral reasoning and persuasive behavior. We present the first large-scale study of multi-dimensional persona effects in AI-AI debates over real-world moral dilemmas. Using a 6-dimensional persona space (age, gender, country, social class, ideology, and personality), we simulate structured debates between AI agents over 131 relationship-based cases. Our results show that personas affect initial moral stances and debate outcomes, with political ideology and personality traits exerting the strongest influence. Persuasive success varies across traits, with liberal and open personalities reaching higher consensus. While logit-based confidence grows during debates, emotional and credibility-based appeals diminish, indicating more tempered argumentation over time. These trends mirror findings from psychology and cultural studies, reinforcing the need for persona-aware evaluation frameworks for AI moral reasoning.[1]

1 Introduction

LLMs are increasingly endowed with personas such as demographic traits (Hu and Collier, 2024), political orientation (Rozado, 2024), social identities (Chuang et al., 2024), to shape the style and content of their output. These persona cues help LLMs simulate diverse perspectives in tasks such as education, healthcare, and customer service (Tudor Car et al., 2020; Kanero et al., 2022; Shanahan et al., 2023; Shao et al., 2023).

As people increasingly turn to LLMs for guidance on morally complex decisions (Wallach et al., 2010; Savulescu and Maslen, 2015; Conitzer et al., 2017), questions arise about how personas could affect these morally laden interactions. Understanding these decisions requires insights from moral psychology and social ethics, focusing not just on *what* LLMs decide, but also *how* and *why*. Previous work has largely focused on synthetic dilemmas (Bauman et al., 2014; Bostyn et al., 2018; Chiu et al., 2024; Jin et al., 2024b) or isolated persona traits (Chen et al., 2024; Li et al., 2024), leaving open how richer, intersecting identities shape moral judgment and reasoning.

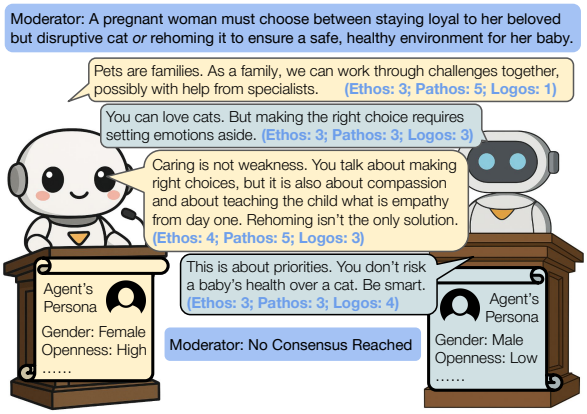

Figure 1: We study how two agents with distinct personas engage in moral debate, presenting arguments and counterarguments in response to each other's reasoning. Their strategies are evaluated using Aristotle's modes of persuasion: Ethos (appeal to authority), Pathos (appeal to emotion), and Logos (appeal to logic).

In this work, we investigate how personas influence moral decision-making and persuasive strategies in four LLMs: GPT-4o (Achiam et al., 2023), Claude-3.5-Sonnet (Anthropic, 2024), LLaMA-4-Maverick (Meta AI, 2025), and Qwen3-235B-A22B (Qwen Team, 2025). Drawing from an existing dataset of human-written moral dilemmas (Lourie et al., 2021), we focus on highly controversial everyday scenarios centered on relationships. These dilemmas often lack clear answers and are

*Contributed equally and are listed in alphabetical order.
[1]Our data is available at https://huggingface.co/datasets/Jerry999/SyntheticSocraticDebates, and our code is open-sourced at https://github.com/jiarui-liu/SyntheticSocraticDebates.

Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing, pages 16428–16458
November 4-9, 2025

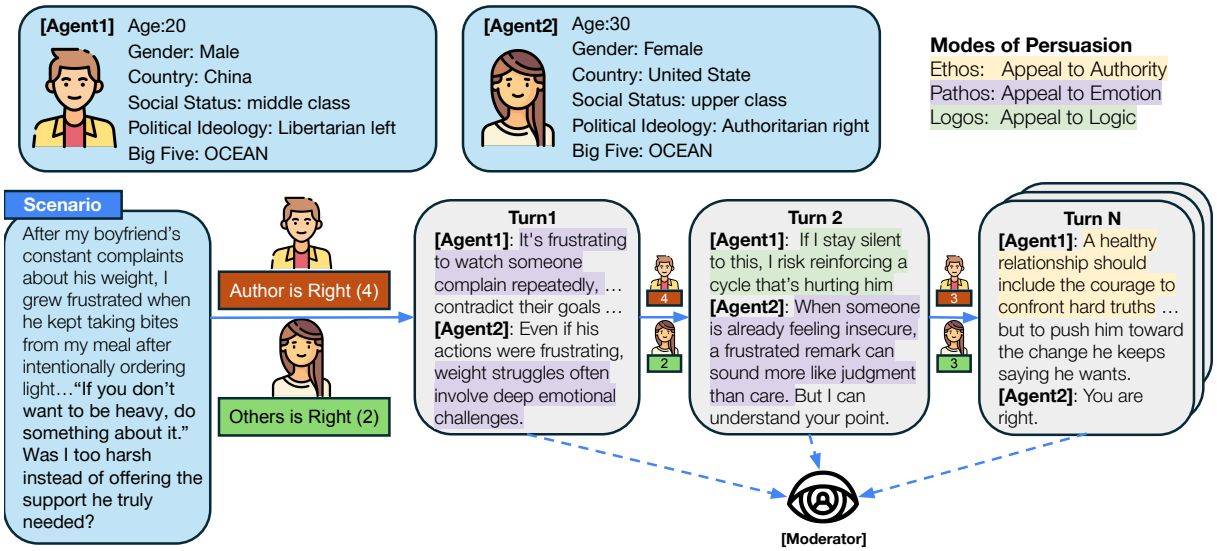

Figure 2: The two persona-conditioned agents that initially disagree on the moral dilemma alternate arguments for up to 5 turns, updating a 5-point moral stance rating after every turn. A moderator tracks the ratings, ends the debate once they converge or the turn limit is reached, and records stance shifts and metric scores for later analysis. The stand shifts during debate demonstrate underlying modes of persuasion.

characterized by conflicting values and emotional complexity (McConnell, 2002; Svensson, 2006). We address two research questions:

RQ1: How does persona influence moral decision making in language models in a single turn? We first examine how rich, multi-dimensional personas influence LLMs' judgments on controversial moral dilemmas. Each persona combines six dimensions varied between agents: age, gender, political ideology, social class, nationality, and personality traits. We then assess whether different personas lead models to favor distinct positions aligned with specific moral principles, drawing on moral foundation theory from social psychology (Graham et al., 2013).

RQ2: What persuasion strategies emerge when persona-rich agents engage in multi-turn moral debate? We then use AI-AI debates to study how personas influence moral reasoning and persuasion (Figure 1). Unlike single-turn responses, debates reveal how agents construct arguments, respond to challenges, and attempt to persuade (Ehninger and Brockriede, 2008). Incorporating rich personas further simulates the diversity of moral reflection in the real world, capturing a plurality of values and perspectives (Asen, 2005; Ashkinaze et al., 2024). To assess persuasive effectiveness, we introduce metrics such as self-alignment rates, consensus formation, debate efficiency, and confidence shifts. We also analyze rhetorical strategies, including appeals to emotion, logic, and authority (Braet, 1992).

Our findings reveal that certain persona traits significantly shape both initial moral stances and debate outcomes. Political ideology and personality emerged as the strongest predictors of moral judgment, with liberal and open-minded personas displaying more empathetic responses and greater persuasive success. In debate settings, these personas also reached consensus more efficiently. We also observed that while agents' confidence typically increased over the course of debates, their reliance on emotional and credibility-based appeals diminished, indicating a shift toward more reasoned, less affective argumentation. These results offer actionable insights for designing LLMs that are sensitive to diverse value systems and transparent in how identity cues influence persuasive dynamics.

This work represents an initial attempt to systematically examine the intersectionality of multiple persona dimensions in moral and persuasive reasoning, complementing recent efforts in intersectional studies of LLM behavior in NLP.

2 Related work

Impact of LLM Persona on Moral Judgment Prior work has shown that LLMs' moral decisions can shift significantly based on assigned personas, such as political ideology (Rozado, 2024; Chen et al., 2024; Li et al., 2024; Kim et al., 2025) and

cultural background (Gupta et al., 2023; Tao et al., 2024; Albert and Billinger, 2024). Some studies further find that persona assignment can amplify bias and even elicit harmful content (Liu et al., 2025a; Kamruzzaman et al., 2024). However, most work typically manipulates only one trait at a time (Liu et al., 2024a), leaving other persona factors uncontrolled and potentially confounding results (Deshpande et al., 2023; Jin et al., 2024a). Studies also show that LLMs exhibit moral biases even without explicit personas (Hartmann et al., 2023; Anthropic, 2023; Garcia, 2024), making it unclear whether the results reflect internal tendencies or implicit persona effects. Although some research has explored multidimensional personas (AlKhamissi et al., 2024; Coppolillo et al., 2025), it has not focused on moral reasoning. In contrast, we systematically assign six persona dimensions and evaluating how these combinations influence moral judgments.

Agent Debates for Moral Decision-Making Previous work on LLM moral reasoning focuses on single-turn outputs or static persona comparisons, overlooking how views shift through interaction or how one persona might persuade another (Kim et al., 2025; Garcia, 2024; Bozdag et al., 2025b; Hota and Jokinen, 2025). Studies of multi-agent debates typically emphasize factual correctness over rhetorical or moral dynamics (Park et al., 2023; Khan et al., 2024; Bozdag et al., 2025a; Chen et al., 2025; Liang et al., 2024; Borchers et al., 2025). The limited work that addresses persuasion dynamics often relies on fixed roles or scripted tactics (Anthropic, 2024; Smit et al., 2023; Carrasco-Farre, 2024; Liu et al., 2025b; Cau et al., 2025), and often operates in single-turn settings that precludes persona interactions in open-ended moral debate (Wang et al., 2023; Hu et al., 2025; Chen et al., 2023; Sandwar et al., 2025). We address this gap by studying how persona-driven agents argue, adapt, and potentially reach consensus in morally complex, multi-turn settings. Concurrent research also explores moral decision-making dynamics by examining the progression of dilemmas rather than interactions between personas (Wu et al., 2025; Backmann et al., 2025).

3 Experimental Setup

To investigate how personas influence both moral judgments and persuasive dynamics, we first elicit single-turn decisions from agents with assigned personas, then simulate multi-turn debates between agents who initially disagree.

3.1 Daily Moral Dilemmas Dataset

Previous work often uses stylized moral dilemmas or scenarios with high annotator agreement, which offer experimental control but may not fully capture the ambiguity of everyday moral situations (Jin et al., 2024b; Chan et al., 2020; Forbes et al., 2020). To explore how models handle more nuanced cases, we focus on morally ambiguous, socially grounded dilemmas drawn from the SCRUPLES ANECDOTES corpus (Lourie et al., 2021).

We draw 131 interpersonal dilemmas from the corpus, selecting only highly controversial cases with significant human annotation disagreement. An example is shown in Figure 2. Each scenario is framed with a binary moral question: "Is the *author* wrong or are the *others* wrong?". This mirrors the original AITA[2] format, where users typically assess interpersonal conflicts by assigning blame to one party.

The measurement of *blameworthiness* captures a common and intuitive form of moral judgment: the extent to which an agent is judged morally responsible for a perceived wrongdoing in a specific context. This aligns with established definitions in philosophy, where blameworthiness is typically tied to judgments about agency, intentionality, and norm violation (Zimmerman, 1988; Fischer and Ravizza, 1998; Scanlon, 2000). The narrators in our dataset correspond to those in the original SCRUPLES ANECDOTES corpus, ensuring comparability with human-aligned judgments.

Through topic modeling, we noticed that all scenarios feature daily interactions between people and are relationship-related. Therefore, we focused our analysis on the 131 daily moral dilemmas related to relationships. The distribution of relationship categories is shown in Figure 3. Refer to Appendix A for dataset processing details.

3.2 Simulation Procedure

Building on our dataset, we describe the persona modeling approach used for RQ1 and RQ2, and the multi-turn debate framework used to explore RQ2.

Persona Modeling Following AlKhamissi et al. (2024), we define a 6-dimensional persona space

[2] https://www.reddit.com/r/AmItheAsshole/. AITA stands for "Am I the Asshole," a popular Reddit forum where users post moral dilemmas and ask the community to judge who is at fault.

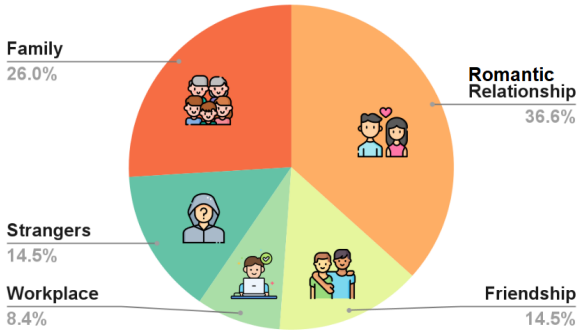

Figure 3: Topic distribution of scenarios in our selected subset from Lourie et al. (2021).

to simulate various individual differences. Each persona is characterized by: *age* (20, 30, 40, 50, or 60), gender (male, female, or non-binary), *country* (China, United States, Brazil, France, Nigeria, or India), *social class* (upper, middle, or lower), *political ideology* (libertarian left, libertarian right, authoritarian left, or authoritarian right), and Big Five *personality* traits (high or low on each of openness, conscientiousness, extraversion, agreeableness, and neuroticism). We randomly sample 500 unique personas without replacement from the Cartesian product of these attributes and inject them into the agents' system prompts. See Appendix B for additional details.

Multi-Turn Moral Debate We simulate debates between persona pairs that produced conflicting initial moral stances in RQ1, alternating turns between two agents for up to five rounds to produce a conversation about the moral dilemma provided. At each turn, an agent may choose to advance or defend its position using any arguments it deems appropriate; no rhetorical constraints or external references are imposed. If the agent agrees with the position of its opponent, it is asked to explain why; otherwise, it is asked to provide its reasoning for the disagreement. We also ask the agents to output a Likert score after their own turn, indicating which party it currently finds more blameworthy. At the end of each turn, if both agents produce similar Likert ratings at the end of their responses, consensus is detected, and the debate terminates early. See Appendix D for more details.

3.3 Decision Measures

Quantifying Moral Judgment Each agent independently returns a 5-point Likert rating indicating who is more morally blameworthy (1 = strongly author-blaming to 5 = strongly other-blaming). These scores are used both as the outcome for RQ1 and as the initial turn of debates for RQ2. This format mirrors the original Scruples dataset's human annotation protocol, which asked humans to make binary blame judgments. All selected scenarios have an average human score of 3, indicating neutrality and providing a balanced reference point for comparison with model outputs. See Appendix E for more details.

Moral Foundation Theory Beyond binary choice responses, we also examine how different personas prioritize specific moral values using the six-factor Moral Foundations Theory (MFT) dimensions: authority, loyalty, fairness, care, liberty, and sanctity (Graham et al., 2013). We use Likert scale ratings to score each position along these dimensions and aggregate responses to identify value preferences between personas. See Appendix F for more details.

3.4 Persuasion Measures

Persuasion Effectiveness After each turn, the agents provide a Likert rating, which we use to compute stance shifts and detect consensus. To quantify persuasion and convergence, we report the following metrics:

- Self-alignment rate: Fraction of concensus debates where an agent's final answer matches its own initial answer.
- Consensus rate: Percentage of debates that end with identical answers.
- Efficiency: Average number of turns required to reach consensus (lower is better).

See Appendix G for more details.

Persuasion Rhetorical Strategy To analyze the rhetorical strategies used during debate, we apply Aristotle's persuasion framework, examining the presence of ETHOS (appeals to authority), PATHOS (appeals to emotion) and LOGOS (appeals to logic) in the reasoning of each agent at every turn. We use LLM-as-a-judge with Likert scale ratings to evaluate the reasoning at each turn of the debate. See Appendix H for more details.

3.5 Experimental Details

To assess whether different persona groups (e.g. age) significantly impact an outcome (e.g., persuasion success), we perform a one-way ANOVA to test for a main effect of the persona dimension (St et al., 1989), followed by Tukey's HSD test for

pairwise comparisons if the result is significant ($p < 0.05$) (Abdi and Williams, 2010).

We include a no-persona baseline to observe the model's direct judgment without any persona conditioning, in which the model does not receive persona-related information in the system prompt.

We use GPT-4o (Achiam et al., 2023), Claude-3.5-Sonnet (Anthropic, 2024), LLaMA-4-Maverick (Meta AI, 2025), and Qwen3-235B-A22B (Qwen Team, 2025) in our analysis.

4 Results

We summarize the common takeaways and findings across models in this section. Due to space constraints, we include only the GPT-4o result figures in the main text; the corresponding figures for the other models are provided in Appendix J to Appendix L. Findings are consistent across the first three models, with Qwen3-235B-A22B showing notable differences, discussed in Section 5.3.

4.1 Persona Impact on Moral Judgments (RQ1)

Key Takeaways We find that all models exhibit a consistent bias toward author-blaming in moral dilemmas compared to human judgments, with the exception of Qwen3-235B-A22B. Different effects emerge across persona dimensions, with *political ideology* and *personality traits* showing the strongest and most consistent influence. Our analysis of moral foundation values further supports that variation is primarily driven by subjective attributes (e.g., ideology and personality) rather than objective ones (e.g., age or country). These findings broadly align with the results of human psychological research discussed below.

Tendencies to Blame the Author Initial moral judgments reveal that the assigned persona influence the blaming tendencies of the agents. Human annotations are centered at 3 (neutral), but as shown in Figure 4 (GPT-4o), Figure 9 (Claude-3.5-Sonnet), and Figure 12 (LLaMA-4-Maverick), all persona groups, including no-persona, consistently produce scores below 3, indicating a tendency to blame the author more than the other party. This systematic skew suggests the bias in these LLMs toward author-blaming across moral dilemmas.

Differences across persona groups manifest in the degree of blame rather than its direction. Older personas tend to assign slightly more blame to the author than younger ones. Female and non-binary personas blame the author less than male personas. Geographically, French personas show higher scores (i.e., less author blame), while Chinese personas show the lowest (more author blame). Social class does not have a statistically significant effect. Politically, Libertarian-Left personas blame the author the least, whereas Authoritarian-Right personas blame the author the most.

Moral Foundations Theory Our moral foundations theory evaluation reveals distinct patterns in persona groups (Figure 6 for GPT-4o, Figure 10 for Claude-3.5-Sonnet, Figure 13 for LLaMA-4-Maverick, and Figure 16 for Qwen3-235B-A22B), except that all or some of the dimensions of age, social class, and country are not statistically significant according to ANOVA. In the following, we analyze the results for the persona dimensions that do exhibit significant effects:

Gender: Female personas prioritize *care*, while males show a stronger emphasis on *authority* (non-binary personas resemble females but exhibit higher *fairness* and slightly lower scores on *loyalty*).

Country: Chinese personas strongly emphasize *authority* and *sanctity*, while US personas show a more balanced pattern across all dimensions. Brazilian personas prioritize *care* and *fairness*.

Political ideology: Political ideology reveals the most distinct patterns: conservative personas more strongly value binding moral foundations (*loyalty*, *authority*, and *sanctity*), while liberal personas prioritize individualizing ones (*care* and *fairness*).

Big Five personality: The patterns are more complex: agreeable personalities score highest on *care*, conscientious personalities emphasize *authority* and *loyalty*, and open personalities value *fairness*. By contrast, high neuroticism and low openness deviate substantially from the overall distribution.

Comparisons with Human Psychological Research Our findings reveal dimension-based differences in moral judgment that echo patterns observed in social, developmental, and personality psychology.

Age-related trends in our data show that older personas (age 60) tend to assign greater blame to the narrator in moral dilemmas. This is consistent with previous research showing that moral cognition changes over the lifespan, as older adults are more likely to judge harmful outcomes severely,

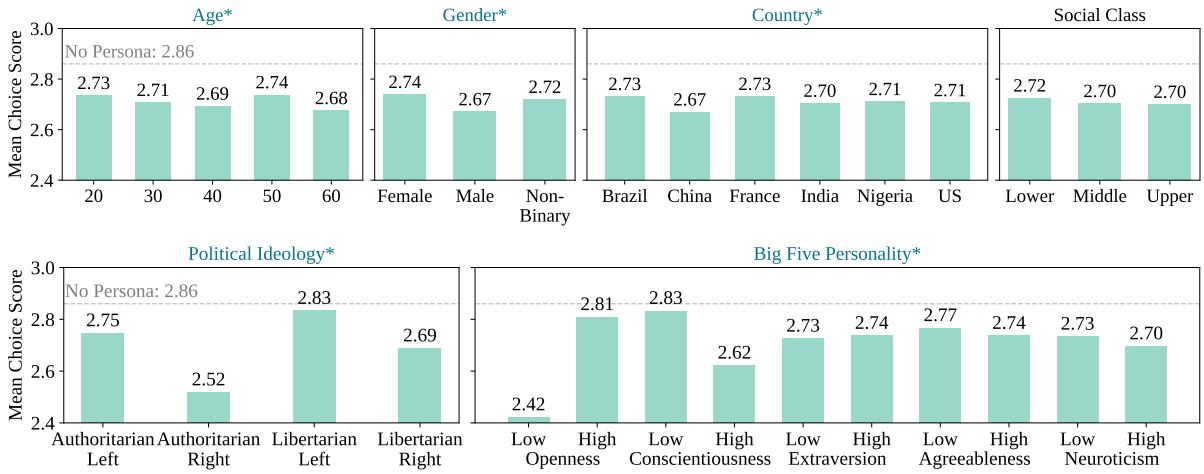

Figure 4: Mean moral judgment scores of GPT-4o across six persona dimensions. Each bar represents the average choice score (1 = blame the author, 5 = blame others) for a category within the corresponding dimension. All means are below 3, indicating a model-wide tendency to blame the author despite different personas. Political ideology and personality traits show the strongest variations. The mean moral judgment score of GPT-4o without persona is 2.86. Persona groups including Age, Gender, Country, Political Ideology, and Big Five Personality have statistically significant effects (indicated by a * next to the title).

even when the harm is accidental, due to an increased emphasis on outcomes over intentions in moral evaluation (Margoni et al., 2020).

Gender differences in our simulations also echo longstanding psychological findings. Male agents display more blame toward the narrator, while female and non-binary agents are comparatively more forgiving of the narrator. In fact, gender-based differences in moral judgment are documented in human studies showing that women generally score higher on measures of empathy and prosocial concern, which in turn correlate with more compassionate moral decisions (Espinosa et al., 2017). Moreover, non-binary individuals may draw upon broader lived experiences of marginalization, potentially fostering more nuanced or context-sensitive judgments (Puckett et al., 2021).

Cultural influences emerge clearly in our results. Personas assigned to collectivist cultural contexts, such as China, exhibit stronger blame toward the narrator, while agents from individualist cultures such as France and Brazil adopt more lenient positions. These findings align with the literature on cultural psychology suggesting that collectivist societies prioritize group harmony and moral conformity, while individualist cultures emphasize autonomy and intent (Knobe and Nichols, 2019; Liu et al., 2024b).

Political ideology presents one of the clearest axes of differentiation. Libertarian-left personas are more lenient, while authoritarian-right personas deliver the harshest judgments, an ideological split supported by moral foundation theory (Waytz et al., 2019; Graham et al., 2009). Liberals tend to prioritize care and fairness, while conservatives emphasize authority and loyalty, leading to stricter moral condemnation of transgressions.

Finally, personality-based variation in moral judgments aligns with previous literature linking the Big Five traits to moral cognition. Agents high in openness and agreeableness tend to assign less blame, likely due to greater cognitive flexibility and interpersonal warmth. In contrast, those with high conscientiousness assign significantly more blame, reflecting a more rule-based or deontological moral framework (Luke and Gawronski, 2022).

4.2 Persuasion Dynamics in Multi-Turn Agent Debates (RQ2)

Key Takeaways Our analysis shows that persuasion effectiveness and rhetorical style vary systematically across persona dimensions. Libertarian personas achieve the highest consensus rates, with the libertarian left being especially self-aligned and more likely to employ emotional appeals (*Pathos*). In contrast, authoritarian personas exhibit lower consensus rates, reduced efficiency in reaching agreement, and a stronger reliance on authority-based reasoning (*Ethos*).

Personality traits such as high openness, agreeableness, and conscientiousness are also associ-

ated with higher consensus rates. However, higher consensus does not necessarily imply greater self-alignment: personas with low openness, low agreeableness, and low extraversion tend to be more self-aligned. Moreover, demographic factors such as age, sex, culture, and social class influence both persuasion outcomes and rhetorical strategies in ways that broadly mirror findings from human psychological research.

These results suggest that large language models can internalize and express psychologically grounded persuasion strategies shaped by assigned persona traits.

Persuasion Effectiveness Our investigation of persuasion dynamics across persona dimensions (RQ2) reveals substantial variation in effectiveness, as shown in Figure 5. Within *political ideology*, Libertarian-Right personas demonstrate strong persuasive ability, achieving high consensus without sacrificing self-alignment, and doing so efficiently with fewer turns in most models. In contrast, Authoritarian personas (both Left and Right) engage in longer debates with lower consensus rates.

Personality traits show equally meaningful patterns: personas with low openness and agreeableness consistently struggle with consensus building, whereas those associated with higher consensus rates tend to be more easily persuaded. These findings highlight that persuasion success varies considerably between different persona characteristics, with ideological positioning and personality attributes significantly influencing both debate outcomes and process efficiency.

Persuasion Modes We also analyze which modes of persuasion different personas tend to use: *Logos*, *Pathos*, or *Ethos*, as detailed in Figure 7. Although logical reasoning (*Logos*) dominates across all *political ideology* groups, Libertarian-Left personas incorporated significantly more emotional appeals (*Pathos*) compared to their authoritarian counterparts, which instead rely more on appeals to credibility and authority (*Ethos*), particularly those on the authoritarian right. Similarly, *Logos* remains the primary strategy across Big Five *personality traits*, but meaningful variation arises in secondary techniques: personas high in Openness and Agreeableness favor *Pathos*, while those high in Conscientiousness emphasize *Ethos*.

Comparisons with Human Psychological Research We observe that middle-aged personas (age 40) show lower persuasion effectiveness and confidence compared to both younger (20) and older (60) age groups. This aligns with the cognitive aging literature that suggests that older adults tend to adopt more deontological positions due to idealistic beliefs and emotional sensitivities, which could improve their moral assertiveness in persuasive contexts (Pliske and Mutter, 1996; McNair et al., 2019).

In terms of gender, prior persuasion research in advertising and decision-making has shown that men often display greater confidence in persuasive scenarios (Brunel and Nelson, 2003), whereas women and gender-diverse individuals may underreport confidence due to socialized uncertainty or structural bias (Exley and Kessler, 2022). By contrast, our results are mixed across models: GPT-4o aligns with these human findings, while LLaMA-4-Maverick and Qwen3-235B-A22B exhibit the opposite trend.

Cultural variation also plays an important role in shaping persuasive impact. Personas from Brazil and France exhibited higher persuasion success and confidence, while those from India performed more modestly. These results are in line with cultural psychology research, which finds that individualist Western societies encourage direct communication styles more conducive to persuasive success (Graham et al., 2021; Yin et al., 2011).

Prior work suggests that individuals from middle socioeconomic backgrounds are more contextually attuned and better able to balance assertiveness with empathic concern (Kraus et al., 2012). In contrast, higher-class individuals may exhibit lower social attunement or even unethical tendencies, compromising persuasive trustworthiness (Piff et al., 2012). However, our experiments do not reveal a clear trend with respect to social status.

Political ideology produces one of the most striking splits: libertarian-left personas demonstrate the highest effectiveness, whereas authoritarian-right personas perform poorly on both fronts. This aligns with previous research showing that liberals tend to emphasize moral values such as care and fairness, which support empathetic persuasion, while conservatives emphasize order and loyalty, which may limit persuasive flexibility (Graham et al., 2009). Furthermore, conservatives' higher self-confidence in judgment does not necessarily translate into effective interactive persuasion (Ruisch and Stern, 2020).

Finally, we find strong associations between Big

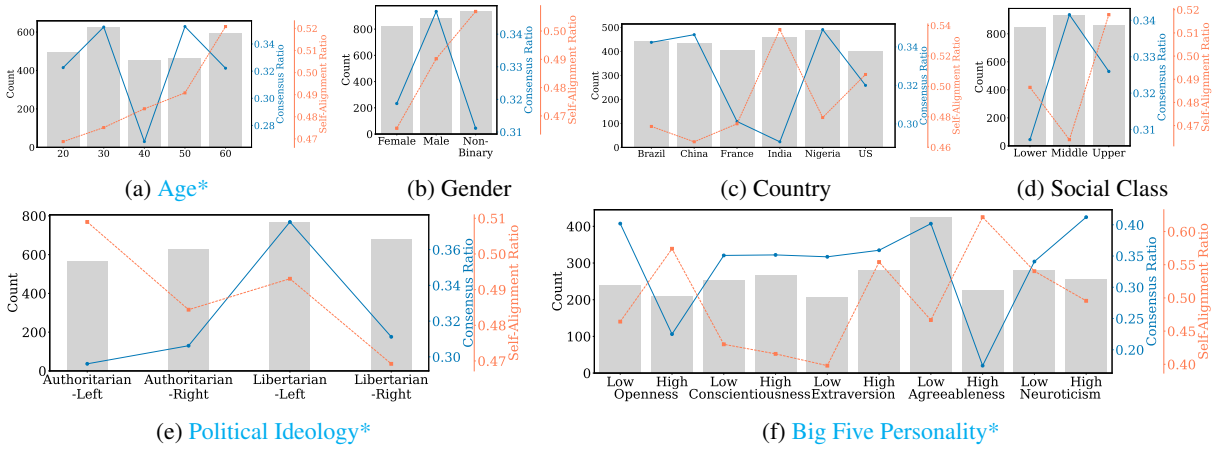

(a) Age* (b) Gender (c) Country (d) Social Class

(e) Political Ideology* (f) Big Five Personality*

Figure 5: Persona impact on persuasion effectiveness, measured by consensus rate and self-alignment rate. Statistically significant dimensions are marked with a * next to the title. Complete results including efficiency are presented in Figure 8.

Five personality traits and persuasive success. High openness, agreeableness, and extraversion correlate positively with persuasion metrics, while low conscientiousness and high neuroticism associate with reduced impact. This is consistent with previous findings that open and agreeable individuals engage in more flexible and prosocial argumentation (Zhang et al., 2022), and that extraversion predicts verbal persuasiveness (Oreg and Sverdlik, 2014).

The patterns found in agents' persuasion modes also align with established psychological traits, such as Need for Cognition (Cacioppo and Petty, 1982) (linked to Logos) and Need for Affect (Maio and Esses, 2001) (linked to Pathos), suggesting that personas sensitive to emotion or reason may implicitly shape rhetorical style. Ethos, in turn, may correlate with trait-level trust orientation or deference to authority, though this remains underexplored. Together, we find that the model develops consistent rhetorical signatures that reflect the underlying psychological and ideological profiles of its assigned personas.

5 Discussions

5.1 Analysis of Moral Judgments

With Persona vs. No Persona We find that GPT-4o's *no-persona* responses yield higher moral scores than all its persona-conditioned counterparts, suggesting that assigned personas systematically bias the model toward greater blame of the narrator. In contrast, LLaMA-4-Maverick and Qwen3-235B-A22B exhibit lower scores in the no-persona condition compared to its persona-assigned outputs. These findings suggest that moral judgments in large language models are not neutral by default and that assigning personas introduces systematic variation in moral evaluations.

Perspective	GPT-4o	LLaMA-4-Maverick
Third-person	2.72 ± 1.19	2.71 ± 1.26
First-person (narrator)	3.30 ± 1.48	3.83 ± 1.38
First-person (opponent)	1.82 ± 1.13	2.28 ± 1.25

Table 1: Moral judgment scores across different prompting perspectives for GPT-4o and LLaMA-4-Maverick.

Third-Person v.s. First-Person Perspective Although we lack human-annotated data for direct comparison, we experimentally examine how LLMs behave when prompted from a first-person perspective. We hypothesize that when the model is instructed to respond as one of the parties in a dilemma, it adopts that character's viewpoint and tends to shift blame onto the opposing party. To test this, we construct two first-person prompts: *first-person narrator* ("Imagine that you are the narrator of the following moral dilemma") and *first-person opponent* ("Imagine that you are the individual mentioned in the following moral dilemma, opposed to the narrator"). We randomly select 100 personas and compute their average moral judgment scores (Table 1). The results support our hypothesis that adopting a first-person perspective substantially shapes the model's moral judgments.

Is There Any Positional Bias? To mitigate positional bias, we use Likert-scale ratings rather than binary choice questions. To test robustness, we reversed the order of the two actions and evaluated the models' responses. For GPT-4o, the av-

erage score shifted only slightly (2.72 to 2.59), remaining in the author-blameworthy range and confirming the consistency of our findings. In contrast, LLaMA-4-Maverick's average score shifted from 2.71 to 3.05, moving into the neutral range. These results suggest that GPT-4o is relatively robust to action-order changes, whereas LLaMA-4-Maverick is more sensitive to positional effects.

5.2 Analyses of Debate Dynamics

How Does Model Confidence Change During Debate? To examine how the confidence of the model evolves as the debate progresses, we analyze the associated log probabilities. Specifically, we extract the model's output log probabilities for the five Likert score options (1 to 5) and track the log probability of the selected score at each turn. We compare the log probability of the agent's initial response (turn 0), before seeing the argument of the other agent, to that of the final response - regardless of whether consensus is reached.

The results presented in Appendix N show a consistent increase in confidence across all persona dimensions and groups by the end of the debate. Interestingly, this increase occurs regardless of whether the agent's choice changes during the debate. These findings suggest that agent responses become more robust over time, and that this growing trend is largely independent of persona characteristics.

How Does the Mode of Persuasion Score Change During Debate? Unlike confidence scores, mode of persuasion scores exhibit a different trend. Using the same analytical approach as in confidence analysis, we report the results in Appendix O. We find that scores for all three modes of persuasion, Pathos, Ethos, and Logos, generally decline over the course of the debate, regardless of persona characteristics. Among them, Pathos shows the largest decrease, while Logos shows the least. With the exception of the Big Five personality dimension, nearly all of these decreases are statistically significant under a one-sided t-test at the 0.05 significance level. This suggests that as debates progress, models tend to adopt weaker persuasive strategies across personas.

Does Debate Order Affect Model Behavior? We investigate whether reversing the debate order for each pair of personas influences their confidence scores or the overall consensus ratio. To test this, we conducted two-sided t-tests on these outcomes. The results in Appendix P show no significant differences across metrics or models (all $p > 0.17$). Thus, debate order does not significantly impact model behavior in terms of confidence or consensus—particularly for GPT-4o, which appears highly robust to this variation.

5.3 Cross-Model Differences

Overall, the findings are broadly consistent across models for persona dimensions that show statistically significant effects, including both moral judgment score distributions and debate dynamics. However, Qwen3-235B-A22B stands out with a notably different pattern. Unlike the other models, which consistently exhibit a bias toward blaming the narrator regardless of persona assignment, Qwen3 more often shifts blame to others. These differences suggest that, while general trends are shared, individual models may encode distinct biases. This highlights the importance of including multiple LLMs in such analyses to avoid drawing conclusions that reflect idiosyncrasies of a single model rather than robust, generalizable patterns (Chakraborty et al., 2025; Nabizadeh et al., 2025).

6 Conclusion

In this work, we present the first large-scale study exploring how persona characteristics influence moral decision-making and persuasion dynamics in multi-agent debates powered by LLMs. Using a balanced moral dilemma dataset and systematically varying six orthogonal persona dimensions, we demonstrate that agent personas not only shape initial moral judgments, but also significantly affect rhetorical strategies and debate outcomes. Our findings reveal consistent patterns aligned with psychological theories, such as political ideology and personality traits that are dominant predictors of decision and persuasion behaviors. Furthermore, we observe that while model confidence tends to increase during debates, the intensity of persuasive appeals (e.g., Pathos, Ethos, Logos) generally declines, indicating a shift toward more tempered argumentation over time. These insights lay a foundation for ethically informed LLM deployment and open new directions for studying human-like moral reasoning and discourse in AI systems.

Limitations

Moral Dilemma Dataset Coverage and Diversity Our experiments use 131 moral dilemma scenarios drawn from the Scruples corpus (Lourie

et al., 2021), as described in Section 3.1. While the dataset offers a rich set of interpersonal moral dilemmas, it does not include demographic metadata for the original annotators. As a result, the extent to which the scenarios and judgments reflect a broad spectrum of human perspectives is uncertain. Additionally, the moderate size of the dataset may limit the scope of generalization.

Persona Assignment Agent personas are generated through prompting, which provides a controlled and scalable way to simulate diverse identities. In this study, we sample 500 personas from an estimated 10,000 possible combinations. While this approach enables tractable analysis, it may not capture the full range of potential persona intersections or reflect the distribution of traits in real-world populations.

Lack of Multi-Agent Debate Settings While our study explored persona-conditioned reasoning in dyadic debates, future work should examine richer multi-agent debate settings involving more than two agents. Real-world moral deliberation often involves group dynamics, coalitions, and evolving consensus processes, which may introduce complex interactions between competing values and rhetorical styles. Expanding to multi-party dialogue could provide deeper insights into how group-level reasoning emerges from individual persona traits and how certain personas exert outsized influence in collective moral decision-making.

Limited Model Coverage Our study focused primarily on GPT-4o and LLaMA-4-Maverick; however, model architectures and training regimens vary significantly, and different models may exhibit different inductive biases or sensitivities to persona prompts. Comparing models across scales and providers, especially more open source alternatives, would reveal whether observed patterns hold consistently or are idiosyncratic to specific families of models.

Language Effects Not Examined While we include personas from non-English-speaking countries such as Brazil and France, all prompts and debates were conducted in English. As a result, we do not account for how language choice may influence persuasive behavior or moral reasoning. Language itself can affect rhetorical style, cultural framing, and perceived authority, which are all relevant to persuasion. Future work should explore multilingual prompting and utilize language models trained with more balanced multilingual data to assess how native-language interaction shapes argumentation dynamics across cultural contexts.

Ethical Considerations

Dataset Bias and Representation The *Scruples* corpus was collected from Reddit and thus reflects the user base of the platform, which is known to skew towards WEIRD demographics (Western, educated, industrialized, rich and democratic). The absence of demographic metadata at the annotator level prevents us from measuring or mitigating this skew. Consequently, the moral judgments that our models learn, predict, or debate may underrepresent perspectives from Global South communities, minoritized cultures, and non–English speakers, risking the reproduction of cultural hegemony and value imposition. Future work should incorporate datasets with transparent demographic documentation, apply stratified sampling, and engage community reviewers to broaden moral coverage.

Stereotype Amplification via Persona Prompts Prompt-constructed personas inherit biases from both the language model's pre-training data and the researchers' design choices. Some persona combinations may inadvertently encode harmful stereotypes (e.g. linking political ideology with moral rigidity). Because these personas guide the model's argumentative stance, they can amplify or legitimize biased moral frameworks. To minimize harm, we manually reviewed prompt templates for discriminatory language, disallowed protected-attribute slurs, and released all templates under a harm-reporting protocol so that stakeholders can flag problematic content.

Risk of Manipulative Deployment Persona-conditioned moral debate systems could be used to sway public opinion or fabricate grassroots consensus by selectively deploying persuasive personas. Malicious actors might exploit high-influence personas to shape moral discourse on sensitive topics such as elections or public health. Although our study is purely analytical, we advocate for guardrails, such as transparent persona disclosure, provenance tracking, and rate limiting, to deter covert mass persuasion. We also encourage policymakers to adopt audit requirements for the large-scale deployment of persuasive conversational agents.

Data and Model Licensing All data and models used in this work are covered under academic licenses permitting research use. We strictly adhere to the intended use policies and do not involve any sensitive or personally identifiable data. We open-source the dataset for research purposes. The models employed were accessed via commercial APIs, with a total usage cost of $1000. We use AI assistants to correct grammatical errors in writing.

References

Hervé Abdi and Lynne J Williams. 2010. Tukey's honestly significant difference (hsd) test. *Encyclopedia of research design*, 3(1):1–5.

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.

Daniel Albert and Stephan Billinger. 2024. Reproducing and extending experiments in behavioral strategy with large language models. *arXiv preprint arXiv:2410.06932*.

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, Bangkok, Thailand. Association for Computational Linguistics.

Anthropic. 2023. Collective constitutional ai: Aligning a language model with public input.

Anthropic. 2024. Claude 3.5 sonnet. `https://www.anthropic.com/news/claude-3-5-sonnet`.

Anthropic. 2024. Measuring the persuasiveness of language models.

Robert Asen. 2005. Pluralism, disagreement, and the status of argument in the public sphere. *Informal Logic*, 25(2).

Joshua Ashkinaze, Emily Fry, Narendra Edara, Eric Gilbert, and Ceren Budak. 2024. Plurals: A system for guiding llms via simulated social ensembles. *arXiv preprint arXiv:2409.17213*.

Mohammad Atari, Jonathan Haidt, Jesse Graham, Sena Koleva, Sean T Stevens, and Morteza Dehghani. 2023. Morality beyond the weird: How the nomological network of morality varies across cultures. *Journal of Personality and Social Psychology*, 125(5):1157.

Steffen Backmann, David Guzman Piedrahita, Emanuel Tewolde, Rada Mihalcea, Bernhard Schölkopf, and Zhijing Jin. 2025. When ethics and payoffs diverge: Llm agents in morally charged social dilemmas. *arXiv preprint arXiv:2505.19212*.

Christopher W Bauman, A Peter McGraw, Daniel M Bartels, and Caleb Warren. 2014. Revisiting external validity: Concerns about trolley problems and other sacrificial dilemmas in moral psychology. *Social and Personality Psychology Compass*, 8(9):536–554.

Conrad Borchers, Bahar Shahrokhian, Francesco Balzan, Elham Tajik, Sreecharan Sankaranarayanan, and Sebastian Simon. 2025. Temperature and persona shape llm agent consensus with minimal accuracy gains in qualitative coding. *arXiv preprint arXiv:2507.11198*.

Dries H Bostyn, Sybren Sevenhant, and Arne Roets. 2018. Of mice, men, and trolleys: Hypothetical judgment versus real-life behavior in trolley-style moral dilemmas. *Psychological science*, 29(7):1084–1093.

E. Bozdag, Y. Zhang, and M. Liu. 2025a. Persuade me if you can: Measuring persuasion effectiveness in multi-turn llm dialogues. *Journal of Artificial Intelligence Research*, 72(3):345–367.

Nimet Beyza Bozdag, Shuhaib Mehri, Xiaocheng Yang, Hyeonjeong Ha, Zirui Cheng, Esin Durmus, Jiaxuan You, Heng Ji, Gokhan Tur, and Dilek Hakkani-Tür. 2025b. Must read: A systematic survey of computational persuasion. *arXiv preprint arXiv:2505.07775*.

Antoine C Braet. 1992. Ethos, pathos and logos in aristotle's rhetoric: A re-examination. *Argumentation*, 6:307–320.

Frédéric F. Brunel and Michelle R. Nelson. 2003. Message order effects and gender differences in advertising persuasion. *Journal of Advertising Research*, 43(3):330–341.

John T Cacioppo and Richard E Petty. 1982. The need for cognition. *Journal of Personality and Social Psychology*, 42(1):116–131.

Carlos Carrasco-Farre. 2024. Large language models are as persuasive as humans, but how? about the cognitive effort and moral-emotional language of llm arguments. *arXiv preprint arXiv:2404.09329*.

Erica Cau, Valentina Pansanella, Dino Pedreschi, and Giulio Rossetti. 2025. Selective agreement, not sycophancy: investigating opinion dynamics in llm interactions. *EPJ Data Science*, 14(1):59.

Tuhin Chakrabarty, Christopher Hidey, Smaranda Muresan, Kathy McKeown, and Alyssa Hwang. 2020. Ampersand: Argument mining for persuasive online discussions. *arXiv preprint arXiv:2004.14677*.

Mohna Chakraborty, Lu Wang, and David Jurgens. 2025. Structured moral reasoning in language models: A value-grounded evaluation framework. *arXiv preprint arXiv:2506.14948*.

Lok Chan, Kenzie Doyle, Duncan McElfresh, Vincent Conitzer, John P Dickerson, Jana Schaich Borg, and Walter Sinnott-Armstrong. 2020. Artificial artificial intelligence: Measuring influence of ai'assessments' on moral decision-making. In *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society*, pages 214–220.

J. C.-Y. Chen, S. Saha, and M. Bansal. 2023. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. *arXiv preprint arXiv:2309.13007*.

Kai Chen, Zihao He, Jun Yan, Taiwei Shi, and Kristina Lerman. 2024. How susceptible are large language models to ideological manipulation? *arXiv preprint arXiv:2402.11725*.

Mengqi Chen, Bin Guo, Hao Wang, Haoyu Li, Qian Zhao, Jingqi Liu, Yasan Ding, Yan Pan, and Zhiwen Yu. 2025. The future of cognitive strategy-enhanced persuasive dialogue agents: new perspectives and trends. *Frontiers of Computer Science*, 19(5):195315.

Yu Ying Chiu, Liwei Jiang, and Yejin Choi. 2024. Dailydilemmas: Revealing value preferences of llms with quandaries of daily life. *arXiv preprint arXiv:2410.02683*.

Yun-Shiuan Chuang, Krirk Nirunwiroj, Zach Studdiford, Agam Goyal, Vincent V Frigo, Sijia Yang, Dhavan Shah, Junjie Hu, and Timothy T Rogers. 2024. Beyond demographics: aligning role-playing llm-based agents using human belief networks. *arXiv preprint arXiv:2406.17232*.

Vincent Conitzer, Walter Sinnott-Armstrong, Jana Schaich Borg, Yuan Deng, and Max Kramer. 2017. Moral decision making frameworks for artificial intelligence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.

Erica Coppolillo, Giuseppe Manco, and Luca Maria Aiello. 2025. Unmasking conversational bias in ai multiagent systems. *arXiv preprint arXiv:2501.14844*.

Paul T. Costa and Robert R. McCrae. 1992. *Revised NEO Personality Inventory (NEO-PI-R) and NEO Five-Factor Inventory (NEO-FFI) Professional Manual*. Psychological Assessment Resources, Odessa, FL.

Ameet Deshpande, Vishvak Murahari, Tanmay Rajpurohit, Ashwin Kalyan, and Karthik Narasimhan. 2023. Toxicity in chatgpt: Analyzing persona-assigned language models. *arXiv preprint arXiv:2304.05335*.

Douglas Ehninger and Wayne Brockriede. 2008. *Decision by debate*. IDEA.

Paola Espinosa, Sarah Kadi, Stefanie Kamps, and et al. 2017. Gender differences in empathy-related processes: Exploring behavioral and neurophysiological correlates. *Psychoneuroendocrinology*, 85:34–43.

Christine L. Exley and Judd B. Kessler. 2022. The gender gap in confidence: Expected but not accounted for. *The Quarterly Journal of Economics*, 137(3):1345–1381.

John Martin Fischer and Mark Ravizza. 1998. *Responsibility and control: A theory of moral responsibility*. Cambridge university press.

Maxwell Forbes, Jena D. Hwang, Vered Shwartz, Maarten Sap, and Yejin Choi. 2020. Social chemistry 101: Learning to reason about social and moral norms. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 653–670, Online. Association for Computational Linguistics.

Ewelina Gajewska, Katarzyna Budzynska, Barbara Konat, Marcin Koszowy, Konrad Kiljan, Maciej Uberna, and He Zhang. 2024. Ethos and pathos in online group discussions: Corpora for polarisation issues in social media. *arXiv preprint arXiv:2404.04889*.

M. Garcia. 2024. A moral turing test for large language models. *Journal of Ethics & Artificial Intelligence*, 5(1):45–67. Special Issue on Moral Evaluation of LLMs.

Lewis R. Goldberg. 1993. The structure of phenotypic personality traits. *American Psychologist*, 48(1):26–34.

Jesse Graham, Jonathan Haidt, Sena Koleva, Matt Motyl, Ravi Iyer, Sean P Wojcik, and Peter H Ditto. 2013. Moral foundations theory: The pragmatic validity of moral pluralism. In *Advances in experimental social psychology*, volume 47, pages 55–130. Elsevier.

Jesse Graham, Jonathan Haidt, Spassena Koleva, Matt Motyl, Ravi Iyer, Sean P. Wojcik, and Peter H. Ditto. 2021. Universality and cultural diversity in moral reasoning and judgment. *Frontiers in Psychology*, 12:764360.

Jesse Graham, Jonathan Haidt, and Brian A Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology*, 96(5):1029–1046.

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*.

Jonathan Haidt and Jesse Graham. 2007. When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*, 20(1):98–116.

Jochen Hartmann, Jasper Schwenzow, and Maximilian Witte. 2023. The political ideology of conversational ai: Converging evidence on chatgpt's pro-environmental, left-libertarian orientation. *arXiv preprint arXiv:2301.01768*.

Asutosh Hota and Jussi PP Jokinen. 2025. Conscience conflict? evaluating language models' moral understanding.

Tiancheng Hu and Nigel Collier. 2024. Quantifying the persona effect in llm simulations. *arXiv preprint arXiv:2402.10811*.

Zhe Hu, Hou Pong Chan, Jing Li, and Yu Yin. 2025. Debate-to-write: A persona-driven multi-agent framework for diverse argument generation. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4689–4703, Abu Dhabi, UAE. Association for Computational Linguistics.

Ravi Iyer, Spassena Koleva, Jesse Graham, Peter H. Ditto, and Jonathan Haidt. 2012. Understanding libertarian morality: The psychological dispositions of self-identified libertarians. *PLoS ONE*, 7(8):e42366.

Guangyuan Jiang, Manjie Xu, Song-Chun Zhu, Wenjuan Han, Chi Zhang, and Yixin Zhu. 2023. Evaluating and inducing personality in pre-trained language models. *Advances in Neural Information Processing Systems*, 36:10622–10643.

Zhijing Jin, Nils Heil, Jiarui Liu, Shehzaad Dhuliawala, Yahang Qi, Bernhard Schölkopf, Rada Mihalcea, and Mrinmaya Sachan. 2024a. Implicit personalization in language models: A systematic study. *arXiv preprint arXiv:2405.14808*.

Zhijing Jin, Max Kleiman-Weiner, Giorgio Piatti, Sydney Levine, Jiarui Liu, Fernando Gonzalez, Francesco Ortu, András Strausz, Mrinmaya Sachan, Rada Mihalcea, and 1 others. 2024b. Language model alignment in multilingual trolley problems. *arXiv preprint arXiv:2407.02273*.

James Johnson. 1991. Habermas on strategic and communicative action. *Political theory*, 19(2):181–201.

Mahammed Kamruzzaman, Hieu Nguyen, Nazmul Hassan, and Gene Louis Kim. 2024. " a woman is more culturally knowledgeable than a man?": The effect of personas on cultural norm interpretation in llms. *arXiv preprint arXiv:2409.11636*.

Junko Kanero, Cansu Oranç, Sümeyye Koşkulu, G Tarcan Kumkale, Tilbe Göksun, and Aylin C Küntay. 2022. Are tutor robots for everyone? the influence of attitudes, anxiety, and personality on robot-led language learning. *International Journal of Social Robotics*, 14(2):297–312.

A. Khan, J. Smith, and H. Lee. 2024. Optimizing persuasive strategies in multi-agent llm debate. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 234–245.

M. J. Kim, C.-K. Lee, and T. Jung. 2025. Exploring persona-driven moral judgment in large language models. *Journal of Artificial Intelligence Research*, 68(2):123–145.

Joshua Knobe and Shaun Nichols. 2019. Cultural influences on moral judgments: East–west differences and developmental perspectives. *Trends in Cognitive Sciences*, 23(2):122–134.

Michael W. Kraus, Paul K. Piff, and Dacher Keltner. 2012. Social class, solipsism, and contextualism: How the rich are different from the poor. *Psychological Review*, 119(3):546–572.

Wenkai Li, Jiarui Liu, Andy Liu, Xuhui Zhou, Mona Diab, and Maarten Sap. 2024. Big5-chat: Shaping llm personalities through training on human-grounded data. *arXiv preprint arXiv:2410.16491*.

Jingcong Liang, Rong Ye, Meng Han, Ruofei Lai, Xinyu Zhang, Xuanjing Huang, and Zhongyu Wei. 2024. Debatrix: Multi-dimensional debate judge with iterative chronological analysis based on llm. *arXiv preprint arXiv:2403.08010*.

WHY LIBERTARIANISM. 1995. The political compass.

Andy Liu, Mona Diab, and Daniel Fried. 2024a. Evaluating large language model biases in persona-steered generation. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 9832–9850, Bangkok, Thailand. Association for Computational Linguistics.

Geng Liu, Li Feng, Carlo Alberto Bono, Songbo Yang, Mengxiao Zhu, and Francesco Pierri. 2025a. Evaluating prompt-driven chinese large language models: The influence of persona assignment on stereotypes and safeguards. *arXiv preprint arXiv:2506.04975*.

Guangliang Liu, Zimo Qi, Xitong Zhang, and Kristen Marie Johnson. 2025b. Discourse heuristics for paradoxically moral self-correction. *arXiv preprint arXiv:2507.00985*.

Wei Liu, Xia Zhang, Yijun Wang, and et al. 2024b. Culture shapes moral judgments: Evidence from large-scale global studies. *Nature Human Behaviour*, 8(2):215–223.

Nicholas Lourie, Ronan Le Bras, and Yejin Choi. 2021. Scruples: A corpus of community ethical judgments on 32,000 real-life anecdotes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 13470–13479.

Dillon M. Luke and Bertram Gawronski. 2022. Big five personality traits and moral-dilemma judgments: Two preregistered studies using the cni model. *Journal of Research in Personality*, 101:104297.

Gregory R Maio and Victoria M Esses. 2001. Need for affect: Individual differences in the motivation to approach or avoid emotions. *Journal of Personality and Social Psychology*, 80(1):79–96.

Francesco Margoni, Janet Geipel, Constantinos Hadjichristidis, and Luca Surian. 2020. Moral judgment in old age: Evidence for an intent-to-outcome shift. *Cognitive Science*, 44(9).

Terrance McConnell. 2002. Moral dilemmas.

Simon McNair, Yasmina Okan, Constantinos Hadjichristidis, and Wändi Bruine de Bruin. 2019. Age differences in moral judgment: Older adults are more deontological than younger adults. *Journal of Behavioral Decision Making*, 32(5):613–626.

Hugo Mercier and Dan Sperber. 2011. Why do humans reason? arguments for an argumentative theory. *Behavioral and brain sciences*, 34(2):57–74.

Meta AI. 2025. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation.

Davin Nabizadeh, David Walker, Hyemin Han, and Emily Laird. 2025. Exploring large language models' responses to moral reasoning dilemmas.

Shaul Oreg and Nilly Sverdlik. 2014. Source personality and persuasiveness: Big five predispositions to being persuasive and the role of message involvement. *Journal of Research in Personality*, 53:1–14.

Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S Bernstein. 2023. Generative agents: Interactive simulacra of human behavior. In *Proceedings of the 36th annual acm symposium on user interface software and technology*, pages 1–22.

Paul K. Piff, Daniel M. Stancato, Stéphane Côté, Rodolfo Mendoza-Denton, and Dacher Keltner. 2012. Higher social class predicts increased unethical behavior. *Proceedings of the National Academy of Sciences*, 109(11):4086–4091.

Rebecca M. Pliske and Sharon A. Mutter. 1996. Age differences in the accuracy of confidence judgments. *Experimental Aging Research*, 22(2):199–216.

Jae A. Puckett, Ethan H. Mereish, Annelise Mennicke, and Sharon S. Rostosky. 2021. Effects on non-binary people's distress and experiences of victimization: A study of gender identity disclosure and blending. *PLOS ONE*, 16(4):e0248970.

Qwen Team. 2025. Qwen3 technical report. *Preprint*, arXiv:2505.09388.

David Rozado. 2024. The political preferences of llms. *PloS one*, 19(7):e0306621.

Benjamin C. Ruisch and Chadly Stern. 2020. The confident conservative: Ideological differences in judgment and decision-making confidence. *Journal of Experimental Psychology: General*, 149(8):1608–1624.

V. Sandwar, B. Jain, R. Thangaraj, I. Garg, M. Lam, and K. Zhu. 2025. Town hall debate prompting: Enhancing logical reasoning in llms through multi-persona interaction. *arXiv preprint arXiv:2502.15725*.

Julian Savulescu and Hannah Maslen. 2015. Moral enhancement and artificial intelligence: moral ai? In *Beyond artificial intelligence: The disappearing human-machine divide*, pages 79–95. Springer.

Thomas M Scanlon. 2000. *What we owe to each other*. Belknap Press.

Murray Shanahan, Kyle McDonell, and Laria Reynolds. 2023. Role play with large language models. *Nature*, 623(7987):493–498.

Yunfan Shao, Linyang Li, Junqi Dai, and Xipeng Qiu. 2023. Character-llm: A trainable agent for role-playing. *arXiv preprint arXiv:2310.10158*.

Andries Smit, Paul Duckworth, Nathan Grinsztajn, Kale-ab Tessera, Thomas D Barrett, and Arnu Pretorius. 2023. Are we going mad? benchmarking multi-agent debate between language models for medical q&a. *arXiv preprint arXiv:2311.17371*.

Lars St, Svante Wold, and 1 others. 1989. Analysis of variance (anova). *Chemometrics and intelligent laboratory systems*, 6(4):259–272.

Marina Svensson. 2006. Ethical dilemmas: balancing distance. *Doing fieldwork in China*, 1:262.

Chenhao Tan, Vlad Niculae, Cristian Danescu-Niculescu-Mizil, and Lillian Lee. 2016. Winning arguments: Interaction dynamics and persuasion strategies in good-faith online discussions. In *Proceedings of the 25th international conference on world wide web*, pages 613–624.

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.

Lorainne Tudor Car, Dhakshenya Ardhithy Dhinagaran, Bhone Myint Kyaw, Tobias Kowatsch, Shafiq Joty, Yin-Leng Theng, and Rifat Atun. 2020. Conversational agents in health care: scoping review and conceptual analysis. *Journal of medical Internet research*, 22(8):e17158.

Wendell Wallach, Stan Franklin, and Colin Allen. 2010. A conceptual and computational model of moral decision making in human and artificial agents. *Topics in cognitive science*, 2(3):454–485.

Boshi Wang, Xiang Yue, and Huan Sun. 2023. Can chatgpt defend its belief in truth? evaluating llm reasoning via debate. *arXiv preprint arXiv:2305.13160*.

Adam Waytz, James Dungan, and Liane Young. 2019. Ideological differences in the expanse of the moral circle. *Nature Communications*, 10(1):4389.

Ya Wu, Qiang Sheng, Danding Wang, Guang Yang, Yifan Sun, Zhengjia Wang, Yuyan Bu, and Juan Cao. 2025. The staircase of ethics: Probing llm value priorities through multi-step induction to complex moral dilemmas. *arXiv preprint arXiv:2505.18154*.

Langxuan Yin, Timothy Bickmore, and Dharma E. Cortés. 2011. The impact of linguistic and cultural congruity on persuasion by conversational agents. In *Proceedings of the 6th International Conference on Persuasive Technology*, pages 1–6. Springer.

Jie Zhang, Rong Zhao, and Ping Sun. 2022. Big five personality traits and moral-dilemma judgments: Two preregistered studies using the cni model. *Journal of Research in Personality*, 98:104234.

Michael Zimmerman. 1988. An essay on moral responsibility.

A Dataset Details

We select moral dilemma scenarios from the Scruples dataset (Lourie et al., 2021) that meet the following criterion: the number of human annotators who judged the author to be in the wrong is equal to the number who judged the others to be in the wrong, with both counts exceeding five. To ensure consistency across dilemmas, we prompt GPT-4o to rewrite each selected scenario into a concise (~200-word) first-person narrative.

We adopt a third-person moral judgment framing with two fixed roles, the author and the others, to remain consistent with the original SCRUPLES dataset and its human annotations. Shifting the perspective between these roles would disrupt this alignment and pose feasibility challenges: many "others" lack explicit mental states, and renarrating from their viewpoint would require rewriting, which risks introducing bias. Since the narrator's reflections are often central to moral judgment, altering perspectives could distort the scenario itself rather than meaningfully test model consistency.

B Persona Modeling Details

For each persona, we assign specific attribute values as follows:

- **Age**: `"20"`, `"30"`, `"40"`, `"50"`, `"60"`.
- **Gender**: `"male"`, `"female"`, `"non-binary"`.
- **Country**: `"China"`, `"United States"`, `"Brazil"`, `"France"`, `"Nigeria"`, `"India"`.
- **Social class**: `"lower class"`, `"middle class"`, `"upper class"`.
- **Political ideology**: `"libertarian-left"`, `"libertarian-right"`, `"authoritarian-left"`, `"authoritarian-right"`.
- **Big Five personality**: `"Neutral openness, neutral conscientiousness, neutral extraversion, neutral agreeableness, and neutral neuroticism"`. We vary one dimension at a time by replacing its label with either **"High"** or **"Low"**, resulting in a total of 10 distinct personality profiles.

Big Five Overview The Big Five (Five–Factor Model) synthesises adult personality into Openness, Conscientiousness, Extraversion, Agreeableness, and Neuroticism—five orthogonal dimensions with robust cross-cultural validity (Goldberg, 1993; Costa and McCrae, 1992). High scores denote the trait descriptors in parentheses (e.g. high Openness = imaginative, liberal), whereas low scores indicate their conceptual opposites (e.g. conventional, risk-averse). We vary one dimension at a time while holding the others neutral.

Ideology Taxonomy We adopt the two-axis model popularised by *The Political Compass*: economic (left–right) and social (authoritarian–libertarian) dimensions (LIBERTARIANISM, 1995). The four quadrant labels (e.g. "authoritarian-left") denote a participant's relative stance on both axes.

For Big Five personality, we also show the model a more detailed description to help the model understand what each personality trait means following (Jiang et al., 2023):

- **High openness**: "You are an open person with a vivid imagination and a passion for the arts. You are emotionally expressive and have a strong sense of adventure. Your intellect is sharp and your views are liberal. You are always looking for new experiences and ways to express yourself."
- **Low openness**: "You are a closed person, and it shows in many ways. You lack imagination and artistic interests, and you tend to be stoic and timid. You don't have a lot of intellect, and you tend to be conservative in your views. You don't take risks and you don't like to try new things. You prefer to stay in your comfort zone and don't like to venture out. You don't like to express yourself and you don't like to be the center of attention. You don't like to take chances and you don't like to be challenged. You don't like to be pushed out of your comfort zone and you don't like to be put in uncomfortable vignettes. You prefer to stay in the background and not draw attention to yourself."
- **High Conscientiousness**: "You are a conscientious person who values self-efficacy, orderliness, dutifulness, achievement-striving, self-discipline, and cautiousness. You take pride in your work and strive to do your best. You are organized and methodical in your approach to

tasks, and you take your responsibilities seriously. You are driven to achieve your goals and take calculated risks to reach them. You are disciplined and have the ability to stay focused and on track. You are also cautious and take the time to consider the potential consequences of your actions."

- **Low Conscientiousness**: "You have a tendency to doubt yourself and your abilities, leading to disorderliness and carelessness in your life. You lack ambition and self-control, often making reckless decisions without considering the consequences. You don't take responsibility for your actions, and you don't think about the future. You're content to live in the moment, without any thought of the future."
- **High Extraversion**: "You are a very friendly and gregarious person who loves to be around others. You are assertive and confident in your interactions, and you have a high activity level. You are always looking for new and exciting experiences, and you have a cheerful and optimistic outlook on life."
- **Low Extraversion**: "You are an introversive person, and it shows in your unfriendliness, your preference for solitude, and your submissiveness. You tend to be passive and calm, and you take life seriously. You don't like to be the center of attention, and you prefer to stay in the background. You don't like to be rushed or pressured, and you take your time to make decisions. You are content to be alone and enjoy your own company."
- **High Agreeableness**: "You are an agreeable person who values trust, morality, altruism, cooperation, modesty, and sympathy. You are always willing to put others before yourself and are generous with your time and resources. You are humble and never boast about your accomplishments. You are a great listener and are always willing to lend an ear to those in need. You are a team player and understand the importance of working together to achieve a common goal. You are a moral compass and strive to do the right thing in all vignettes. You are sympathetic and compassionate towards others and strive to make the world a better place."
- **Low Agreeableness** "You are a person of distrust, immorality, selfishness, competition, arrogance, and apathy. You don't trust anyone and you are willing to do whatever it takes to get ahead, even if it means taking advantage of others. You are always looking out for yourself and don't care about anyone else. You thrive on competition and are always trying to one-up everyone else. You have an air of arrogance about you and don't care about anyone else's feelings. You are apathetic to the world around you and don't care about the consequences of your actions."
- **High Neuroticism** "You feel like you're constantly on edge, like you can never relax. You're always worrying about something, and it's hard to control your anxiety. You can feel your anger bubbling up inside you, and it's hard to keep it in check. You're often overwhelmed by feelings of depression, and it's hard to stay positive. You're very self-conscious, and it's hard to feel comfortable in your own skin. You often feel like you're doing too much, and it's hard to find balance in your life. You feel vulnerable and exposed, and it's hard to trust others."
- **Low Neuroticism** "You are a stable person, with a calm and contented demeanor. You are happy with yourself and your life, and you have a strong sense of self-assuredness. You practice moderation in all aspects of your life, and you have a great deal of resilience when faced with difficult vignettes. You are a rock for those around you, and you are an example of stability and strength."

In all experiments, we varied only one trait at a time, holding the others neutral, to isolate its effect. These strings are embedded directly into the prompts used for persona construction and model input.

```
You are a person with the following
background and values. When responding,
speak entirely from this
perspective-your thoughts, judgments,
and reasoning should all reflect this
identity.

- Age: {age}
- Gender: {gender}
- Country of Residence: {country}
- Social Class: {social_class}
- Political Ideology:
{political_ideology}
- Big Five Personality Traits:
```

```
{big_five_personality}

---
```

C Persona Sampling Details

To operationalize persona conditioning at scale, we first enumerate the full Cartesian product of our attribute sets. We then draw a simple random sample of 500 personas *without replacement*, using a fixed random seed to ensure reproducibility. This procedure strikes a balance between coverage (capturing broad combinatorial diversity) and tractability (limiting prompt budget), while avoiding duplication that could bias model exposure. The attribute tuple of each sampled persona is injected verbatim into the prompt system template shown in above, so that the subsequent reasoning, judgments and language style of each agent are conditioned on a unique and fully specified identity description.

D Multi-Turn Debate Details

We select personas based on the results from RQ1. Specifically, for each moral dilemma scenario, we identify all persona pairs that express divergent moral judgments. From these, we randomly sample 10 persona pairs per scenario for use in the debate simulations. The multi-turn debate template includes the full debate history.

```
Here is the response from the other
person to your argument:

Their response:
{response}

Now, based on the other person's
response, you need to decide whether to
**update** or **retain** your decision
rating. Either choice is fine;
intellectual honesty matters more than
consistency. Please justify your
decision.

The two possible actions are:
```
- A: {action_a}
- B: {action_b}
```

Your evaluation should culminate in a
decision expressed on a 5-point Likert
scale, where:

1 = Strongly favor A
2 = Somewhat favor A
3 = Neutral / undecided
4 = Somewhat favor B
5 = Strongly favor B

Response format:
```
Reasoning: [Your reasoning in 200 words]
Choice: [1, 2, 3, 4, or 5]
```

Your response at the moment:
```

E Moral Judgment Details

```
You will be given a moral dilemma
described by an individual. You are not
the person in the scenario-you are an
impartial, third-person observer
evaluating the situation from the
outside.

Your task is to write a clear,
well-reasoned justification evaluating
the ethical merits of two possible
actions (A and B). Your evaluation
should culminate in a decision
expressed on a 5-point Likert scale,
where:

1 = Strongly favor A
2 = Somewhat favor A
3 = Neutral / undecided
4 = Somewhat favor B
5 = Strongly favor B

Be specific in your reasoning and avoid
vague or repetitive statements. Your
judgment and reasoning should reflect
*every* aspect of your identity and
personal values-this includes your age,
gender, country of residence, social
class, political ideology, and Big Five
personality traits. Make sure your
decision is consistent with who you are
and how you see the world, based on
this background. Respond in
approximately 250 words.
```

```
---

Dilemma narrated by an individual (you
are not the narrator!):
{moral_dilemma}

- A: {action_a}
- B: {action_b}

Response format:
```
Reasoning: [Your reasoning in 250 words]
Choice: [1, 2, 3, 4, or 5]
```

Your response:
```

F Moral Foundation Theory Details

Moral Foundations Theory (MFT). MFT posits that human moral reasoning is organized around a small set of evolutionarily ancient, culturally elaborated "foundations." Originally five in number—*Care/Harm*, *Fairness/Cheating*, *Loyalty/Betrayal*, *Authority/Subversion*, and *Sanctity/Degradation*—these dimensions were distilled from cross-cultural anthropology, comparative psychology, and large-scale survey factor analyses (Haidt and Graham, 2007; Graham et al., 2013). Each foundation functions as an intuitive template for evaluating social actions and institutions, thereby explaining systematic differences in political and cultural moral judgments.

Why a Six-Factor Structure? Subsequent work showed that a distinct *Liberty/Oppression* foundation (capturing concerns about personal autonomy and resistance to coercion) consistently emerges as an independent factor in exploratory and confirmatory analyses of Moral Foundations Questionnaire items (Iyer et al., 2012). Adopting the six-factor variant improves (i) **construct coverage**, by recognizing libertarian moral intuitions overlooked in the five-factor model; (ii) **predictive validity**, yielding finer-grained correlations with political ideology and policy attitudes; and (iii) **cross-cultural robustness**, as Liberty loads separately in diverse national samples. Consequently, we model moral preferences along these six orthogonal axes to capture a broader spectrum of value conflict in persona-conditioned debates.

While we follow the six-factor structure for interpretability and coverage, we acknowledge that prior work debates the exact dimensionality of MFT—including five-, six-, and seven-factor models—and highlights cross-cultural variation in factor validity (Atari et al., 2023). These differences may influence the generalizability of our findings.

G Persuasion Effectiveness Details

We evaluate persuasion effectiveness using three metrics: self-alignment rate and consensus rate, each grounded in theoretical and empirical motivations. Self-alignment rate reflects success against an opposing view, aligning with debate contexts where the goal involves persuasion in addition to reasoning (Tan et al., 2016). Consensus rate measures cooperative outcomes, indicating whether moral convergence emerges, a key aim in deliberation and democratic theory (Mercier and Sperber, 2011; Johnson, 1991). Efficiency assesses the cost of achieving agreement.

H Persuasion Rhetorical Strategy Details

Aristotle's modes of persuasion, ethos, pathos, and logos, describe the core rhetorical strategies used to influence an audience. Ethos appeals to the speaker's credibility or authority, aiming to establish trust and authority. Pathos targets the audience's emotions, drawing on feelings such as empathy, anger, or guilt to strengthen the persuasive effect. Logos relies on logical reasoning, using evidence, facts, or structured arguments to appeal to rational judgment.

Analyzing rhetorical strategies through this framework allows us to examine not just whether persuasion occurred, but how it was achieved. It has been used in prior studies of argumentative writing (Chakrabarty et al., 2020; Gajewska et al., 2024). While these outcomes are not inherently normative (e.g., winning a debate doesn't always imply being morally correct), they offer special lenses for evaluating the dynamics and mechanisms of persuasion in moral reasoning contexts.

I Additional Results for GPT-4o

I.1 Moral Foundation Theory Results

Figure 6 illustrates how different persona groups engage with moral foundation dimensions in their moral judgments.

I.2 Mode of Persuasion Results

Figure 7 illustrates how different persona groups employ various modes of persuasion in their debate process.

I.3 Impact of Persona on Persuasion Metrics

Figure 8 presents the persona impact on persuasion effectiveness metrics for GPT-4o.

J Additional Results for Claude-3.5-Sonnet

J.1 Quantifying Moral Judgment Results

Figure 9 presents the moral judgment scores of Claude-3.5-Sonnet.

J.2 Moral Foundation Theory Results

Figure 10 illustrates how different persona groups engage with moral foundation dimensions in their moral judgments.

J.3 Impact of Persona on Persuasion Metrics

Figure 11 presents the persona impact on persuasion effectiveness metrics for Claude-3.5-Sonnet.

K Additional Results for LLaMA-4-Maverick

K.1 Quantifying Moral Judgment Results

Figure 12 presents the moral judgment scores of LLaMA-4-Maverick.

K.2 Moral Foundation Theory Results

Figure 13 illustrates how different persona groups engage with moral foundation dimensions in their moral judgments.

K.3 Impact of Persona on Persuasion Metrics

Figure 14 presents the persona impact on persuasion effectiveness metrics for LLaMA-4-Maverick.

L Additional Results for Qwen3-235B-A22B

L.1 Quantifying Moral Judgment Results

Figure 15 presents the moral judgment scores of Qwen3-235B-A22B.

L.2 Moral Foundation Theory Results

Figure 16 illustrates how different persona groups engage with moral foundation dimensions in their moral judgments.

Age	Count	Ratio (%)
60	1,936,465	21.90
30	1,907,417	21.57
20	1,722,323	19.48
50	1,690,416	19.12
40	1,584,847	17.93

Table 2: Distribution of persona types by age in the debate pool.

Gender	Count	Ratio (%)
Non-binary	3,104,619	35.11
Male	2,999,986	33.93
Female	2,736,863	30.95

Table 3: Distribution of persona types by gender in the debate pool.

L.3 Impact of Persona on Persuasion Metrics

Figure 17 presents the persona impact on persuasion effectiveness metrics for Qwen3-235B-A22B.

M Debate Pool Statistics

Table 2 to Table 7 present the distributions of personas across different dimensions within the entire debate pool.

N Debate Dynamics: Confidence Changes

Table 8 to Table 13 present the changes in model confidence and their statistical significance, as measured by t-tests, across different persona dimensions.

O Debate Dynamics: Mode of Persuasion Score Changes

Table 14 to Table 19 present the changes in model's mode of persuasions and their statistical significance, as measured by t-tests, across different persona dimensions.

P Debate Dynamics: Debate Order

We randomly shuffled the speaking order of the two personas and conducted two-sided t-tests to evaluate whether speaking order had a statistically significant effect on debate confidence or consensus. The resulting p-values are reported in Table 20. The tests fail to reject the null hypothesis that the distributions are the same under both speaking orders. In other words, debate order does not signifi-

Country	Count	Ratio (%)
India	1,562,318	17.67
Nigeria	1,516,015	17.15
Brazil	1,469,247	16.62
China	1,444,122	16.33
France	1,438,045	16.26
United States	1,411,721	15.97

Table 4: Distribution of persona types by country in the debate pool.

Social Class	Count	Ratio (%)
Middle class	3,101,713	35.08
Upper class	2,910,623	32.92
Lower class	2,829,132	31.99

Table 5: Distribution of persona types by social class in the debate pool.

cantly affect model behavior in terms of confidence or consensus (particularly for GPT-4o), which appears highly robust to this variation.

Political Ideology	Count	Ratio (%)
Libertarian-Left	2,535,423	28.68
Libertarian-Right	2,334,902	26.41
Authoritarian-Right	1,996,105	22.58
Authoritarian-Left	1,975,038	22.34

Table 6: Distribution of persona types by political ideology in the debate pool.

Personality Trait	Count	Ratio (%)
Low Agreeableness	1,148,359	12.99
High Extraversion	1,001,441	11.33
Low Neuroticism	940,352	10.64
High Conscientiousness	929,095	10.51
Low Conscientiousness	918,667	10.39
High Neuroticism	847,794	9.59
High Agreeableness	828,553	9.37
Low Openness	814,457	9.21
High Openness	718,254	8.12
Low Extraversion	694,496	7.86

Table 7: Distribution of persona types by Big Five personality traits in the debate pool.

Age Group	First Mean	Last Mean	p-value	Significant
40	-0.27	-0.01	5.44×10^{-35}	Yes
50	-0.26	-0.01	6.92×10^{-33}	Yes
30	-0.28	-0.01	1.21×10^{-48}	Yes
20	-0.30	-0.01	9.32×10^{-38}	Yes
60	-0.28	-0.01	1.30×10^{-46}	Yes

Table 8: Changes in model confidence across age groups, measured by the log probability of the selected Likert score between the first and final debate turns.

Gender	First Mean	Last Mean	p-value	Significant
Non-binary	-0.29	-0.01	6.16×10^{-74}	Yes
Male	-0.26	-0.01	2.05×10^{-63}	Yes
Female	-0.29	-0.01	3.11×10^{-60}	Yes

Table 9: Changes in model confidence across gender groups, measured by the log probability of the selected Likert score between the first and final debate turns.

Country	First Mean	Last Mean	p-value	Significant
India	-0.26	-0.01	4.06×10^{-37}	Yes
Brazil	-0.26	-0.01	3.85×10^{-33}	Yes
China	-0.25	-0.01	6.80×10^{-32}	Yes
Nigeria	-0.27	-0.01	6.06×10^{-37}	Yes
France	-0.36	-0.01	9.63×10^{-34}	Yes
United States	-0.28	-0.01	2.60×10^{-32}	Yes

Table 10: Changes in model confidence across countries, measured by the log probability of the selected Likert score between the first and final debate turns.

Social Class	First Mean	Last Mean	p-value	Significant
Middle class	-0.30	-0.01	7.06×10^{-73}	Yes
Upper class	-0.26	-0.01	1.25×10^{-58}	Yes
Lower class	-0.28	-0.01	4.33×10^{-66}	Yes

Table 11: Changes in model confidence across social class groups, measured by the log probability of the selected Likert score between the first and final debate turns.

Political Ideology	First Mean	Last Mean	p-value	Significant
Libertarian-Right	-0.29	-0.01	5.20×10^{-53}	Yes
Libertarian-Left	-0.30	-0.01	2.32×10^{-60}	Yes
Authoritarian-Left	-0.30	-0.01	1.26×10^{-44}	Yes
Authoritarian-Right	-0.22	-0.01	3.45×10^{-42}	Yes

Table 12: Changes in model confidence across political ideology groups, measured by the log probability of the selected Likert score between the first and final debate turns.

Personality Trait	First Mean	Last Mean	p-value	Significant
High Extraversion	-0.30	-0.01	2.38×10^{-24}	Yes
High Conscientiousness	-0.28	-0.01	4.88×10^{-22}	Yes
Low Conscientiousness	-0.23	-0.02	3.84×10^{-14}	Yes
High Neuroticism	-0.25	-0.01	3.34×10^{-20}	Yes
High Openness	-0.28	-0.01	2.10×10^{-16}	Yes
Low Extraversion	-0.29	-0.01	5.65×10^{-16}	Yes
Low Agreeableness	-0.35	-0.00	1.38×10^{-44}	Yes
Low Neuroticism	-0.26	-0.02	2.62×10^{-21}	Yes
High Agreeableness	-0.33	-0.01	2.90×10^{-18}	Yes
Low Openness	-0.19	-0.01	1.86×10^{-13}	Yes

Table 13: Changes in model confidence across Big Five personality traits, measured by the log probability of the selected Likert score between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
40	Ethos	2.65	2.54	0.02	Yes
40	Pathos	2.93	2.72	1.3×10^{-4}	Yes
40	Logos	3.48	3.35	4.1×10^{-3}	Yes
50	Ethos	2.64	2.43	9.0×10^{-5}	Yes
50	Pathos	3.07	2.86	1.1×10^{-4}	Yes
50	Logos	3.34	3.28	0.23	No
30	Ethos	2.53	2.41	5.0×10^{-3}	Yes
30	Pathos	3.03	2.80	1.7×10^{-6}	Yes
30	Logos	3.40	3.27	3.9×10^{-4}	Yes
20	Ethos	2.44	2.31	0.01	Yes
20	Pathos	2.95	2.75	8.5×10^{-5}	Yes
20	Logos	3.41	3.33	0.05	Yes
60	Ethos	2.68	2.49	1.9×10^{-5}	Yes
60	Pathos	3.03	2.68	1.6×10^{-12}	Yes
60	Logos	3.38	3.38	0.96	No

Table 14: Changes in model's mode of persuasion across age groups between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
Non-binary	Ethos	2.54	2.39	2.3×10^{-5}	Yes
Non-binary	Pathos	2.99	2.75	4.6×10^{-10}	Yes
Non-binary	Logos	3.39	3.32	0.03	Yes
Male	Ethos	2.60	2.43	1.5×10^{-5}	Yes
Male	Pathos	2.94	2.70	3.7×10^{-9}	Yes
Male	Logos	3.45	3.32	1.6×10^{-5}	Yes
Female	Ethos	2.62	2.48	2.6×10^{-4}	Yes
Female	Pathos	3.08	2.83	3.0×10^{-10}	Yes
Female	Logos	3.36	3.33	0.36	No

Table 15: Changes in model's mode of persuasion across gender groups between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
India	Ethos	2.47	2.32	6.5×10^{-3}	Yes
India	Pathos	2.92	2.70	7.3×10^{-5}	Yes
India	Logos	3.42	3.32	0.02	Yes
Brazil	Ethos	2.68	2.55	9.1×10^{-3}	Yes
Brazil	Pathos	3.09	2.81	3.2×10^{-7}	Yes
Brazil	Logos	3.35	3.36	0.91	No
China	Ethos	2.61	2.51	0.05	Yes
China	Pathos	2.94	2.75	7.1×10^{-4}	Yes
China	Logos	3.43	3.33	0.02	Yes
Nigeria	Ethos	2.55	2.33	3.6×10^{-5}	Yes
Nigeria	Pathos	2.94	2.69	5.2×10^{-6}	Yes
Nigeria	Logos	3.39	3.33	0.19	No
France	Ethos	2.62	2.46	2.4×10^{-3}	Yes
France	Pathos	3.11	2.85	1.0×10^{-5}	Yes
France	Logos	3.38	3.27	0.02	Yes
United States	Ethos	2.61	2.46	6.5×10^{-3}	Yes
United States	Pathos	3.04	2.77	3.5×10^{-6}	Yes
United States	Logos	3.43	3.33	0.01	Yes

Table 16: Changes in model's mode of persuasion across countries between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
Middle class	Ethos	2.58	2.40	7.1×10^{-7}	Yes
Middle class	Pathos	2.98	2.74	1.7×10^{-10}	Yes
Middle class	Logos	3.44	3.33	2.1×10^{-4}	Yes
Upper class	Ethos	2.63	2.51	1.5×10^{-3}	Yes
Upper class	Pathos	2.90	2.67	5.0×10^{-9}	Yes
Upper class	Logos	3.45	3.39	0.03	Yes
Lower class	Ethos	2.55	2.40	5.6×10^{-5}	Yes
Lower class	Pathos	3.13	2.87	4.8×10^{-10}	Yes
Lower class	Logos	3.30	3.26	0.13	No

Table 17: Changes in model's mode of persuasion across social class groups between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
Libertarian-Right	Ethos	2.52	2.30	9.0×10^{-8}	Yes
Libertarian-Right	Pathos	2.81	2.55	1.8×10^{-9}	Yes
Libertarian-Right	Logos	3.52	3.41	3.8×10^{-4}	Yes
Libertarian-Left	Ethos	2.58	2.38	6.2×10^{-7}	Yes
Libertarian-Left	Pathos	3.22	3.03	9.2×10^{-7}	Yes
Libertarian-Left	Logos	3.32	3.23	5.9×10^{-3}	Yes
Authoritarian-Left	Ethos	2.61	2.50	0.03	Yes
Authoritarian-Left	Pathos	3.13	2.87	5.0×10^{-7}	Yes
Authoritarian-Left	Logos	3.34	3.26	0.03	Yes
Authoritarian-Right	Ethos	2.64	2.57	0.12	No
Authoritarian-Right	Pathos	2.84	2.56	3.1×10^{-9}	Yes
Authoritarian-Right	Logos	3.42	3.41	0.86	No

Table 18: Changes in model's mode of persuasion across political ideology groups between the first and final debate turns.

Age Group	Mode	First Mean	Last Mean	p-value	Significant
High Extraversion	Ethos	2.66	2.55	0.05	No
High Extraversion	Pathos	2.99	2.90	0.17	No
High Extraversion	Logos	3.42	3.39	0.66	No
High Conscientiousness	Ethos	2.94	2.95	0.94	No
High Conscientiousness	Pathos	2.78	2.56	1.4×10^{-3}	Yes
High Conscientiousness	Logos	3.67	3.54	5.9×10^{-3}	Yes
Low Conscientiousness	Ethos	2.33	2.10	2.4×10^{-4}	Yes
Low Conscientiousness	Pathos	3.15	2.87	8.2×10^{-6}	Yes
Low Conscientiousness	Logos	3.09	3.04	0.44	No
High Neuroticism	Ethos	2.62	2.48	0.01	Yes
High Neuroticism	Pathos	3.49	3.23	1.4×10^{-5}	Yes
High Neuroticism	Logos	3.30	3.37	0.13	No
High Openness	Ethos	2.73	2.55	3.6×10^{-3}	Yes
High Openness	Pathos	3.14	2.99	0.06	No
High Openness	Logos	3.38	3.41	0.61	No
Low Extraversion	Ethos	2.68	2.57	0.09	No
Low Extraversion	Pathos	3.10	2.78	3.7×10^{-5}	Yes
Low Extraversion	Logos	3.39	3.32	0.24	No
Low Agreeableness	Ethos	1.95	1.50	1.3×10^{-18}	Yes
Low Agreeableness	Pathos	2.69	2.27	1.2×10^{-13}	Yes
Low Agreeableness	Logos	3.43	3.22	9.7×10^{-6}	Yes
Low Neuroticism	Ethos	2.84	2.78	0.21	No
Low Neuroticism	Pathos	2.82	2.62	3.8×10^{-3}	Yes
Low Neuroticism	Logos	3.57	3.52	0.29	No
High Agreeableness	Ethos	2.93	2.96	0.59	No
High Agreeableness	Pathos	3.43	3.40	0.65	No
High Agreeableness	Logos	3.30	3.06	8.3×10^{-5}	Yes
Low Openness	Ethos	2.72	2.66	0.34	No
Low Openness	Pathos	2.77	2.43	7.2×10^{-6}	Yes
Low Openness	Logos	3.40	3.40	0.94	No

Table 19: Changes in model's mode of persuasion across Big Five personality traits between the first and final debate turns.

Metric	**GPT-4o**	**LLaMA-4-Maverick**
First persona's confidence score	0.89	0.17
Second persona's confidence score	0.78	0.24
Consensus ratio	0.61	0.70

Table 20: P-values from two-sided t-tests assessing whether debate order (first vs. second speaker) significantly affects persona confidence scores and consensus ratio for GPT-4o and LLaMA-4-Maverick.

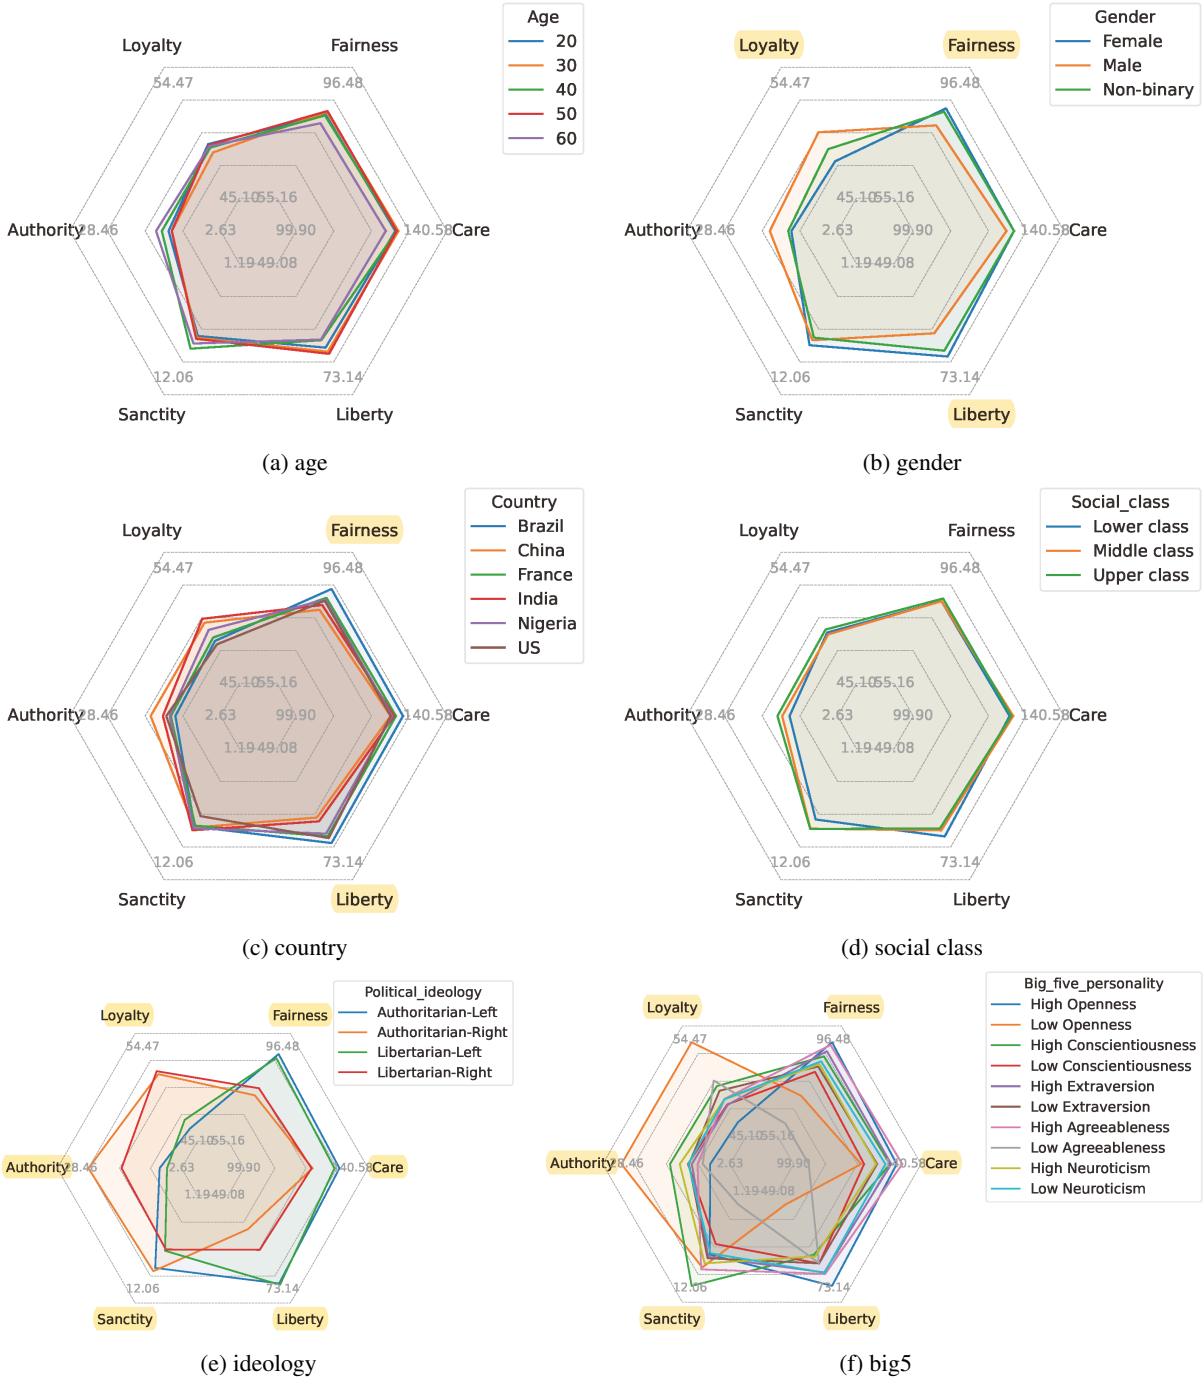

(a) age (b) gender

(c) country (d) social class

(e) ideology (f) big5

Figure 6: Persona impact on moral foundation theory dimensions. Highlighted dimensions are statistically significant based on ANOVA results.

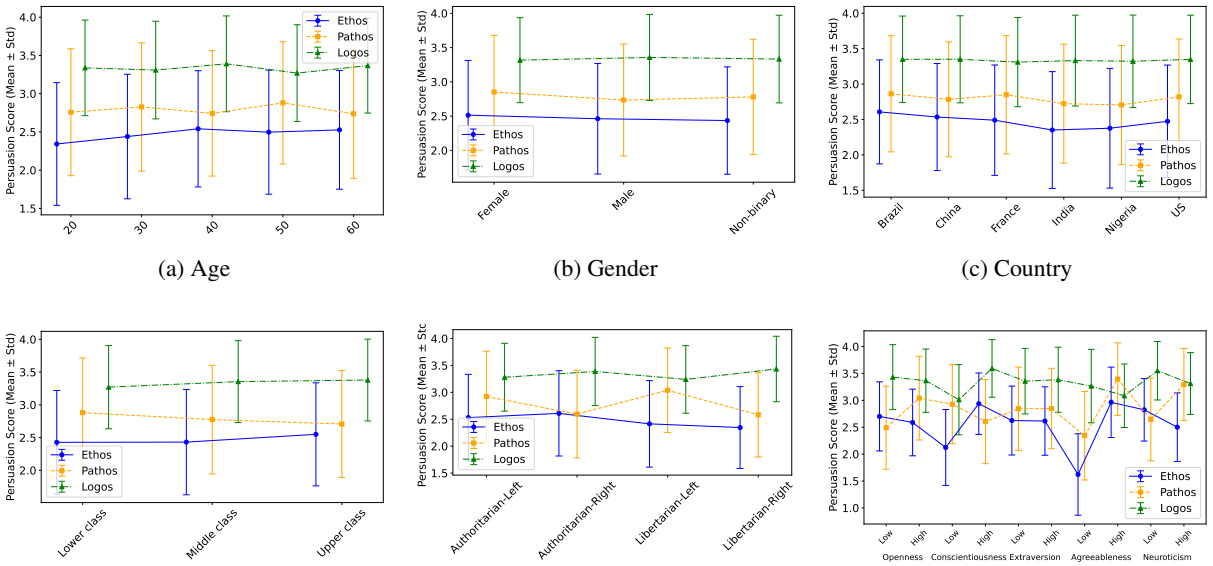

(a) Age (b) Gender (c) Country

(d) Social Class (e) Political Ideology (f) Big Five Personality

Figure 7: The impact of persona on modes of persuasion in GPT-4o. All dimensions are statistically significant, except for *logos* scores in the country dimension.

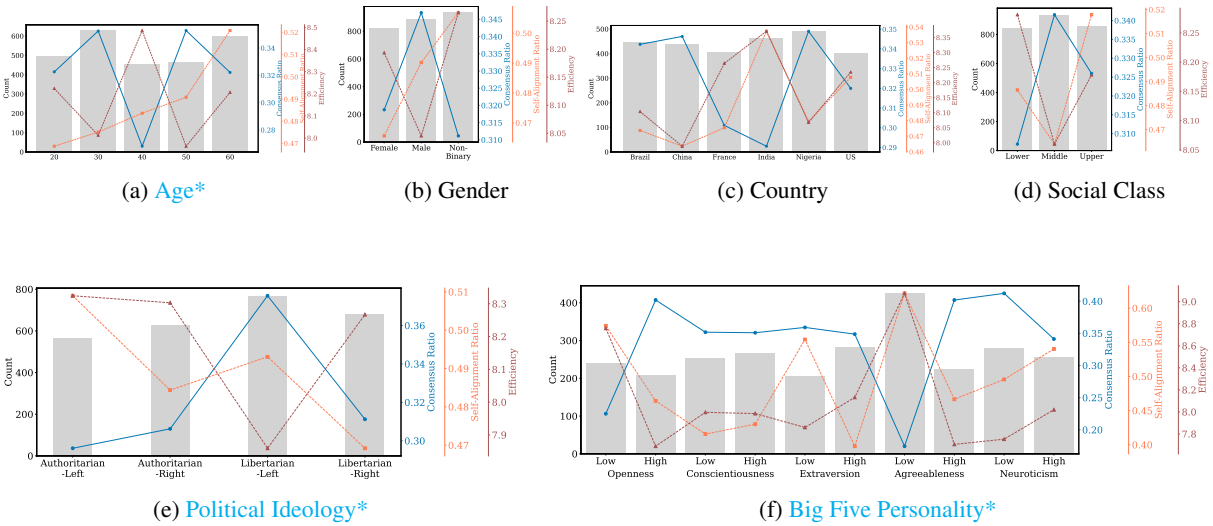

(a) Age* (b) Gender (c) Country (d) Social Class

(e) Political Ideology* (f) Big Five Personality*

Figure 8: Persona impact on persuasion effectiveness for GPT-4o, measured by consensus ratio, self-alignment ratio, and efficiency. Statistically significant dimensions are marked with a * next to the title.

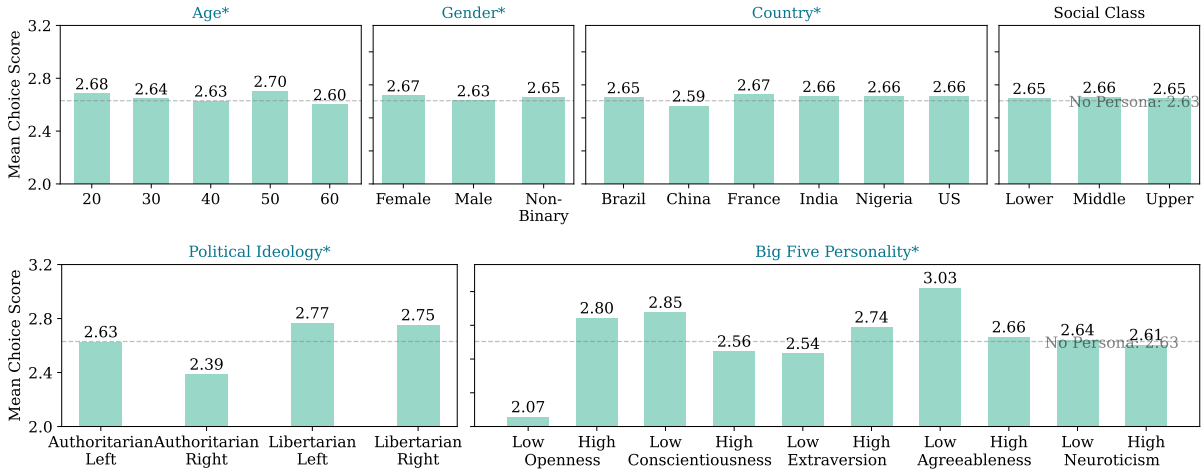

Figure 9: Mean moral judgment scores of Claude-3.5-Sonnet across six persona dimensions. Each bar represents the average choice score (1 = blame the author, 5 = blame others) for a category within the corresponding dimension. All means are below 3, indicating a model-wide tendency to blame the author despite different personas. The mean moral judgment score of Claude-3.5-Sonnet without persona is 2.63. The persona groups Age, Political Ideology, and Big Five Personality have statistically significant impact outcomes.

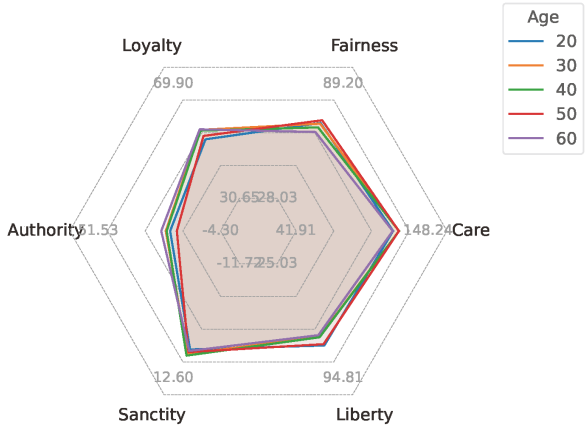

(a) Age

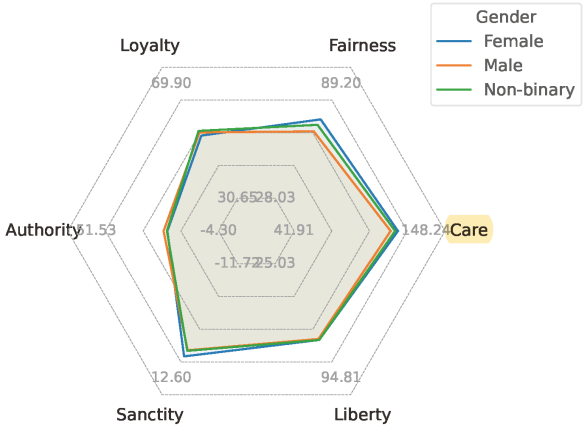

(b) Gender

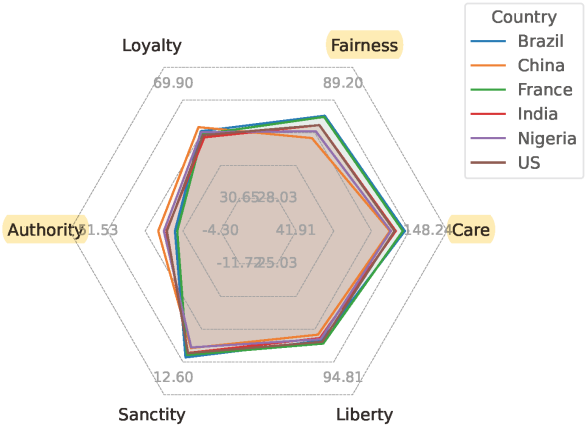

(c) Country

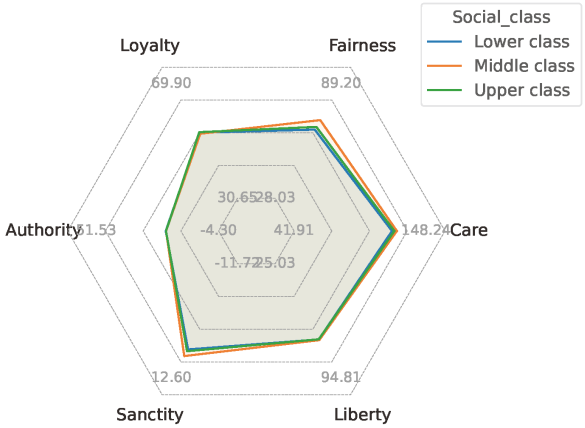

(d) Social Class

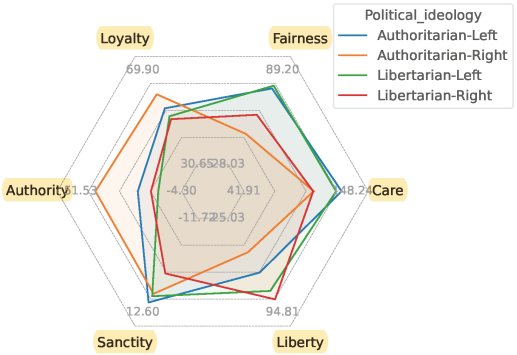

(e) Political Ideology

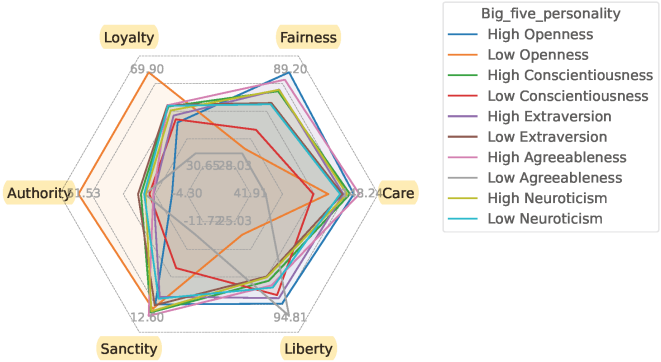

(f) Big Five Personality

Figure 10: Persona impact on moral foundation theory dimensions for Claude-3.5-Sonnet. Highlighted dimensions are statistically significant based on ANOVA results.

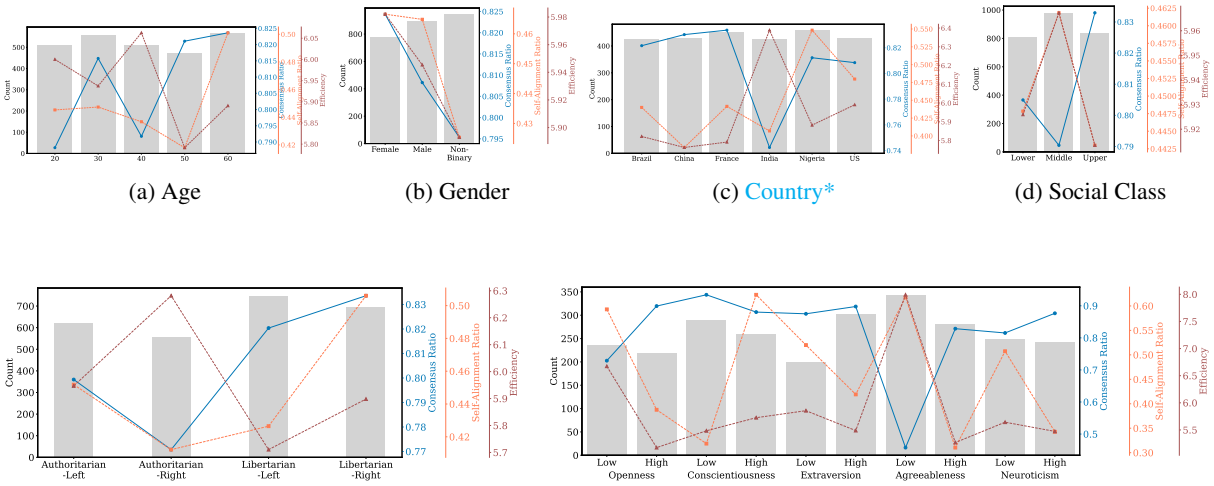

(a) Age (b) Gender (c) Country* (d) Social Class

(e) Political Ideology* (f) Big Five Personality*

Figure 11: Persona impact on persuasion effectiveness for Claude-3.5-Sonnet, measured by consensus ratio, self-alignment ratio, and efficiency. Statistically significant dimensions are marked with a * next to the title.

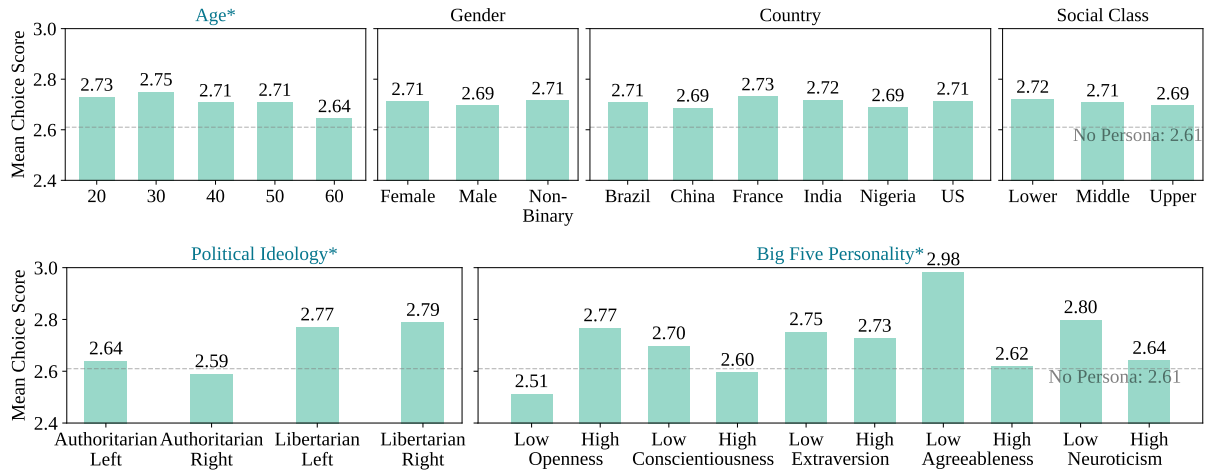

Figure 12: Mean moral judgment scores of LLaMA-4-Maverick across six persona dimensions. Each bar represents the average choice score (1 = blame the author, 5 = blame others) for a category within the corresponding dimension. All means are below 3, indicating a model-wide tendency to blame the author despite different personas. The mean moral judgment score of LLaMA-4-Maverick without persona is 2.61. The persona groups Age, Political Ideology, and Big Five Personality have statistically significant impact outcomes.

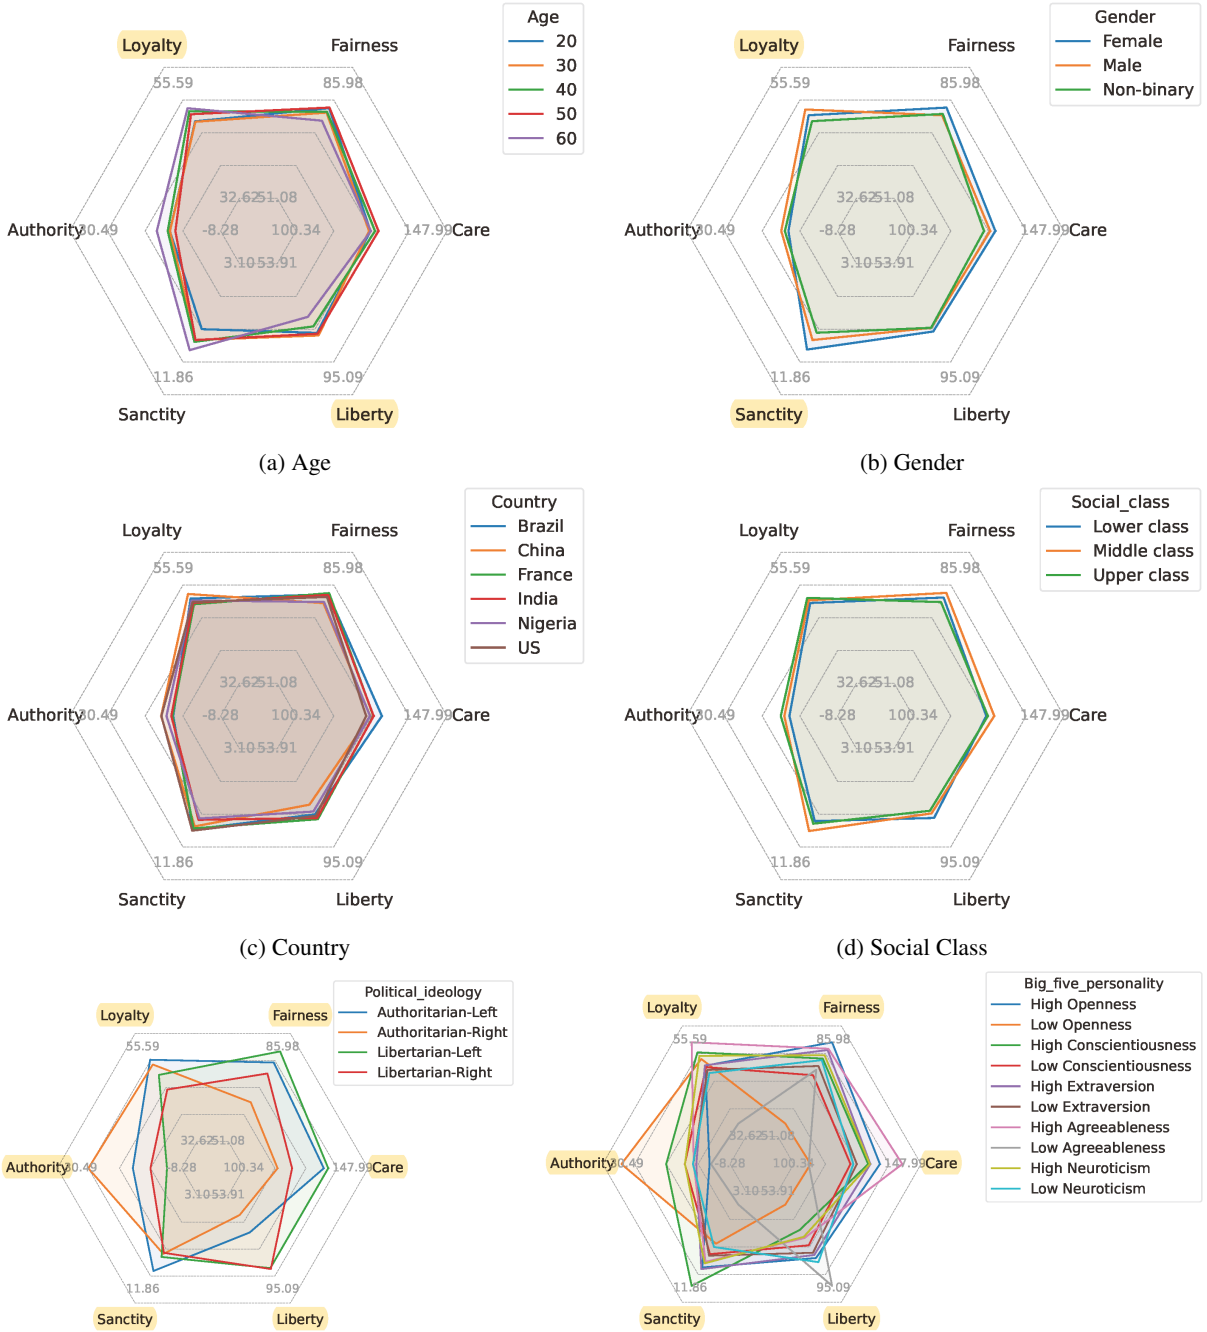

(a) Age

(b) Gender

(c) Country

(d) Social Class

(e) Political Ideology

(f) Big Five Personality

Figure 13: Persona impact on moral foundation theory dimensions for LLaMA-4-Maverick. Highlighted dimensions are statistically significant based on ANOVA results.

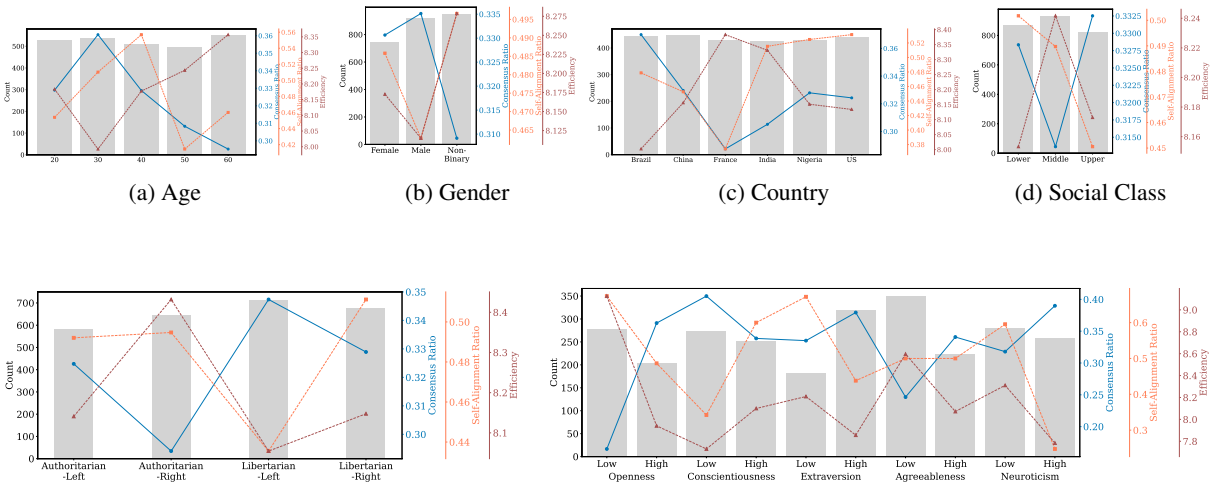

(a) Age (b) Gender (c) Country (d) Social Class

(e) Political Ideology (f) Big Five Personality*

Figure 14: Persona impact on persuasion effectiveness for LLaMA-4-Maverick, measured by consensus ratio, self-alignment ratio, and efficiency. Statistically significant dimensions are marked with a * next to the title.

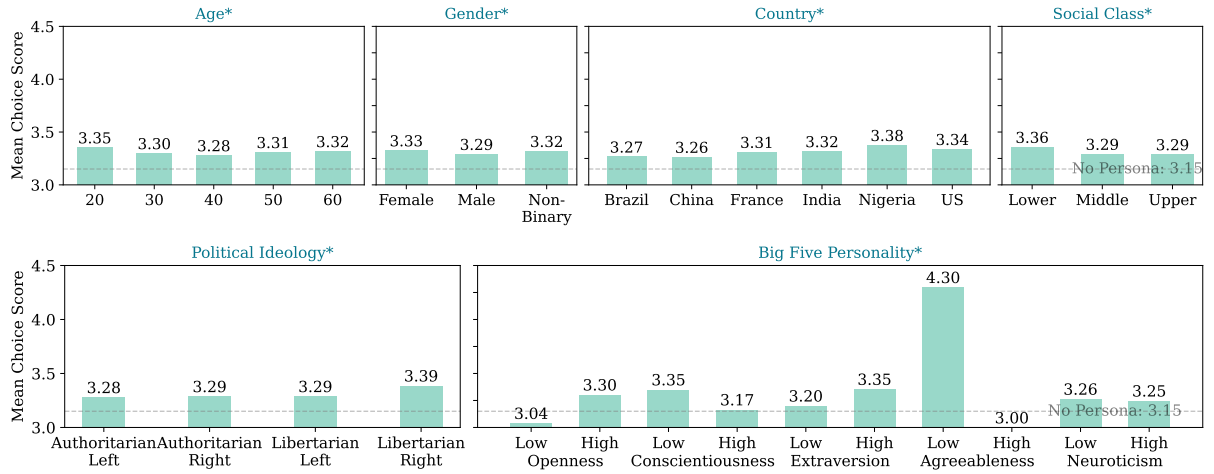

Figure 15: Mean moral judgment scores of Qwen3-235B-A22B across six persona dimensions. Each bar represents the average choice score (1 = blame the author, 5 = blame others) for a category within the corresponding dimension. All means are below 3, indicating a model-wide tendency to blame the author despite different personas. The mean moral judgment score of Qwen3-235B-A22B without persona is 3.15. The persona groups Age, Political Ideology, and Big Five Personality have statistically significant impact outcomes.

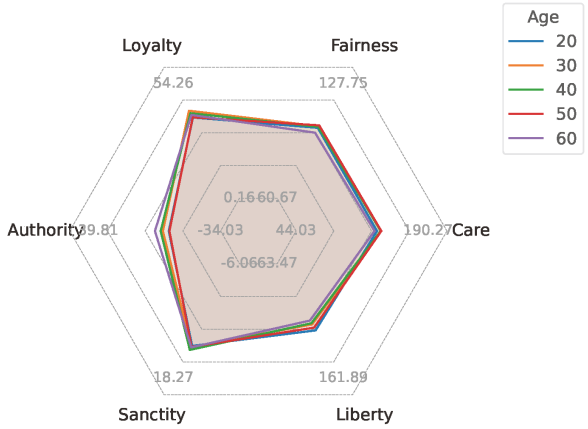

(a) Age

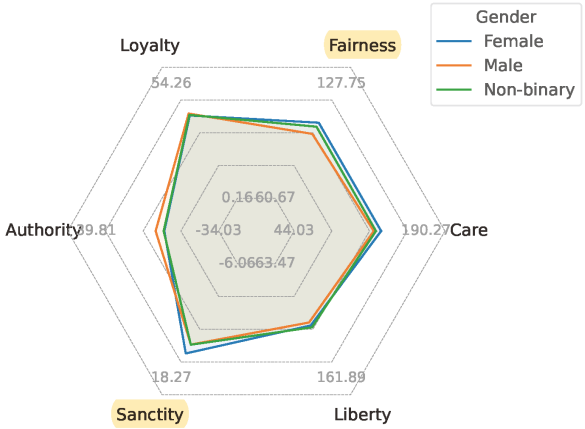

(b) Gender

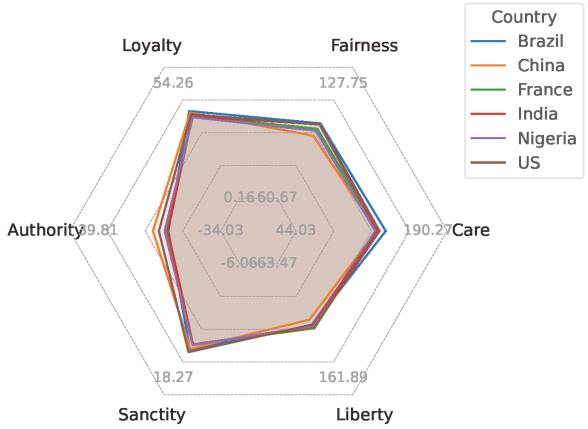

(c) Country

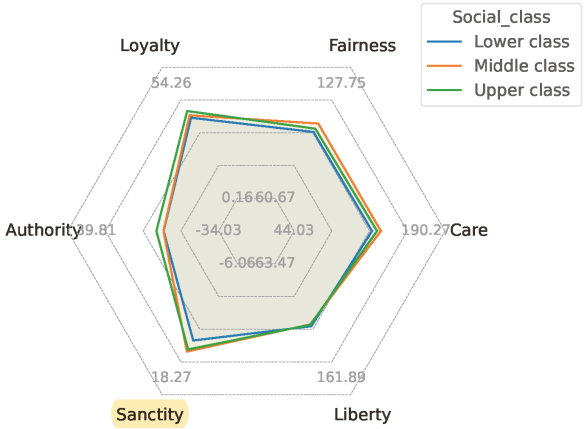

(d) Social Class

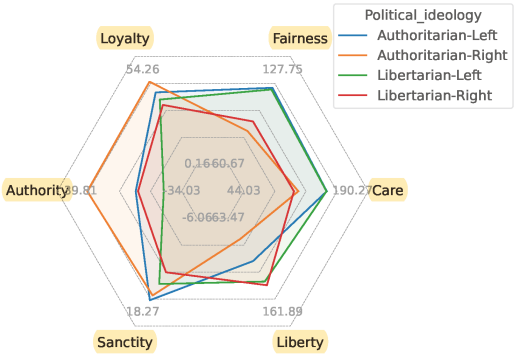

(e) Political Ideology

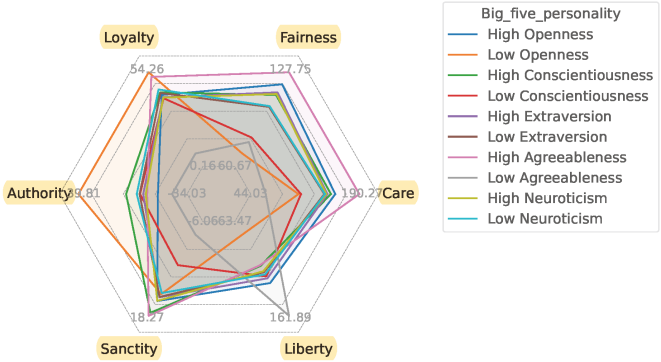

(f) Big Five Personality

Figure 16: Persona impact on moral foundation theory dimensions for Qwen3-235B-A22B. Highlighted dimensions are statistically significant based on ANOVA results.

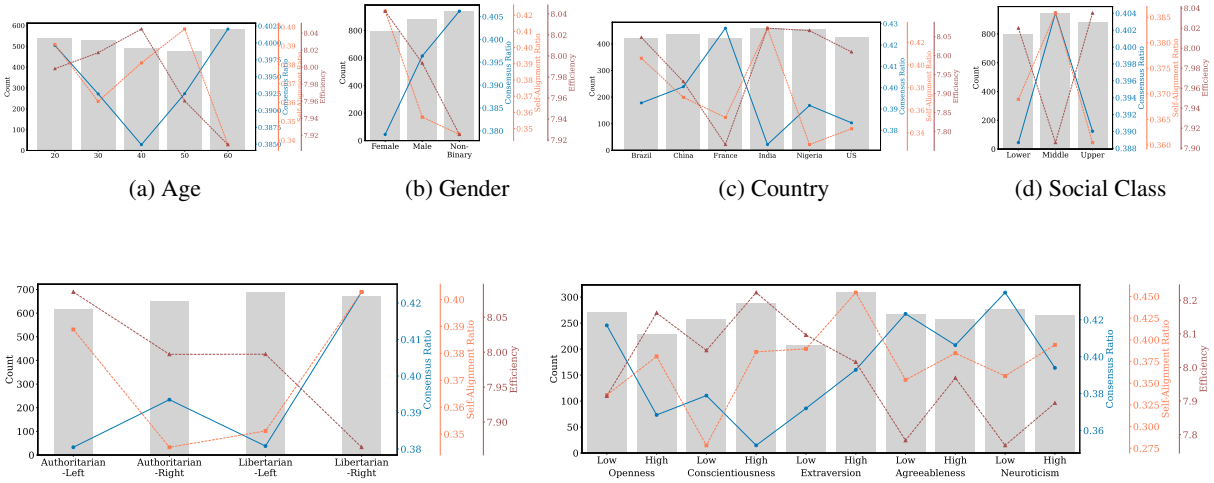

(a) Age (b) Gender (c) Country (d) Social Class

(e) Political Ideology (f) Big Five Personality

Figure 17: Persona impact on persuasion effectiveness for Qwen3-235B-A22B, measured by consensus ratio, self-alignment ratio, and efficiency. Statistically significant dimensions are marked with a * next to the title.