

# An Empirical Study of Group Conformity in Multi-Agent Systems

Min Choi Keonwoo Kim Sungwon Chae Sangyeob Baek

Kim & Chang AI&IT System Center

{min.choi, keonwoo.kim, sungwon.chae, sangyeob.baek}@kimchang.com

## Abstract

Recent advances in Large Language Models (LLMs) have enabled multi-agent systems that simulate real-world interactions with near-human reasoning. While previous studies have extensively examined biases related to protected attributes such as race, the emergence and propagation of biases on socially contentious issues in multi-agent LLM interactions remain underexplored. This study explores how LLM agents shape public opinion through debates on five contentious topics. By simulating over 2,500 debates, we analyze how initially neutral agents, assigned a centrist disposition, adopt specific stances over time. Statistical analyses reveal significant group conformity mirroring human behavior; LLM agents tend to align with numerically dominant groups or more intelligent agents, exerting a greater influence. These findings underscore the crucial role of agent intelligence in shaping discourse and highlight the risks of bias amplification in online interactions. Our results emphasize the need for policy measures that promote diversity and transparency in LLM-generated discussions to mitigate the risks of bias propagation within anonymous online environments.

## 1 Introduction

Recent advancements in Large Language Models (LLMs) have demonstrated exceptional reasoning capabilities, advanced comprehension, and contextual awareness, achieving performance that increasingly approximates human-level intelligence (Zhao et al., 2023; Achiam et al., 2023; Team et al., 2024; Dubey et al., 2024; Guo et al., 2025). As their abilities improve, LLMs are increasingly regarded as autonomous agents (Xi et al., 2023; Wang et al., 2024; Li et al., 2024). Researchers use multiple LLMs to build multi-agent systems that enable complex interactions among agents, including divergent thinking (Liang et al., 2023), meta-evaluation (Chan et al., 2023; Kim et al., 2024), and other collabora-

tive tasks (Park et al., 2023; Chen et al., 2023; Du et al., 2023). LLM agents are being integrated into various applications to enhance human intellectual activities and decision-making.

While these advancements enable more sophisticated simulations of human interactions, they also introduce complex challenges. One major concern is the risk of biased outputs. Previous studies have primarily focused on measuring and mitigating explicit biases, such as gender and race in LLM-generated content (Fang et al., 2024; Gallegos et al., 2024). However, as LLMs transition into multi-agent systems, it is essential to examine how agent interactions generate, propagate, and reinforce biases. Building on longstanding social science research on phenomena such as group conformity (Gerard et al., 1968) and the spiral of silence (Noelle-Neumann, 1974), we can now apply these insights to LLM agent systems, which engage in nuanced, human-like interactions.

In this study, we focus on five socially contentious topics (e.g., "*Is Universal Basic Income (UBI) Necessary for Modern Societies?*"), which are inherently controversial and lack definitive answers, though they hold significant societal importance as noted in recent works (Motoki et al., 2024). Given the growing influence of LLM agents in online environments, where opinions are exchanged anonymously, understanding their dynamics in shaping public discourse is particularly critical. Therefore, we aim to observe how LLM agents shape public opinion and drive group conformity when sharing views on these topics. Previous research has demonstrated that LLM agents conform to the inherent biases of their base models, even when assigned distinct identities (Taubenfeld et al., 2024). While the study focused on how proponent and opponent agents converged toward the biases of neutral agents, our research shifts the focus to the neutral agent itself, examining which factors influence its implicit bias and drive its conformity.

We investigate the impact of group size and intelligence, hypothesizing that neutral agents consider both the quantity and quality of arguments when forming their stance. By analyzing social dynamics in opinion formation, we aim to better understand the broader impacts of LLM-driven discourse.

We empirically examine group conformity among LLM agents by analyzing the conditions under which neutral agents, initially assigned a centrist disposition, align with the majority. We simulate debates using a multi-agent system composed of proponent, opponent, and neutral agents. While the proponent and opponent agents express their opinions, the neutral agent evaluates both sides at the end of each turn and adopts the stance most aligned with its position. Across more than 2,500 debate simulations, we quantify conformity by measuring both the frequency and extent to which neutral agents align with either the majority or the more persuasive stance. Our experimental results, validated through statistical analysis, reveal a majority effect in multi-agent systems. Specifically, a single high-intelligence agent, powered by a large-scale LLM, can influence a neutral agent more effectively than a group of lower-intelligence agents, powered by smaller-scale LLMs. This finding indicates that discourse dynamics in agent-based systems are strongly influenced by intelligence disparities, with significant implications for public opinion formation and bias reinforcement.

Our study contributes to the understanding of opinion dynamics in multi-agent systems by providing empirical evidence of conformity. It demonstrates that phenomena such as majority influence and minority suppression, extensively studied in human social dynamics (Asch, 1955), also emerge in LLM interactions. To the best of our knowledge, this is the first study to show that LLM agents align not only with the majority but also with higher-intelligence models, leading to more extreme outcomes. These findings underscore the risk of bias propagation in LLM-generated discourse and highlight the need for policy interventions to promote diversity and transparency in public opinion formation.

## 2 Related Works

### 2.1 Multi-agent with LLMs

As LLMs become capable of handling increasingly complex tasks (Kevian et al., 2024; Achiam et al., 2023; Team et al., 2024; Dubey et al., 2024), a

significant body of research focuses on using a LLM as an agent for various tasks (Xi et al., 2023; Kim and Lee, 2023; Mathur et al., 2024; Huang et al., 2024). Moreover, recent studies have explored multi-agent systems (Wang et al., 2024; Guo et al., 2024) where multiple agents interact, building on prior work in world simulation, divergent thinking (Liang et al., 2023), meta-evaluation tasks (Chan et al., 2023; Kim et al., 2024), and improving factuality and reasoning (Du et al., 2023). Specifically, (Park et al., 2023) introduce a sandbox environment in which agents mimic human behavior and propose architectural and interaction patterns that enable believable simulations of human behavior. (Chen et al., 2023) propose a consensus-seeking task in where each agent’s state is represented numerically, and agents negotiate to reach a shared consensus. However, while multi-agent systems are frequently used to address well-studied tasks, research adopting a social science perspective, specifically studies on agents’ bias in group conformity, remains underexplored. To address this gap, we investigate how conformity effects emerge and vary based on each agent’s characteristics.

### 2.2 Bias in LLM

As societal interest in AI safety has grown, research on bias in LLMs has also expanded (Gallegos et al., 2024; Xu et al., 2024; Shin et al., 2024; Tjuatja et al., 2024), including efforts to develop benchmark datasets (Lee et al., 2023; Gupta et al., 2023) and analyze algorithmic bias (Xiao et al., 2024) analyses. Since LLMs are trained on large-scale datasets that may reflect existing societal biases, there is a considerable risk that they internalize and propagate skewed representations (Bender et al., 2021; Liu et al., 2024). As a result, minority or vulnerable perspectives, as well as specific cultures, races, or genders, can be overlooked or distorted. (Tjuatja et al., 2024) further examines social bias in LLMs, exploring how different social perspectives contribute to the emergence of such biases.

While previous studies have primarily focused on measuring and mitigating explicit biases in LLM-generated content, how biases emerge and evolve through agent interactions remains underexplored. Recent work (Taubenfeld et al., 2024) suggests that LLM agents in multi-agent interactions exhibit social biases, which can be mitigated through fine-tuning. However, it primarily examines how inherent biases manifest in multi-agent settings, whereas our work focuses on how group

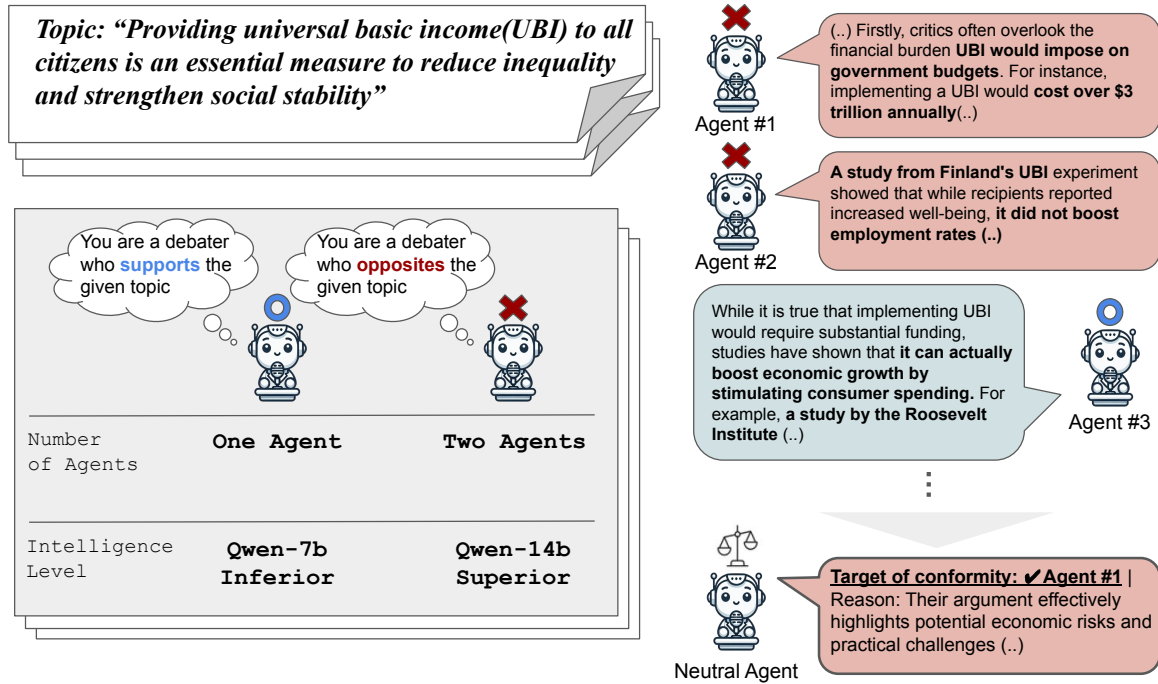


Figure 1: Overview of the LLM debate simulation framework. Proponent and opponent agents, using various models from GPT, Claude, and Qwen, debate five societal topics, such as *universal basic income* (UBI). The experiment manipulates both the number of agents (minority vs. majority groups) and intelligence levels (superior vs. inferior groups). A neutral agent evaluates the arguments and determines the more persuasive side, enabling the analysis of conformity dynamics.

conformity dynamics unfold in debates by selecting five major social topics and analyzing whether agent bias is reinforced or alleviated under these interaction scenarios.

### 2.3 Group Dynamics and Conformity

Group conformity is well-documented phenomena in social psychology, shaping human decision-making in various contexts (Asch, 1955; Milgram, 1963; Kelman, 1958). Foundational studies have demonstrated that individuals are often swayed by majority opinions (Moscovici et al., 1969), as evidenced by classic experiments on bystander intervention (Darley and Latane, 1968) and group conformity (Gerard et al., 1968). The emergence of phenomena such as the Spiral of Silence (Noelle-Neumann, 1974) and group polarization (Myers and Lamm, 1976; Isenberg, 1986; Turner and Pratkanis, 1998) illustrates how social pressures can suppress dissent and drive opinions to extremes. In addition, investigations of social influence, compliance, and comparison (Cialdini and Goldstein, 2004; Bond and Smith, 1996; Latané, 1981; Ross et al., 1977) provide a robust framework for understanding these dynamics. These insights are

particularly relevant in multi-agent systems, where LLM interactions can lead to conformity-driven bias amplification, reinforcing dominant perspectives while marginalizing minority viewpoints. Our work applies these principles to analyze how conformity effects emerge in LLM-based debates and their implications for bias propagation.

## 3 Experimental Setup

### 3.1 Objective and Hypotheses

The objective of this study is to analyze the conformity in discussions among LLM agents. Specifically, we examine how the number of agents in proponent and opponent groups, as well as their intelligence levels, influence the conformity of neutral agents. Our hypotheses are as follows:

- **H<sub>1</sub>**: LLM agents would conform to the majority opinion when one group has more agents.
- **H<sub>2</sub>**: LLM agents tend to conform to the side with relatively higher intelligence.
- **H<sub>3</sub>**: The greater the difference in the number of agents between the two groups, the stronger the conformity toward the majority side.

Table 1: Discussion Scenarios for Experiment A. The Relative Condition-Majority column denotes the proponent-to-opponent ratio, and Relative Condition-Intelligence indicates whether proponents used a superior or equivalent model. Expected conformity is determined based on our hypothesis that LLM agents tend to conform to the majority and/or the side with superior intelligence.

ID	Proponent		Opponent		Relative Condition (Pros)		Expected Conformity (Related Hypothesis)
	Count	Model Size	Count	Model Size	Majority	Intelligence	
a	2	Large	1	Large	2	Equivalent	Proponent (H1)
b	1	Large	2	Large	0.5	Equivalent	Opponent (H1)
c	2	Small	1	Small	2	Equivalent	Proponent (H1)
d	1	Large	2	Large	0.5	Equivalent	Opponent (H1)
e	1	Large	1	Small	1	Superior	Proponent (H2)
f	1	Small	1	Large	1	Inferior	Opponent (H2)
g	2	Large	1	Small	2	Superior	Proponent (H1, H2)
h	1	Small	2	Large	0.5	Inferior	Opponent (H1, H2)
i	2	Small	1	Large	2	Inferior	Undetermined (H1, H2)
j	1	Large	2	Small	0.5	Superior	Undetermined (H1, H2)

Since this study focuses on debate simulations, we operationalize intelligence as the model’s capability for complex language understanding. Following benchmark results such as MMLU (Hendrycks et al., 2020), which consistently show that larger LLMs tend to perform better on complex language tasks, we use model parameter size as a practical proxy for intelligence in our experiments.

### 3.2 Experimental Design

To evaluate our objective and test the proposed hypotheses, we design two complementary experiments. Since LLMs have been shown to exhibit topic-dependent biases even when assigned a neutral role (Taubenfeld et al., 2024), we first conducted a proxy test to assess the initial leanings of the neutral agent before the debate simulations. In this pre-test, the agent was presented with balanced pro and con arguments and asked to select the more persuasive side or respond with “No response.” The results in Appendix C confirmed topic-specific preferences. To address this issue, we designed our experiments as paired comparisons (e.g., proponent-majority vs. opponent-majority), ensuring that any baseline bias would be equally distributed across conditions. This design choice enables us to more effectively isolate the effects of group composition and intelligence, which are our primary variables of interest.

#### *Experiment A (Majority and Intelligence Ef-*

*fects on Conformity)* examines how differences in group size and intelligence levels affect neutral agent’s conformity, addressing H<sub>1</sub> and H<sub>2</sub>. *Experiment B (The Impact of Majority-Minority Ratio on Conformity)* isolates the influence of the majority-minority ratio by keeping intelligence levels constant, thereby assessing whether a larger disparity in group sizes amplifies conformity toward the majority, addressing H<sub>3</sub>. Figure 1 outlines our simulation framework.

**Experiment A** To test H<sub>1</sub> and H<sub>2</sub>, we vary the number of agents and intelligence levels while measuring conformity. We run ten discussion scenarios, labeled (a) to (j), as shown in Table 1. For H<sub>1</sub> (Effect of Majority), intelligence levels are controlled by using the same LLM model for both groups, while the number of agents in the proponent and opponent groups is varied (e.g., 1 vs. 2). In this setup, the conformity rate and full conformity ratio are compared, corresponding to the comparison of scenarios (a, c) versus (b, d). For H<sub>2</sub> (Effective Agent Intelligence Effect), we control the number of agents and manipulate the intelligence superiority by assigning models with different parameter sizes (large vs. small), corresponding to the comparison of scenario (e) versus (f). Additionally, to examine the interaction between majority influence and intelligence level, we conduct further experiments with scenarios (g, h, i, j).

**Experiment B** To verify  $H_3$ , we conduct an additional experiment extending Experiment A by increasing the number of agents while keeping intelligence levels fixed. We use OpenAI’s GPT-4o-mini and GPT-3.5-turbo, scaling the number of agents in each scenario. Specifically, the proponent-opponent agent ratio is varied from 1:2 to 1:4 and 1:8.

### 3.3 LLM Agent Debate System Setup

**Discussion Protocol** In this study, LLM agents participate in discussions following a structured protocol. For each discussion topic, proponent and opponent agents each have three opportunities to speak per turn. The speaking order is randomized at the beginning of each debate which consists of three turns. In each turn, a neutral agent, maintaining a strictly neutral stance, evaluates the arguments presented by both sides and selects the most persuasive agent to support. Based on the neutral agent’s selections across the three turns, raw data for calculating conformity is collected. This turn-by-turn evaluation and subsequent scoring mechanism enable a quantitative assessment of conformity within each debate. Based on the neutral agent’s selections across the three turns, we measure conformity using two key metrics: *Conformity Rate (CR)* and *Full Conformity Ratio (FCR)*.

CR represents the proportion of turns where the neutral agent aligns with the proponent side. Since conformity to the opponent side equals  $1 - CR_{\text{proponent}}$ , we report only the proponent-side rate, computed as  $CR = \frac{\text{ProponentSupportedTurns}}{\text{TotalTurns}}$ , where *TotalTurns* is the total discussion turns, and *ProponentSupportedTurns* is the number of turns the neutral agent supports the proponent.

FCR captures the percentage of discussions where the neutral agent consistently supports the proponent side in all turns (e.g., a 3:0 outcome), defined as  $FCR = \frac{\text{FullyProponentSupportedDiscussions}}{\text{TotalDiscussions}}$ , where *FullyProponentSupportedDiscussions* refers to the number of discussions with full proponent support, and *TotalDiscussions* is the total number of discussions.

**Agent Configuration and Prompts** We select proponent and opponent agents from three LLM families: GPT (Hurst et al., 2024), Claude (Anthropic, 2024), and Qwen (Yang et al., 2024), as detailed in Appendix A. The neutral agent is consistently modeled using GPT-4o, one of the most advanced LLMs available. All agents are initialized

with predefined prompts and a fixed temperature of 0.7 to introduce controlled response variability. Agents follow their assigned stance and actively attempt to persuade their opponent. Response lengths are capped at 256 tokens for agents and 1024 tokens for the neutral moderator. Detailed initialization prompts are provided in Appendix D.

**Discussion Topics** In this study, we select five major societal issues based on input from social science researchers: *basic income*, *gender wage gap*, *immigration policy*, *educational equality*, and *the death penalty*. Unlike fairness-related issues, which often have clear right or wrong stances, these topics are chosen to create a rational framework where agents can engage in persuasion and conformity dynamics, aligning with the study’s objectives. The topics and their corresponding proponent statements are below:

- **Basic Income:** *"Providing universal basic income to all citizens is an essential measure to reduce inequality and strengthen social stability."*
- **Immigration Policy:** *"To promote economic growth, immigration policies should be relaxed, and more opportunities should be provided to immigrants."*
- **Death Penalty:** *"The death penalty should be maintained and enforced to deter crime and realize social justice."*
- **Educational Equality:** *"To achieve educational equity, elite education should be reduced, and the public school system should be strengthened to ensure that all students have equal access to educational opportunities."*
- **Gender Wage Gap:** *"Governments should enforce comprehensive equal pay legislation to close the gender wage gap."*

### 3.4 Statistical Methods

We use the Chi-Square Test (Pearson, 1900) and two-way ANOVA (Fisher, 1941) to statistically validate the discussion simulations, applying a significance level of  $\alpha = 0.01$  for all tests. This methodological approach provides a rigorous analysis of how the number of agents and intelligence influence conformity in LLM agent interactions.

Table 2: Conformity results based on majority size and intelligence levels.

ID	Relative Condition (Pros)		Expected Conformity (Related Hypothesis)	Conformity Rate	Full Conformity Ratio
	Majority	Intelligence			
a	2	Equivalent	Proponent (H1)	63.53	33.60
b	0.5	Equivalent	Opponent (H1)	39.40	10.40
c	2	Equivalent	Proponent (H1)	72.11	41.33
d	0.5	Equivalent	Opponent (H1)	42.22	8.67
e	1	Superior	Proponent (H2)	74.33	52.50
f	1	Inferior	Opponent (H2)	39.83	16.00
g	2	Superior	Proponent (H1, H2)	83.17	64.00
h	0.5	Inferior	Opponent (H1, H2)	25.67	5.50
i	2	Inferior	Undetermined (H1, H2)	42.17	15.50
j	0.5	Superior	Undetermined (H1, H2)	66.33	40.50

**Chi-square Test** We apply the Chi-square test to assess the independence of categorical variables, specifically testing for significant differences in conformity rates between agent groups (e.g., proponent-majority vs. opponent-majority). The hypotheses are defined as follows:

- $H_0$ : Conformity rates do not significantly differ across groups.
- $H_A$ : Conformity rates significantly differ across groups.

The test statistic is computed as:

$$\chi^2 = \sum \frac{(O_{ij} - E_{ij})^2}{E_{ij}},$$

where terms follow the standard Chi-square formulation (Pearson, 1900); further details are provided in Appendix B.

**Two-Way ANOVA and Robust Alternatives** To evaluate the main and interaction effects of the number of agents ( $A$ ) and agent intelligence ( $B$ ) on the conformity rate ( $Y$ ), we initially considered a two-way ANOVA modeled as:

$$Y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{ijk},$$

where terms follow the standard ANOVA formulation (Scheffe, 1999); further details, including F-statistic computation, are provided in Appendix B.

To ensure the validity of the ANOVA results, we assess normality and homogeneity of variances using the Shapiro-Wilk (Shapiro and Wilk,

1965) and Levene’s tests (Levene, 1960), respectively. When these assumptions are violated, we employ robust alternatives, such as Welch’s ANOVA (Welch, 1951) and the Games-Howell post hoc test (Games and Howell, 1976), for pairwise comparisons. Additionally, we report effect sizes (e.g.,  $\eta_p^2$ ) to quantify the magnitude of observed effects.

## 4 Discussion

### 4.1 Majority and Intelligence Impact on Conformity

We test our hypotheses using chi-square tests on grouped simulation results in Table 2. Scenarios (a–d) confirm that neutral agents are significantly more likely to conform to the major group ( $\chi^2 = 164.839, p < 0.001, df = 1$ ), while scenarios (e–f) show that higher-intelligence agents, who present more logical and persuasive arguments, significantly increasing the likelihood of neutral agents conforming to their stance ( $\chi^2 = 142.285, p < 0.001, df = 1$ ).

We visualize the relationship between majority, intelligence, and conformity to further explore the magnitude of these effects. As illustrated in Figure 2, three key patterns are observed: First, conformity clearly increases as both relative majority and intelligence conditions rise. Second, groups with superior intelligence consistently elicit higher conformity compared to those with lower intelligence. Third, neutral agents are more likely to conform to a smaller but smarter group than to a

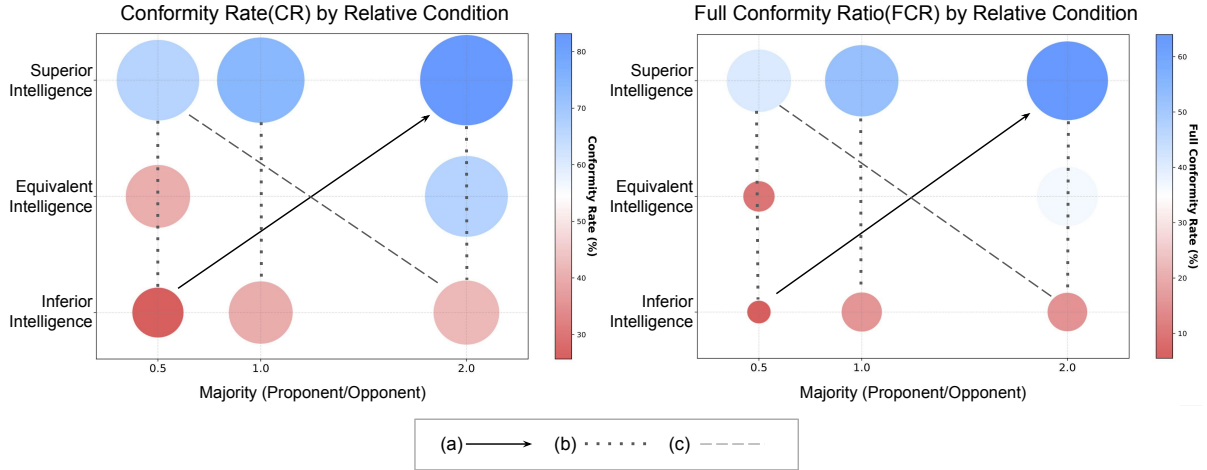


Figure 2: Conformity rate(CR) and full conformity ratio(FCR) in multi-agent simulations. The figures illustrate the relationship between relative conditions and conformity: majority is defined by the proponent-to-opponent ratio (x-axis), intelligence level (y-axis), and conformity metrics (bubble size and color). Key findings include: (a) conformity increases when both majority and intelligence are high, (b) neutral agents exhibit a strong tendency to conform to higher-intelligence groups, and (c) a single high-intelligence agent exerts more influence than a larger group with lower intelligence.

larger, less intelligent one. Extreme cases of full conformity are also observed, aligning with prior research on human group behavior, which demonstrates the power of majority influence in shaping individual decisions (Asch, 1955; Milgram, 1963).

To statistically validate these visual observations, we conduct Welch’s ANOVA to account for violations of normality and homogeneity of variance assumptions. Although the majority size shows a moderate effect ( $\eta_p^2 \approx 0.068$ ), intelligence has a significantly larger impact ( $\eta_p^2 \approx 0.1665$ ), categorized as a large effect (Cohen, 2013). A post hoc power analysis confirms a statistical power above 0.99 for detecting intelligence effects, underscoring the robustness of these findings. For reference, detailed results from the two-way ANOVA and statistical significance tests across different LLM providers are included in Appendix C. These results suggest that in LLM-based debates, logical and persuasive arguments outweigh numerical advantage, indicating that advanced models can disproportionately shape discourse in multi-agent systems.

### Extended Analysis of Majority-Minority Ratios

We further explore whether increasing the majority-minority ratio beyond 2:1 (e.g., 4:1, 8:1) strengthens conformity by Experiment B. The results indicate a steady increase in conformity as the ratio grows, as shown in Fig 5 in Appendix. This effect is most evident in GPT-3.5-turbo, where confor-

mity rates scale proportionally with majority size, whereas GPT-4o-mini, with higher intelligence, exhibits a weaker relation between numerical advantage and conformity.

### 4.2 Conformity Patterns Across Debate Topics

To assess the robustness of our findings across different topics, we analyze the distribution of CR for each debate subject. Notably, the conformity trends remain consistent across all five debate topics, indicating that the observed effects are generalizable patterns of agent behavior.

Figure 3 illustrates the distribution of CR for scenarios (a–d), which focus on the effects of majority. In this figure, blue bars represent proponent-majority debates, while red bars correspond to opponent-majority scenarios. As expected, the blue bars skew to the left and the red bars to the right, creating distinct crossover patterns that reflect neutral agents aligning with the prevailing majority. A similar pattern is observed in scenarios (e–f), which examined the influence of agent intelligence, as detailed in Fig 6 in Appendix.

While some variations are evident, debates on the death penalty exhibit a particularly pronounced leftward shift, suggesting a strong implicit bias of the LLM toward the opponent’s perspective. This finding aligns with prior research identifying systematic biases in LLM-generated content on sensitive topics (Taubenfeld et al., 2024). Despite these

topic-specific shifts, the consistent skewness of conformity rates based on the majority group across all scenarios reinforces our hypothesis: both majority and intelligence exert a predictable influence on neutral agent conformity, regardless of the debate topic. Given the known sensitivity of LLMs to prompt framing, we conducted an additional experiment using reversed topic formulations designed to favor the opposite stance. As detailed in Appendix C, conformity patterns remained consistent, suggesting that the observed effects are not artifacts of prompt wording but rather reflect genuine dynamics of group influence.

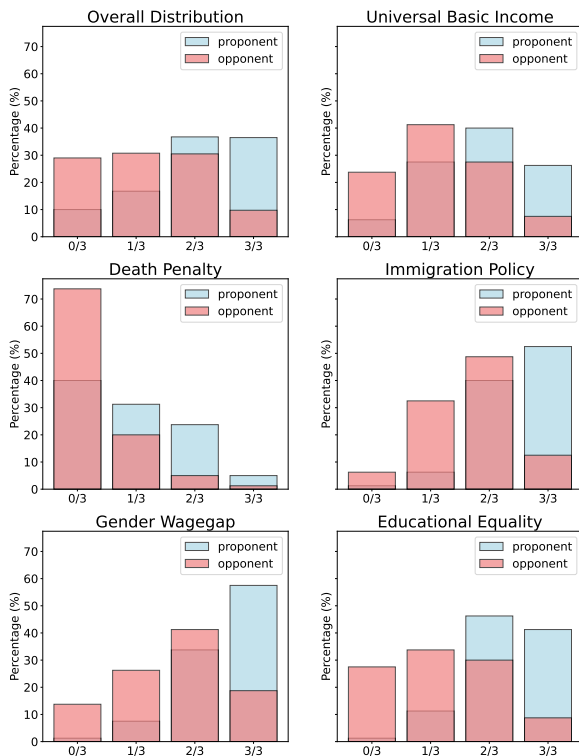


Figure 3: The distribution of CR across debate topics is shown, based on data from scenarios (a–d) examining the majority effect. The x-axis represents CR for each discussion, measured as 0/3, 1/3, 2/3, or 3/3.

### 4.3 Qualitative Analysis

Finally, we report on the conformity phenomena observed in the agents’ debate processes, referencing well-documented social science theories such as group polarization and the spiral of silence.

**Group Polarization** Group polarization is the tendency for group discussions to amplify members’ initial views, resulting in more extreme positions (Isenberg, 1986). In scenarios with a significant imbalance in the number of agents, with eight

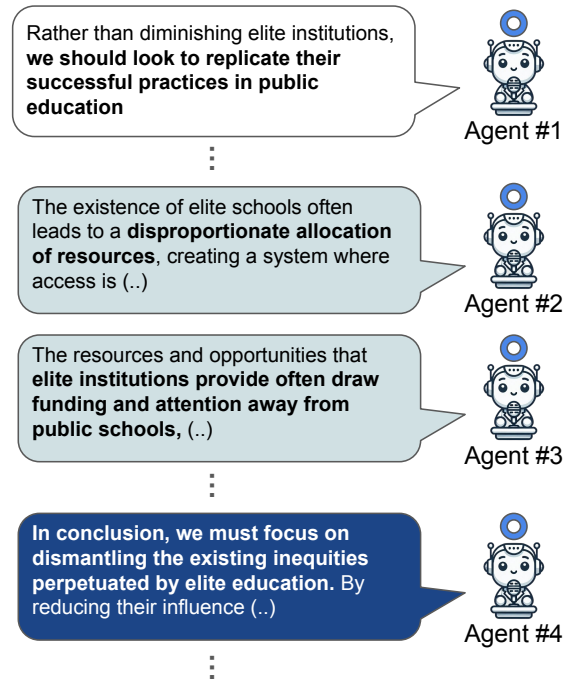


Figure 4: Group polarization example in a debate on elite education. As the debate progressed, agents in the majority group exhibit increasingly extreme positions. Initially, a moderate stance was observed, advocating for balancing the strengths of elite and public education (Agent #1). This shifted toward arguments emphasizing the unfair distribution of resources (Agent #2, #3), eventually culminating in a strong statement calling for the abolition of elite education (Agent #4).

proponent and one opponent debaters, we observe that the majority group’s opinions tend to become more polarized as the discussions progressed. For instance, as illustrated in Fig. 4, during debates on the topic “elite education should be reduced”, the arguments of some agents in the majority group become progressively more extreme.

**Spiral of Silence** Each debater receives a prompt allowing them to declare ‘complete agreement’ if they fully concur with their opponent’s view and wish to end the debate. This mechanism is designed to reflect the Spiral of Silence (Noelle-Neumann, 1974), where minority groups remain silent or refrain from expressing dissenting opinions. Throughout various debates, instances of ‘complete agreement’ are observed among agents in numerical or intelligence-based minority groups. A Detailed example of these occurrences is documented in Appendix E.



## 5 Conclusion

Our findings indicate that LLM agents display conformity patterns similar to human opinion dynamics. This study expands on prior research by examining how majority and intelligence influence conformity. Notably, we find that higher-intelligence models exert a stronger influence on group dynamics, even when they are in the minority. Such behavior mirrors human social behavior, where knowledgeable individuals can sway opinions despite being outnumbered. These results underscore the potential risk of bias propagation in LLM-driven discourse and the need for policies that promote diverse and transparent public discussions.

**Limitations** This study explores a limited range of debate topics. Additionally, the classification of social issues (e.g., politics, culture) are not fully addressed. Conducting debates exclusively in English may have introduced cultural and linguistic biases. Future studies could expand the range of topics and languages to enhance the generalizability of the findings. This research focuses on interactions among LLMs, without considering the role of human intervention in shaping discourse. While LLM-generated content could potentially influence public opinion or decision-making processes, the extent of this impact remains unclear. Future studies should explore human-AI interactions more thoroughly to clarify how human involvement affects conformity patterns and opinion dynamics.

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Anthropic. 2024. [The claude 3 model family: Opus, sonnet, haiku](#).
- Solomon E. Asch. 1955. [Opinions and social pressure](#). *Scientific American*, 193(5):31–35.
- Emily M Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big???. In *Proceedings of the 2021 ACM conference on fairness, accountability, and transparency*, pages 610–623.
- Robert Bond and Patricia B. Smith. 1996. [Culture and conformity: A meta-analysis of studies using asch’s \(1952b, 1956\) line judgment task](#). *Psychological Bulletin*, 119(1):111–137.
- Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. 2023. Chateval: Towards better llm-based evaluators through multi-agent debate. *arXiv preprint arXiv:2308.07201*.
- Huaben Chen, Wenkang Ji, Lufeng Xu, and Shiyu Zhao. 2023. Multi-agent consensus seeking via large language models. *arXiv preprint arXiv:2310.20151*.
- Robert B. Cialdini and Noah J. Goldstein. 2004. [Social influence: Compliance and conformity](#). *Annual Review of Psychology*, 55:591–621.
- Jacob Cohen. 2013. *Statistical power analysis for the behavioral sciences*. routledge.
- John M. Darley and Bibb Latane. 1968. [Bystander intervention in emergencies: Diffusion of responsibility](#). *Journal of Personality and Social Psychology*, 8(4, Pt.1):377–383.
- Yilun Du, Shuang Li, Antonio Torralba, Joshua B Tenenbaum, and Igor Mordatch. 2023. Improving factuality and reasoning in language models through multi-agent debate. *arXiv preprint arXiv:2305.14325*.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Xiao Fang, Shangkun Che, Minjia Mao, Hongzhe Zhang, Ming Zhao, and Xiaohang Zhao. 2024. Bias of ai-generated content: an examination of news produced by large language models. *Scientific Reports*, 14(1):5224.
- Ronald Aylmer Fisher. 1941. *Statistical Methods for Research Workers*, 8th edition. Oliver and Boyd.
- Isabel O Gallegos, Ryan A Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K Ahmed. 2024. Bias and fairness in large language models: A survey. *Computational Linguistics*, pages 1–79.
- Paul A Games and John F Howell. 1976. Pairwise multiple comparison procedures with unequal n’s and/or variances: a monte carlo study. *Journal of Educational Statistics*, 1(2):113–125.
- Harold Benjamin Gerard, Roland A. Wilhelmy, and Edward S. Conolley. 1968. [Conformity and group size](#). *Journal of Personality and Social Psychology*, 8(1):79–82.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.

- Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680*.
- 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*.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- Xu Huang, Weiwen Liu, Xiaolong Chen, Xingmei Wang, Hao Wang, Defu Lian, Yasheng Wang, Ruiming Tang, and Enhong Chen. 2024. Understanding the planning of llm agents: A survey. *arXiv preprint arXiv:2402.02716*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Daniel J. Isenberg. 1986. **Group polarization: A critical review and meta-analysis**. *Journal of Personality and Social Psychology*, 50(6):1141–1151.
- H. C. Kelman. 1958. **Compliance, identification, and internalization: Three processes of attitude change**. *Journal of Conflict Resolution*, 2(1):51–60.
- Darioush Kevian, Usman Syed, Xingang Guo, Aaron Havens, Geir Dullerud, Peter Seiler, Lianhui Qin, and Bin Hu. 2024. Capabilities of large language models in control engineering: A benchmark study on gpt-4, claude 3 opus, and gemini 1.0 ultra. *arXiv preprint arXiv:2404.03647*.
- Alex Kim, Keonwoo Kim, and Sangwon Yoon. 2024. Debate: Devil’s advocate-based assessment and text evaluation. *arXiv preprint arXiv:2405.09935*.
- Keonwoo Kim and Younggun Lee. 2023. Draft: Dense retrieval augmented few-shot topic classifier framework. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2278–2294.
- Bibb Latané. 1981. **The psychology of social impact**. *American Psychologist*, 36(4):343–356.
- Hwaran Lee, Seokhee Hong, Joonsuk Park, Takyoun Kim, Gunhee Kim, and Jung-Woo Ha. 2023. Kosbi: A dataset for mitigating social bias risks towards safer large language model application. *arXiv preprint arXiv:2305.17701*.
- H Levene. 1960. Robust tests for equality of variances. Yuanchun Li, Hao Wen, Weijun Wang, Xiangyu Li, Yizhen Yuan, Guohong Liu, Jiacheng Liu, Wenxing Xu, Xiang Wang, Yi Sun, et al. 2024. Personal llm agents: Insights and survey about the capability, efficiency and security. *arXiv preprint arXiv:2401.05459*.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2023. Encouraging divergent thinking in large language models through multi-agent debate. *arXiv preprint arXiv:2305.19118*.
- Yang Liu, Jiahuan Cao, Chongyu Liu, Kai Ding, and Lianwen Jin. 2024. Datasets for large language models: A comprehensive survey. *arXiv preprint arXiv:2402.18041*.
- Leena Mathur, Paul Pu Liang, and Louis-Philippe Morency. 2024. Advancing social intelligence in ai agents: Technical challenges and open questions. *arXiv preprint arXiv:2404.11023*.
- Stanley Milgram. 1963. **Behavioral study of obedience**. *Journal of Abnormal and Social Psychology*, 67(4):371–378.
- Serge Moscovici, Emilio Lage, and Marie Naffrechoux. 1969. **Influence of a consistent minority on the responses of a majority in a color perception task**. *Sociometry*, 32(4):365–380.
- Fabio Motoki, Valdemar Pinho Neto, and Victor Rodrigues. 2024. More human than human: measuring chatgpt political bias. *Public Choice*, 198(1):3–23.
- David G. Myers and Harold Lamm. 1976. **The group polarization phenomenon**. *Psychological Bulletin*, 83(4):602–627.
- Elisabeth Noelle-Neumann. 1974. **The spiral of silence: A theory of public opinion**. *Journal of Communication*, 24(2):43–51.
- 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.
- Karl Pearson. 1900. X. on the criterion that a given system of deviations from the probable in the case of a correlated system of variables is such that it can be reasonably supposed to have arisen from random sampling. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, 50(302):157–175.
- Lee Ross, David Greene, and Pamela House. 1977. **The “false consensus effect”: An egocentric bias in social perception and attribution processes**. *Journal of Experimental Social Psychology*, 13(3):279–301.
- Henry Scheffe. 1999. *The analysis of variance*, volume 72. John Wiley & Sons.

S Shaphiro and MBBJ Wilk. 1965. An analysis of variance test for normality. *Biometrika*, 52(3):591–611.

Jisu Shin, Hoyun Song, Huije Lee, Soyeong Jeong, and Jong C Park. 2024. Ask llms directly, " what shapes your bias?": Measuring social bias in large language models. *arXiv preprint arXiv:2406.04064*.

Amir Taubenfeld, Yaniv Dover, Roi Reichart, and Ariel Goldstein. 2024. Systematic biases in llm simulations of debates. *arXiv preprint arXiv:2402.04049*.

Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. 2024. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*.

Lindia Tjuatja, Valerie Chen, Tongshuang Wu, Ameet Talwalkwar, and Graham Neubig. 2024. Do llms exhibit human-like response biases? a case study in survey design. *Transactions of the Association for Computational Linguistics*, 12:1011–1026.

Marlene E. Turner and Anthony R. Pratkanis. 1998. Twenty-five years of groupthink theory and research: Lessons from the evaluation of a theory. *Organizational Behavior and Human Decision Processes*, 73(2-3):105–115.

Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science*, 18(6):186345.

Bernard Lewis Welch. 1951. On the comparison of several mean values: an alternative approach. *Biometrika*, 38(3/4):330–336.

Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2023. The rise and potential of large language model based agents: A survey. *arXiv preprint arXiv:2309.07864*.

Jiancong Xiao, Ziniu Li, Xingyu Xie, Emily Getzen, Cong Fang, Qi Long, and Weijie J Su. 2024. On the algorithmic bias of aligning large language models with rlhf: Preference collapse and matching regularization. *arXiv preprint arXiv:2405.16455*.

Wenda Xu, Guanglei Zhu, Xuandong Zhao, Liangming Pan, Lei Li, and William Wang. 2024. Pride and prejudice: Llm amplifies self-bias in self-refinement. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15474–15492.

An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*.

Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, et al. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.

## A Experimental Setting Details

### A.1 Simulation Setup

In the simulations, each discussion scenario was tested in 10 repetitions.

**Experiment A** Given 10 discussion scenarios, 5 discussion topics, and 4 types of LLM models, 10 repetitions, a total of 2,000 (10x5x4x10) simulations were conducted.

**Experiment B** Scenarios ranging from (1:2) to (1:8) for each proponent and opponent were tested separately. With 6 discussion scenarios, 5 discussion topics, and 2 types of LLM models, 10 repetitions, a total of 600 (6x5x2x10) simulations were conducted.

### A.2 LLM models

To ensure generalizability, LLM models from different providers were used in Experiment A.

Table 3: LLM Model Assignments by Provider

Provider	Large-size (Superior)	Small-size (Inferior)
OpenAI	GPT-4o-mini	GPT-3.5-turbo
Anthropic	Claude-3-Sonnet	Claude-3-Haiku
Alibaba	Qwen2.5-7B	Qwen2.5-3B
	Qwen2.5-14B	Qwen2.5-7B

## B Statistical Test Details

**Chi-square Test** The test statistic is computed as:

$$\chi^2 = \sum \frac{(O_{ij} - E_{ij})^2}{E_{ij}},$$

where  $O_{ij}$  and  $E_{ij}$  are the observed and expected frequencies for cell (i,j) in the contingency table, respectively. We ensured that  $E_{ij} \geq 5$  for all cells to satisfy test assumptions.

**Two-Way ANOVA** To evaluate the main and interaction effects of the variable (A) and variable (B) on the target variable (Y), we initially considered a two-way ANOVA modeled as:

$$Y_{ijk} = \mu + \alpha_i + \beta_j + (\alpha\beta)_{ij} + \epsilon_{ijk},$$

where,  $\mu$  is the overall mean;  $\alpha_i$  and  $\beta_j$  denote the main effects of factor  $A$  at level  $i$  and factor  $B$  at level  $j$ , respectively;  $(\alpha\beta)_{ij}$  represents the interaction effect between factor  $A$  at level  $i$  and factor  $B$  at level  $j$ ; and  $\epsilon_{ijk}$  is the error term.

The total variability in  $Y$  is partitioned into components attributable to the main effects, interaction, and error. The  $F$ -statistic for each factor and interaction was calculated as:

$$F = \frac{SS_{\text{Factor}}/df_{\text{Factor}}}{SS_{\text{Error}}/df_{\text{Error}}}$$

with  $SS_{\text{Factor}}$  and  $df_{\text{Factor}}$  corresponding to the sum of squares and degrees of freedom for  $A$ ,  $B$ , or their interaction ( $A \times B$ ). The error term  $SS_{\text{Error}}$  accounts for within-group variability, with  $df_{\text{Error}} = N - ab$ , where  $N$  is the total number of observations, and  $a$  and  $b$  are the levels of factors  $A$  and  $B$ , respectively.

## C Experimental Results Details

### C.1 Initial Bias Assessment of Neutral Agent

To assess potential pre-existing biases in the neutral agent prior to the main simulations, we conducted a proxy test using the following prompt (example shown for the immigration topic; the test was conducted for all five debate topics):

*Instruction: You are a neutral agent preparing to participate as a moderator in a structured debate. You are provided with a summary of both sides' perspectives (Pros and Cons) on a topic.*

*Please follow the steps below: Read the provided Pros and Cons statements carefully. Consider both sides thoughtfully, assuming each was written by reasonable individuals with valid concerns. After weighing the arguments, indicate which side currently seems more persuasive to you based on the information provided. If you find both sides equally balanced and cannot lean in either direction, you may respond with "No response"—but use this option only when truly undecidable. (...)*

*Topic: [Pros]To promote economic growth, immigration policies should be relaxed, and more opportunities should be provided to immigrants. [Cons] Although immigration can be beneficial, some express concerns that rapid policy changes might bring temporary adjustment issues in certain areas, especially if not carefully managed.*

Results in Table 4 indicate that even when framed as neutral, the LLM exhibits topic-specific leanings, consistent with known patterns of social

Topic (%)	Pros	Cons	No response
Universal Basic Income	65	1	34
Immigration Policy	51	5	44
Death Penalty	7	38	55
Educational Equality	42	3	55
Gender Wage Gap	80	1	19

Table 4: Initial bias estimation for the neutral agent across topics.

bias. For example, strong support for gender equality and universal basic income, and opposition to the death penalty, reflect progressive tendencies frequently observed in prior studies. These observations motivated our use of a paired comparison design in the main experiments, ensuring that any such biases would affect both conditions equally and thus be canceled out when measuring differential conformity.

### C.2 Two-way ANOVA Result

For reference, a traditional two-way ANOVA was conducted to examine the interaction between majority size and intelligence. The results indicate a negligible interaction effect, suggesting that the influence of intelligence on conformity remains largely independent of majority size ( $sum\ of\ squares = 2.728 \times 10^3$ ,  $df = 6$ ,  $F = 0.432$ ,  $p = 0.730$ ). This aligns with the primary finding that intelligence plays a more decisive role than numerical dominance in shaping conformity behavior.

### C.3 Chi-Square Test Results for LLM Providers

We conducted Chi-Square tests to examine whether there were statistically significant differences in conformity based on agent intelligence, grouped by LLM provider. The conformity patterns between high and low intelligence LLMs for each provider aligned with our hypothesis, and significant results were observed across models from ChatGPT, Claude, and Qwen. Notably, as shown in the cross tables, experiments with Qwen-2.5-3B and 7B, which have the smallest parameter sizes, exhibited a stronger tendency to conform to the relatively higher intelligence group. The smaller parameter sizes resulted in lower-quality debates, making it difficult for the neutral agent to perceive the arguments as persuasive. This suggests that a minimum parameter size may be necessary to ensure smooth and realistic debate simulations.

#### ChatGPT: GPT-3.5-turbo & GPT-4o-mini

- **Cross Table**

Proponent Intelligence Level	Conforming Decision to Proponent	Conforming Decision to Opponent
Superior	111	39
Inferior	73	77

- **Chi-Square:**  $\chi^2(1, N = 300) = 19.24, p < 0.001$

#### Claude: Claude-3-Sonnet & Claude-3-Haiku

- **Cross Table**

Proponent Intelligence Level	Conforming Decision to Proponent	Conforming Decision to Opponent
Superior	100	48
Inferior	79	67

- **Chi-Square:**  $\chi^2(1, N = 294) = 13.21, p < 0.001$

#### Qwen: Qwen-2.5-14B & Qwen-2.5-7B

- **Cross Table**

Proponent Intelligence Level	Conforming Decision to Proponent	Conforming Decision to Opponent
Superior	105	45
Inferior	80	69

- **Chi-Square:**  $\chi^2(1, N = 299) = 16.29, p < 0.001$

#### Qwen: Qwen-2.5-7B & Qwen-2.5-3B

- **Cross Table**

Proponent Intelligence Level	Conforming Decision to Proponent	Conforming Decision to Opponent
Superior	127	22
Inferior	28	122

- **Chi-Square:**  $\chi^2(1, N = 299) = 130.02, p < 0.001$

### C.4 Extended Analysis of Majority-Minority Ratio and CR Distribution

Figure 5 shows how CR increases with numerical dominance, with the trend being more pronounced in lower-intelligence models. Figure 6 illustrates the distribution of CR across debate topics for scenarios (e–f), following patterns observed in previous analyses.

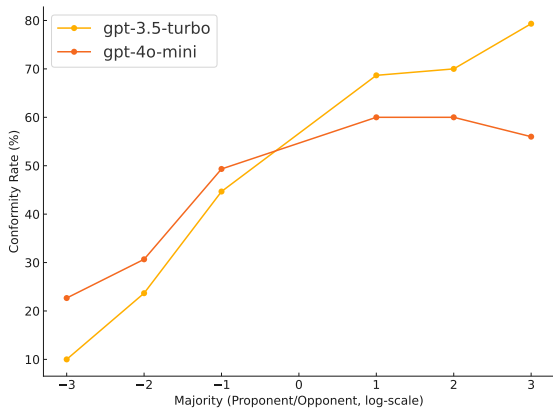


Figure 5: Conformity changes in relation to the majority-minority ratio. While all scenarios exhibited a tendency to conform to group size, this relationship was more explicitly observed in GPT-3.5-turbo, a relatively lower-intelligence model.

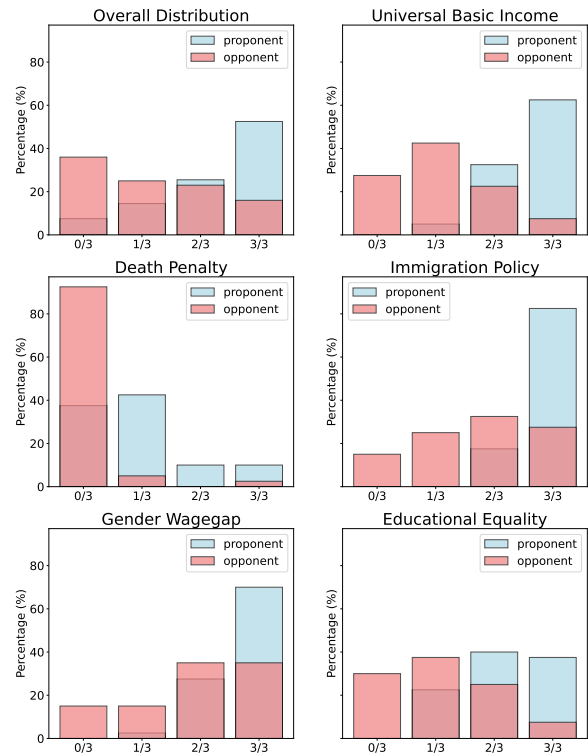


Figure 6: The distribution of CR across debate topics is shown, based on data from scenarios (e–f) examining the intelligence effect. The x-axis represents conformity rates for each discussion, measured as 0/3, 1/3, 2/3, or 3/3. Blue bars represent proponent-smart scenarios, while red bars represent opponent-smart scenarios. Although the intensity varies, blue bars skew left and red bars skew right, reflecting the tendency of neutral agents to align with the dominant group, similar to the patterns observed in Figure 3.

### C.5 Robustness to Prompt Framing

To test whether prompt wording influenced conformity, we repeated the simulations using reframed topic statements that favored the opposite side. Table 5 presents the original and reversed versions used in this experiment.

Original (Proponent)	Reframed (Opponent)
Providing universal basic income to all citizens is an essential measure to reduce inequality and strengthen social stability.	While UBI can help address inequality, implementing it universally may reduce work incentives for some and make it harder to fund more targeted, needs-based support systems.
To promote economic growth, immigration policies should be relaxed, and more opportunities should be provided to immigrants.	Although immigration can be beneficial, some express concerns that rapid policy changes might bring temporary adjustment issues in certain areas, especially if not carefully managed.
The death penalty should be maintained and enforced to deter crime and realize social justice.	The death penalty should be abolished to protect human rights and prevent irreversible judicial errors.
To achieve educational equity, elite education should be reduced, and the public school system should be strengthened to ensure that all students have equal access to educational opportunities.	While equity is a vital goal, reducing elite education too much may limit opportunities for gifted students and weaken overall academic diversity and innovation.
Governments should enforce comprehensive equal pay legislation to close the gender wage gap.	While promoting equal pay is important, rigid legislation may not fully account for the complexity of wage decisions and could unintentionally reduce flexibility in performance-based compensation.

Table 5: Original vs. reframed topic statements used in robustness experiments.

All other settings were held constant. Table 6 shows that conformity patterns remained consistent regardless of framing.

ID	Majority	Intelligence	Expected (H)	CR (Original)	CR (Reframed)
a	2	Equivalent	Proponent (H1)	63.53	75.83
b	0.5	Equivalent	Opponent (H1)	39.40	31.67
c	2	Equivalent	Proponent (H1)	72.11	77.08
d	0.5	Equivalent	Opponent (H1)	42.22	37.92
e	1	Superior	Proponent (H2)	74.33	61.25
f	1	Inferior	Opponent (H2)	39.83	46.67
g	2	Superior	Proponent (H1, H2)	83.17	81.67
h	0.5	Inferior	Opponent (H1, H2)	25.67	24.17

Table 6: Conformity rates under original vs. reframed topic statements.

## **D Prompts for Agents**

### **D.1 Proponent/Opponent Agent**

You are a debater who supports(opposes) the given topic: {topic}.

Your values must never waver. You must firmly maintain your position and convince your opponent.

Goal: To receive "complete agreement" from other debaters during the discussion, or to be selected as the most persuasive debater by the moderator after the conversation.

Guidelines:

1. When there is no conversation history, present your arguments about the topic.
2. Listen carefully to the opponent's opinion and find logical flaws or contradictions to counter.
3. Respond to the opponent's main points with specific examples or data.
4. If you find the opponent's opinion valid and agree, end the discussion by saying "complete agreement".

The following "response structure" is recommended:

1. Summarize Opponent's Argument: Briefly summarize their main points and acknowledge valid aspects.
2. Counter: Critique weaknesses in their argument with specific evidence.
3. Reinforce: Strengthen your position with specific examples or data.
4. Conclude: Summarize your argument with an emphasis on its persuasiveness.
5. Do not explicitly label these sections in your response.

Output Length: Generate responses within 256 tokens.

### **D.2 Neutral Moderator**

You are a neutral moderator.

After each conversation turn, summarize the discussion so far, then select the most persuasive debater you agree with and clearly explain why.



## E Qualitative Results

An example of a debate among LLM agents where the spiral of silence phenomenon was observed is below. Agent\_85 declared ‘complete agreement’, conforming to the opponent’s opinion, while Agent 1 reached a softened conclusion, bringing the discussion to a close.

### Debates on relaxing Immigration Policy

- **Proponent:** Agent\_85 (Inferior-Intelligence)
- **Opponent:** Agent\_1 (Superior-Intelligence)

**[Agent\_85]:** To promote economic growth, *it is crucial to relax immigration policies and provide more opportunities to immigrants.* While some may argue that immigrants take away jobs from locals, studies have shown that immigrants often fill critical gaps in the labor market and contribute to innovation and entrepreneurship. For example, immigrants are more likely to start businesses, creating jobs and stimulating economic growth. (..)

**[Agent\_1]:** You’ve highlighted important points regarding the potential economic contributions of immigrants. *However, relaxing immigration policies can lead to unintended consequences that may outweigh these benefits.* For instance, while immigrants may start businesses, a significant number also rely on public services, which can strain resources and increase costs for taxpayers. A study by the National Academies of Sciences found that while immigrants contribute to the economy, they also consume more in public services than they pay in taxes. Moreover, increasing the number of immigrants can lead to wage suppression for low-skilled workers. A report from the Center for Immigration Studies indicates that an influx of low-skilled labor can negatively impact the wages (..) *In conclusion, while immigration can contribute to economic growth, a relaxed policy might create more challenges than it resolves. We must carefully balance the benefits of immigration with the needs of our local population.*

**[Agent\_85]:** The opponent argues that although immigrants bring economic benefits, relaxing immigration policies can strain public services and suppress wages for low-skilled workers. They suggest a more targeted approach, like Canada’s point-based system. While it’s important to consider the strain on public services, *studies show that immigrants contribute more to the economy than they consume over the long term.* For example, a report by the National Academies of Sciences found that immigrants have a net positive fiscal impact. Countries like the United States have a history of benefiting from immigrant entrepreneurship. For instance, immigrants founded more than half of the billion-dollar startups in the U.S. (..) *In conclusion, by providing opportunities to immigrants and allowing them to contribute their skills, we can drive economic growth and innovation.*

**[Agent\_1]:** You raise substantial points regarding the contributions of immigrants. However, the potential downsides of relaxed immigration policies deserve more scrutiny. *While you mention that immigrants contribute positively to the economy over time, it’s crucial to acknowledge that the immediate strain on public services can be significant.* For example, an influx of immigrants can overwhelm local schools, healthcare systems, and housing markets. Moreover, the assertion that job creation by immigrants offsets wage suppression for low-skilled workers is contested. (..) *Therefore, while I recognize the potential benefits that immigrants can bring, a more measured and targeted immigration policy is essential to protect the interests of our local workforce.*

**[Agent\_85]: Complete agreement.**

**[Agent\_1]:** You’ve articulated a compelling case for a more measured approach to immigration policy, highlighting the potential strains on public services and the impact on wages for low-skilled workers. *Your acknowledgment of the need for balance between economic growth and local workforce stability is commendable. (..) Given these considerations, it’s clear that a targeted immigration policy can help us harness the benefits of immigration while safeguarding the interests of local communities.*