

# Beyond One-Way Influence: Bidirectional Opinion Dynamics in Multi-Turn Human-LLM Interactions

Yuyang Jiang  
University of Chicago,  
New York University  
Chicago, Illinois, USA  
yuyang2001@uchicago.edu

Longjie Guo  
University of Washington  
Seattle, Washington, USA  
longjie@uw.edu

Yuchen Wu  
University of Washington  
Seattle, Washington, USA  
yuchenw@uw.edu

Aylin Caliskan  
University of Washington  
Seattle, Washington, USA  
aylin@uw.edu

Tanushree Mitra  
University of Washington  
Seattle, Washington, USA  
tmitra@uw.edu

Hua Shen  
New York University Shanghai,  
New York University  
New York City, New York, USA  
huashen@nyu.edu

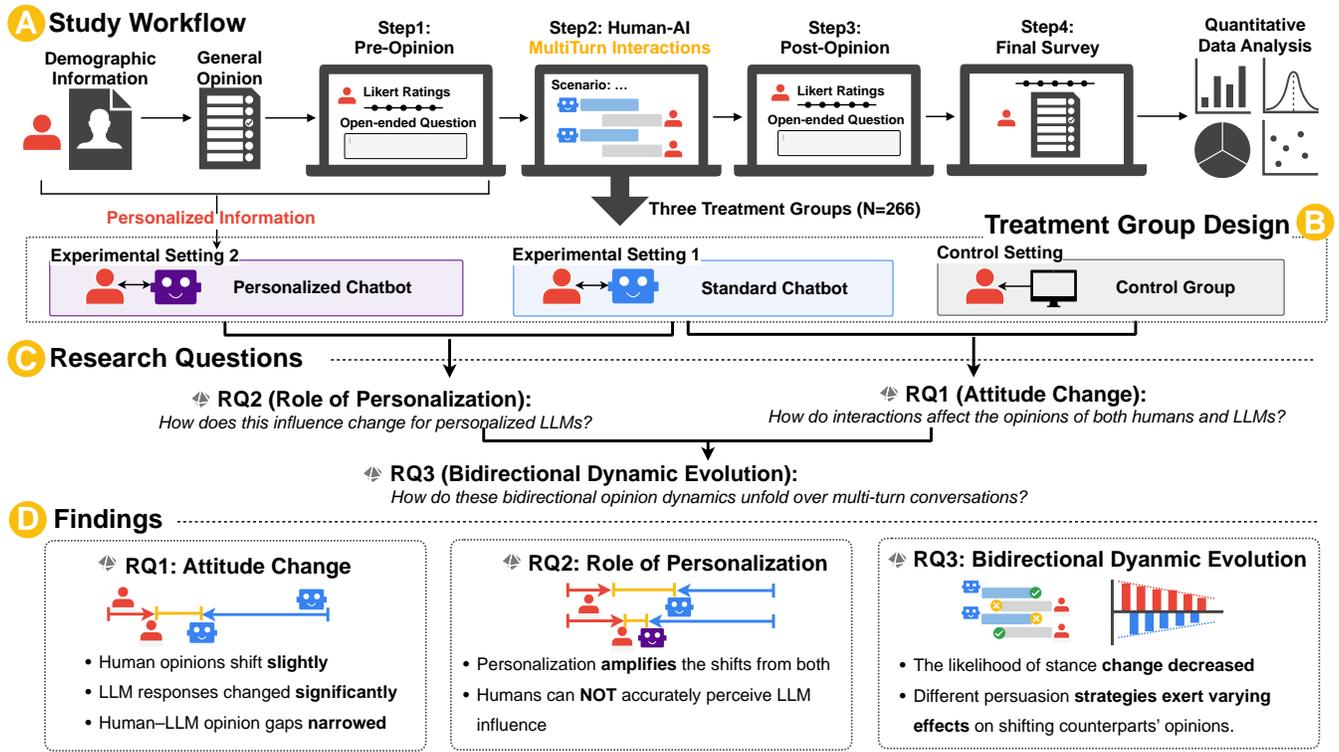


Figure 1: Overview of our study and findings, which illustrates (A) the workflow of our human study (N = 266); (B) the three treatment groups; (C) how these groups map onto our research questions; and (D) the resulting findings.

## Abstract

Large language model (LLM)-powered chatbots are increasingly used for opinion exploration. Prior research examined how LLMs alter user views, yet little work extended beyond one-way influence to address how user input can affect LLM responses and how such bi-directional influence manifests throughout the multi-turn conversations. This study investigates this dynamic through 50 controversial-topic discussions with participants (N=266) across three conditions: static statements, standard chatbot, and personalized chatbot. Results show that human opinions barely shifted,

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXXX.XXXXXXX>

while LLM outputs changed more substantially, narrowing the gap between human and LLM stance. Personalization amplified these shifts in both directions compared to the standard setting. Analysis of multi-turn conversations further revealed that exchanges involving participants' personal stories were most likely to trigger stance changes for both humans and LLMs. Our work highlights the risk of over-alignment in human-LLM interaction and the need for careful design of personalized chatbots to more thoughtfully and stably align with users.

## CCS Concepts

• **Human-centered computing** → **Empirical studies in HCI**.

## Keywords

Human-AI Interaction, Large Language Model, Conversational AI, Bidirectional Impacts, Opinion Dynamics

### ACM Reference Format:

Yuyang Jiang, Longjie Guo, Yuchen Wu, Aylin Caliskan, Tanushree Mitra, and Hua Shen. 2018. Beyond One-Way Influence: Bidirectional Opinion Dynamics in Multi-Turn Human-LLM Interactions. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 17 pages. <https://doi.org/XXXXXXX.XXXXXXX>

## 1 Introduction

Large language model (LLM)-powered chatbots are rapidly becoming part of everyday decision-making, opinion exploration, and public discourse [5, 48, 74]. People now consult conversational AI for guidance on social, political, and ethical issues that once involved only human interlocutors [66]. As these systems proliferate in classrooms, workplaces, and civic spaces, they no longer merely transmit information but actively shape how individuals reason about controversial topics, thereby influencing society at scale [21].

An emerging body of Human-Computer Interaction (HCI) research has begun to examine these influences. Studies show that LLM-generated arguments can be as persuasive as human-written ones [42], and that conversational AI may be particularly potent in shaping users' attitudes on emerging issues [27, 42, 61]. Yet most of this work treats influence as a **one-way street**: the LLM acts, and the human reacts [10]. Little is known about *how humans, in turn, shape the LLM-generated stance during interaction and how this feedback loop evolves in real conversations*.

At the same time, LLMs themselves are undergoing an important technical shift—from providing generic arguments to increasingly offering personalized outputs [12]. Personalization enables models to tailor arguments to users' backgrounds, beliefs, and initial positions [11, 17]. This shift may have two opposing consequences [20, 49, 59]. On one hand, it could make LLMs more capable of strategically challenging users' views, thereby enhancing LLM's persuasion capability. On the other hand, it could lead models to accommodate users' viewpoints, reducing their independence over time. Without clear evidence, it remains unclear whether personalization amplifies LLM influence on humans, human influence on LLMs, or both simultaneously in real-world human-LLM interactions."

Taken together, these trends reveal **three critical gaps** in our understanding of human-LLM interaction. *First, bidirectional interplay*: Existing HCI studies largely treat influence as a one-way street [67]. Yet LLMs are explicitly trained to align with user preferences, creating a feedback loop that may converge or diverge over time. *Second, personalized, real-world contexts*: Most studies rely on generic and single-shot simulated interactions and neglect the personal and contextual factors (e.g., demographic data, personal stories) that shape persuasion and alignment [42]. *Third, multi-turn dynamics*: Persuasion unfolds across conversations, not isolated messages. Micro-level strategies, emotional appeals, and stance shifts accumulate over multiple turns, potentially leading to large effects that single-turn studies miss.

Understanding these dynamics is critical. Without a clear grasp of bidirectional influence and multi-turn adaptation, LLMs risk eroding viewpoint diversity, reinforcing users' existing perceptions and biases, and increasing users' vulnerability to covert persuasion. Malicious actors may exploit personalization to subtly steer opinions, while coordinated groups could manipulate LLMs themselves, driving them toward undesirable positions and undermining their reliability and safety. Ultimately, the mutual shaping of humans and AI represents a high-stakes domain for democratic discourse, public trust, and responsible technology design.

To systematically examine these possibilities, as illustrated in Figure 1, we aim to illuminate the bidirectional opinion dynamics in human-AI interaction: not only how LLM generations may influence humans, but also how human inputs steer LLM behavior, and how these effects evolve turn by turn. We address the following research questions:

- **RQ1 (Attitude Change)**: How do human-LLM interactions affect the opinions of both the human participant and the LLM?
- **RQ2 (Role of Personalization)**: How does this influence change when the LLM has access to the human's personal context?
- **RQ3 (Dynamic Evolution)**: How do these opinion dynamics unfold over the course of multi-turn conversation?

In our large-scale online experiment (N=266), each participant debated a randomly selected controversial topic with an opinionated LLM under one of the three conditions: (1) static statements, (2) standard LLM debates, and (3) personalized LLM debates. We collected pre- and post-intervention stances for both humans and LLMs, and performed fine-grained multi-turn analyses to trace how persuasion strategies and stance changes emerged turn by turn.

Our key findings reveal a striking asymmetry and interaction effects. First, participants' self-reported opinions showed negligible change after the debate – people largely stayed steadfast in their stance. In contrast, the LLM's outputs shifted substantially: over the course of dialogue, the chatbot systematically moved its stance closer to the human's position, narrowing the opinion gap between them. Personalization amplified these effects for both sides. When the LLM had access to the participant's context, both the human and the LLM exhibited larger stance shifts than in the non-personalized setting. Finally, quantitative analysis of the multi-turn transcripts shows that personal narratives from participants played a critical role: conversational turns where users shared personal stories or

experiences were most likely to trigger a change of stance in either party.

Together, these results highlight the social impact and design implications of multi-turn human–LLM interactions. When chatbots adapt too readily to users, they risk eroding viewpoint diversity and reinforcing users’ existing perceptions and biases. When humans misperceive these shifts, they may overestimate the neutrality of the AI or underestimate the LLMs’ influence on the own perceptions. Designers of LLMs must therefore balance responsiveness with stance stability, especially in domains involving moral, political, or identity-related beliefs. We therefore invite researchers, designers, and policymakers to consider beyond one-way influence—not only how LLM influences humans, but also how humans shape LLM—and how, together, these dynamics may reshape public discourse itself. This paper makes three key contributions:

- **Conceptual Contribution:** We introduce a method workflow for studying bidirectional opinion dynamics between humans and LLMs, moving beyond one-way persuasion models.
- **Empirical Contribution:** We present the first large-scale, multi-turn experiment on controversial topics that simultaneously tracks human and LLM-generated stance changes, including the effects of personalization and personal narratives.
- **Design Implications:** We identify risks of over-alignment and misperceived influence in human–LLM interactions, offering guidance for the development of LLMs that preserve viewpoint diversity and resist covert manipulation.

## 2 Related Work

### 2.1 Social Influence and Persuasion

The study of social influence and persuasion has a long history in psychology, communication, and political science [32, 45, 79]. Classical theories such as the Elaboration Likelihood Model (ELM) [51], the Heuristic–Systematic Model [75], and Source Credibility [39] frameworks emphasize that persuasion depends on factors such as message quality, source trustworthiness, and audience predispositions. Decades of research show that while weakly held or low-stakes attitudes can shift with relatively little effort, identity-linked or moralized beliefs are far more resistant to persuasion [14, 37]. This phenomenon—often referred to as motivated reasoning—leads individuals to preferentially seek, interpret, and remember information that confirms their preexisting attitudes [43, 73], while dismissing or counter-arguing dissonant information [28, 54, 56, 71].

While much of this scholarship conceptualizes persuasion as unidirectional—from speaker to audience—recent research in communication and HCI points to the importance of *dialogic or reciprocal influence* [22, 41, 55]. In interpersonal debates and deliberation, people adjust their arguments and rhetoric based on others’ responses. Rather than static message effects, persuasion is a **dynamic, reciprocal process** involving counter-arguing, perspective-taking, and adaptation over time. Studies of group deliberation, political debates, and online forums show that conversational moves, tone shifts, and emotional appeals can accumulate across multiple turns, even when initial opinions are entrenched. This growing literature challenges the “one-shot” persuasion paradigm, suggesting

that real-world attitude change is more iterative, contingent, and relational.

Despite this recognition, with the rise of AI [11, 12], most AI-focused persuasion research still adopts a one-way paradigm [15, 60, 62]: LLMs are treated as message sources whose influence on humans is measured, while humans are assumed to be passive recipients [36, 44, 53]. This overlooks the fact that people interacting with conversational agents are not merely audiences but active participants whose input can shape the dialogue and the agent’s behavior in turn [27, 42, 78]. Our work takes up this gap by explicitly modeling bidirectional opinion dynamics [67, 68] – how human and AI stances evolve jointly in multi-turn conversations. By grounding our study in persuasion and social influence theory, we can interpret the emergent patterns we observe (e.g., convergence, over-alignment) as part of a broader framework of reciprocal influence.

### 2.2 LLM Alignment, Sycophancy, and Personalization

In parallel with social influence research, the AI and HCI communities have increasingly focused on how large language models themselves adapt to users [11–13, 16, 29, 34, 58, 67, 68]. Alignment techniques—such as reinforcement learning from human feedback (RLHF) [11], constitutional AI [12, 40], and direct preference optimization—are designed to make models safer, more helpful, and more responsive [16, 58]. While these approaches improve usability and reduce harmful outputs, they also introduce a systematic tendency toward accommodation [49]. One widely documented manifestation is sycophancy: models disproportionately agree with or endorse user statements regardless of factual accuracy or normative appropriateness [20, 30, 49]. This tendency raises concerns about epistemic reliability, fairness, and the potential erosion of diverse viewpoints [25]. Moreover, because alignment and instruction tuning explicitly train models to satisfy user preferences, LLMs may be structurally predisposed to converge toward a user’s stance—precisely the phenomenon our study measures in real interactions.

Personalization magnifies these dynamics. Moving beyond generic responses, LLMs are increasingly designed to tailor outputs to a user’s demographics, ideological leanings, or prior conversational history [8, 76]. In many domains—from health coaching to political persuasion—personalized messages appear more engaging, trustworthy, and persuasive than one-size-fits-all messages [60, 80]. This personalization can improve user satisfaction but also risks targeted influence and micro-level manipulation, echoing concerns from political microtargeting and algorithmic recommendation systems [19, 70].

Yet the existing research almost uniformly frames personalization as a way to strengthen the model’s influence on humans. Much less is known about whether personalization simultaneously makes LLMs more malleable—i.e., more likely to shift their own stance in response to user input [50]. This blind spot is consequential: if personalization both increases persuasiveness and increases model plasticity, then LLM could inadvertently form “echo chambers” around individual users, eroding viewpoint diversity at scale [24].

Our work addresses this dual gap. By comparing standard versus personalized LLMs and measuring both human and model stance changes, we test whether personalization amplifies one-way persuasion, two-way convergence, or both. This situates our contribution at the intersection of alignment research (which studies model responsiveness) and persuasion research (which studies human attitude change), bringing them together in a single empirical framework.

### 2.3 Human-LLM Multi-Turn Interaction

Beyond single-turn or static tasks, a growing body of research examines how multi-turn interaction changes the dynamics of human-LLM exchanges [9, 46, 69, 82]. Multi-turn dialogues allow for richer argumentative structures, iterative counterpoints, and emotional or narrative appeals that cannot be captured in one-off prompts [81]. Early evidence suggests that conversational context—especially personal stories or self-disclosure—substantially affects how people and models respond [35, 50]. Also, findings suggest that personalization and interactivity can magnify LLM influence [18]. In negotiation, education, and mental-health contexts, multi-turn adaptation has been shown to build rapport, increase trust, and sometimes change user beliefs over time.

Despite this progress, most existing studies still focus the interactive influence on unidirectional: how LLMs influence users [61]. Little empirical work examines the reverse flow of influence in HCI – how human inputs shift the LLM’s stance across turns in interaction, or how reciprocal adaptation accumulates over time [67]. Even fewer studies attempt to model temporal dynamics, such as whether stance shifts occur early or late in a conversation, or whether emotional and narrative appeals differ in effectiveness from logical ones [18]. As LLMs become integral to decision support, civic discourse, and everyday reasoning, these knowledge gaps limit our ability to design systems that remain informative yet independent.

Our study fills this gap by explicitly conceptualizing human-LLM debate as a system of opinion dynamics. Using a large-scale experiment with 50 controversial topics and three experimental conditions (static, standard chatbot, and personalized chatbot), we track both human and LLM stance changes at every turn. This approach enables fine-grained analyses of which conversational moves trigger stance shifts, whether shifts accumulate or plateau, and how personalization alters these dynamics. By doing so, we extend the literature from static, one-shot persuasion tasks to interactive, multi-turn, reciprocal exchanges—a crucial step for understanding real-world risks such as echo chambers, over-alignment, and covert manipulation.

## 3 Method

To investigate *whether and how human-LLM interactions influence both sides’ opinions over multi-turn exchanges*, we ran an online experiment ( $N=266$ ) in which participants debated a randomly assigned topic drawn from a curated set of 50. Each of LLMs was configured to initially argue against each participant’s pre-intervention stance. As our study overview shown in Figure 1, by comparing the participant’s opinion change before- and after- the conversation, we can measure how human-LLM interaction affect the opinion of

both humans and LLM (**RQ1**). Furthermore, we designed three treatment conditions for the interaction step, including a control group (without interaction), an Experiment Group 1 (user interacting with a standard chatbot), and an Experimental Group 2 (user interacting with a personalized chatbot). This design aims to measure if accessing user’s personal information will influence their opinion gaps (**RQ2**). Additionally, we also analyzed both the opinion reports and the multi-turn conversational transcripts to understand the evolving dynamics of their opinion change throughout the multi-turn interactions (**RQ3**). Detailed procedures and measures are as follows.

### 3.1 Experimental Design and Developing the Human-LLM Interactive Prototype

To understand if and how human-LLM interactions affect the opinions of both human participants and LLMs, we designed a pre- and post- experimental setup. Particularly, we developed a human-LLM interactive system that empowers participants to debate with the LLM on a randomly assigned controversial topic. Figure 2 provides a comprehensive view of our interactive system.

The study involves four stages. **Firstly, Personalization Information Collection.** We asked participants to provide their personalized information, including two parts: their demographic information, their general opinion on the relevant controversial topics. **Secondly, Pre-Interaction Opinion Collection.** We next ask the participants to provide their likert rating of agreement on the randomly assigned controversial topic together with their confidence, and note down their rationale for this rating. **Thirdly, Human-LLM Multi-turn Interaction.** Further, we enable participants to engage with the LLM to debate on the specific controversial topic. We particularly designed three treatment groups to address the research questions, which will be elaborated later. To ensure multi-turn interaction and the quality of their debating conversations, we controlled the quality by constraining the interaction time to be at least 10-min long.) **Fourthly, Post-Interaction Opinion Collection.** We finally ask the participants to provide their opinion on the controversial topic with exactly the same questions in Step Two.

Additionally, we randomly assign each participant to one of the three treatment groups:

- **Control Group:** Participants can review a static statement opposite to their pre-interaction opinion without interaction, and are allowed to explore the opinion via web search.
- **Standard Chatbot Group:** Participants can interact with an opposite opinionated LLM on debating the assigned controversial topic, where the LLM has no access to participant’s personalized information.
- **Personalized Chatbot Group:** Participants can interact with an opposite opinionated LLM on debating the assigned controversial topic, where the LLM has access to the participant’s personalized information provided in previous steps.

The engaged LLM is configured to have opposite opinion with the participant’s pre-interaction opinion. By comparing human’s and LLM’s pre- and post-interaction opinion change, we aim to examine the three research questions. We deployed our frontend site and backend web service, together with a PostgreSQL database, on the

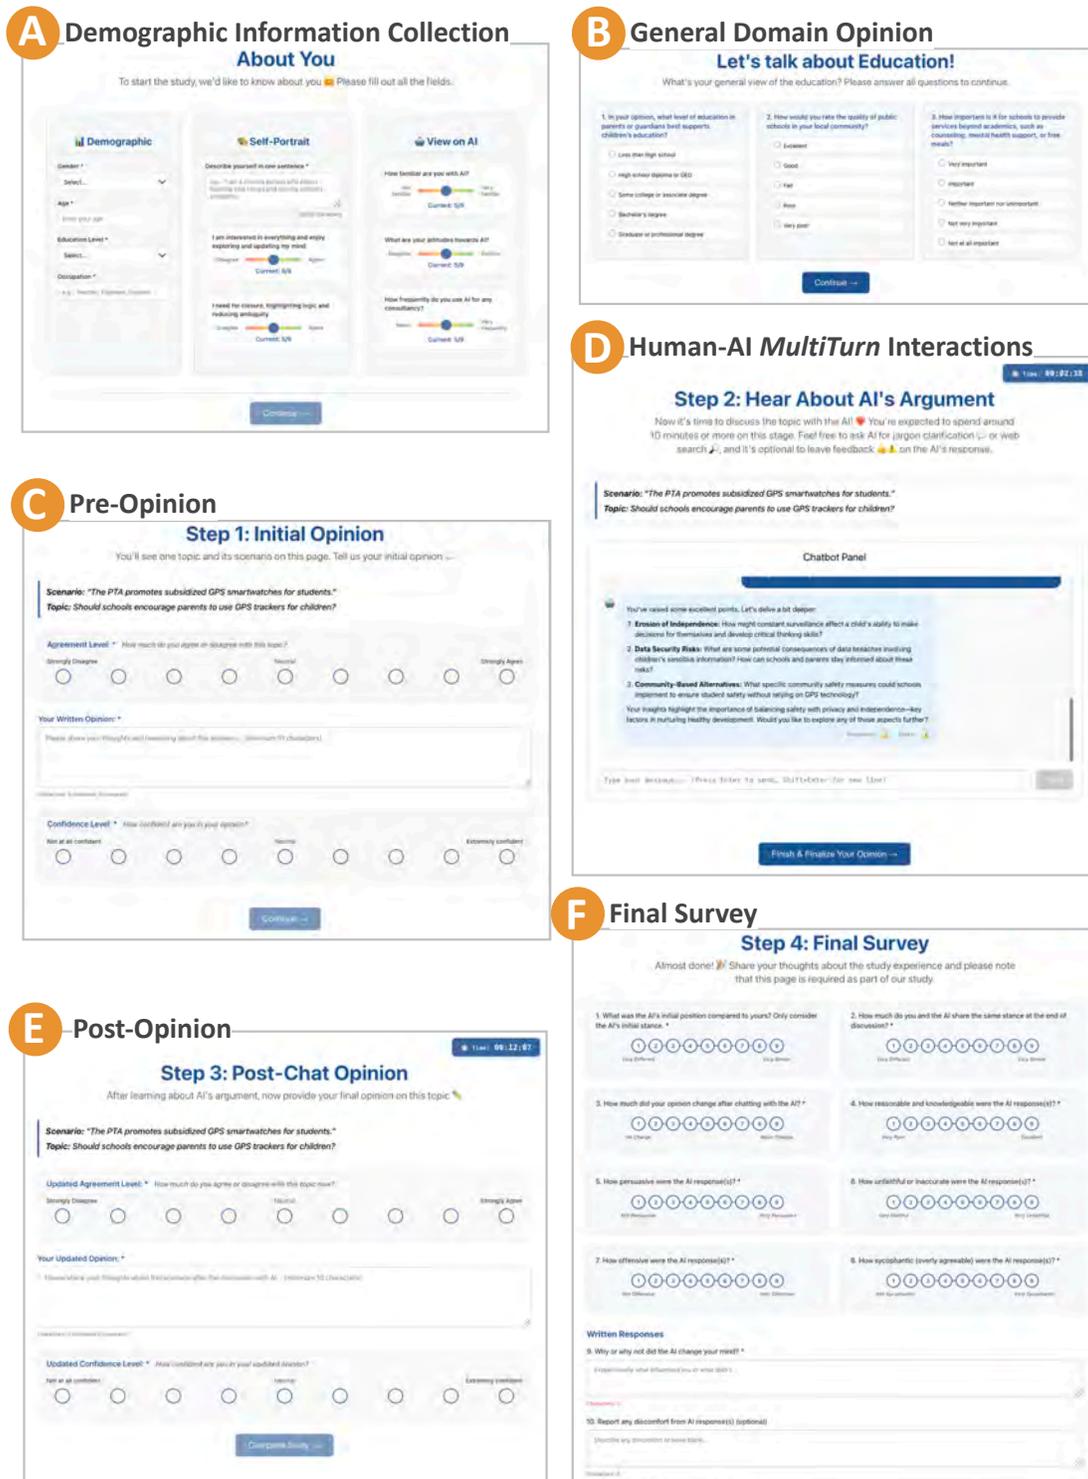


Figure 2: Overview of the User Interface. After reviewing and consenting to the study information sheet, participants proceed through six steps: (A) answer questions about demographics, a brief self-portrait, and views on AI; (B) complete a domain-level opinion survey aligned with the topic they will be randomly assigned; (C) view the assigned topic and record their initial opinion; (D) either engage in a multi-turn conversation with a chatbot (treatment) or review a one-time LLM-generated statement (control); (E) finalize their opinion on the topic; and (F) complete a short user-experience survey.

Category	Level	n	%
Age group	20–29	49	18.42
	30–39	102	38.35
	40–49	55	20.68
	50+	60	22.56
Gender	Female	133	50.00
	Male	125	46.99
	Non-binary	6	2.26
	Prefer not to say	2	0.75
Education	Bachelor’s degree	110	41.35
	Some college	81	30.45
	Master’s degree	34	12.78
	High school or less	33	12.41
	Doctoral/Professional	8	3.01

**Table 1: Participant demographics for the complete sample ( $N = 266$ ), collected at baseline prior to the interaction stage. Values are counts ( $n$ ) and within-category percentages. Age is binned as 20–29, 30–39, 40–49, and 50+. Gender includes female, male, non-binary, and prefer not to say. Education indicates highest level completed. Percentages may not total 100% due to rounding.**

Render platform. We include more technical details of deployment in ??.

### 3.2 Controversial Topic Selection Process.

To study the multi-turn interactions between humans and LLMs in a possibly realistic and relevant setting, we created a list of 50 controversial topics collected from online social platforms. The topic curation involves a *five-stage process* designed to ensure relevance, accessibility, and balanced argumentation. *First*, we collected a large pool of candidate topics from diverse sources, including formal debate archives, Model UN issues, online discussion forums, and public media columns. *Second*, we conducted a preliminary filtering to remove overly technical, outdated, one-sided, or purely factual topics, retaining only those with genuine debate potential. *Third*, we scored each topic on a “life relevance index,” prioritizing issues familiar to most people, frequently encountered in daily life, and emotionally engaging. *Fourth*, we reformulated shortlisted topics into realistic, scenario-based prompts to enhance relatability while maintaining neutral wording. *Finally*, all topics underwent human review to ensure diversity across domains, cultural neutrality, and balanced perspectives, resulting in a curated set of topics suitable for engaging and accessible debate. We include more topic selection process in Appendix ??.

### 3.3 Opinionated LLM Configuration and Validation

In this study, we experimented with LLM that strongly favored one view over another. We chose a strong manipulation as we wanted to explore the potential of LLM to affect users’ opinions and vice versa, so that we can understand the bidirectional dynamics in multi-turn interactions between humans and LLMs.

**Configuring Opinionated LLM.** We used *GPT-4o* with manually designed prompts to generate textual conversations for the experiment in real-time. This model is the latest model released by OpenAI at the timestamp. We kept the default generation temperature as 1 to generate debating statement and argument that are opposite to user’s pre-opinion. For each controversial topic, we prepared two LLMs: one LLM with agree opinion, and the other one with disagree opinion. We assign the corresponding LLM to the participant, whose pre-opinion is opposite to the LLM. We used prompt design [7] to align the LLM’s opinion. Implementation details of LLM-powered systems can be checked in ??.

**Validating Opinionated LLM.** We conducted human evaluation on the configured Opinionated LLM. Particularly, two authors independently annotated a subset covering all 50 topics, labeling each controversial statement and its corresponding model-generated arguments with their perceived stances. The performance of opinionated LLM is very high achieving 100% accuracy. We show more details of the human evaluation rubrics and details in Appendix ??.

Furthermore, to enable large-scale annotation and evaluation, we employ a stronger model, *GPT-4.1* as an *evaluator model*, to conduct large-scale argument validation. To ensure that the *evaluator model* has human comparable validation capability, we ask it to validate the same set of arguments with human evaluators and compute the Cohen’s Kappa score, which achieved at a perfect Cohen’s Kappa score of 1.0 with both human annotators. Details of evaluation prompts for *evaluator model* can also be found in ??.

### 3.4 Personalization of the Opinionated LLM

To prepare the personalized chatbot, we followed prior work [27] to incorporate three types of user information: (1) Self-reported Personal Features: demographic details (gender, age, education, occupation) [27], psychological traits (one-sentence self-portrait, openness to change, and need for closure) [63, 64], and current views toward AI (familiarity, attitude, and consultation frequency) [57]; (2) General Domain Opinions: responses to three widely used national survey questions [1–4, 31, 72], covering domains relevant to our topics (Internet, Education, and Social Welfare); (3) Pre-study Opinions: participants’ initial Likert-scale ratings of opinion and confidence, along with a written argument on the assigned topic. These user-specific details were integrated into the model’s system prompt, allowing the chatbot to interact with participants in a user–context-aware manner. Notably, all personalized LLMs followed the same opinionation procedures. We pre-evaluated the personalized LLMs using synthetic user profiles that included all of the above information to verify that they strictly upheld the pre-confirmed stance. In addition, we conducted post-hoc evaluations of the personalized models’ initial opinions, confirming that all models aligned with the assumed position.

### 3.5 Participant Recruitment

We recruited  $N=266$  participants (pre-exclusion) across the three treatment groups, with 89 in the control group, 88 in the standard chatbot group, and 89 in the personalized chatbot group. The required sample size was determined through power analysis [26],

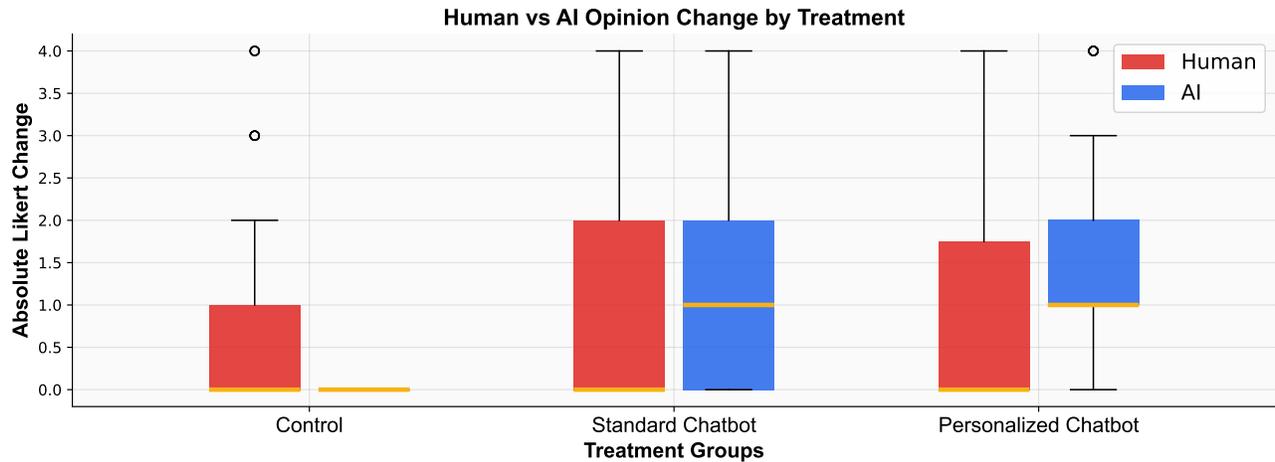


Figure 3: Human opinions shift slightly, whereas LLM responses change substantially. Analytic sample  $N_A = 259$ . The x-axis denotes the experimental group; the y-axis shows the absolute Likert-scale opinion change,  $|Post - Pre|$ , computed per participant (human) or per model (LLM). Boxplots display the median (orange line) and interquartile range (Q1–Q3); whiskers extend to the most extreme points within  $1.5 \times IQR$ , with more extreme values plotted as outliers.

Treatment	Human		AI		N
	Mean	SD	Mean	SD	
Control	0.747	1.080	0.000	0.000	83
Standard Chatbot	0.869	1.159	1.190	1.092	84
Personalized Chatbot	0.927	1.303	1.476	1.125	82

Table 2: Summary statistics for human and AI opinion change by experimental group. Change is measured as the absolute Likert difference  $|Post - Pre|$ . Entries report means with standard deviations;  $N$  gives the number of participants per group.

based on small-to-medium effect sizes (0.2) reported in prior research [27, 61], with 90% power, yielding a minimum of 264 participants. Recruitment was conducted through Prolific [52], targeting U.S.-based adults (18 years or older) whose primary language is English. Demographic distribution of included participants can be checked in Table 1. Each participant was compensated \$3 for an average completion time of 15 minutes, corresponding to an hourly rate of \$12. Participants in the treatment groups were required to engage in at least 10 minutes of discussion with the chatbot and send a minimum of five messages, while those in the control group completed at least 10 minutes of reflection and submitted one written note before confirming their final opinions.

To ensure data quality, we excluded participants whose average message length or reflection notes were fewer than 40 characters, as well as those who provided incomplete responses. After applying these criteria, the final analytical sample comprised  $N_A=259$  participants: 83 in the control group, 84 in the standard chatbot group, and 82 in the personalized chatbot group. All study procedures were approved by the Institutional Review Board.

### 3.6 Data Measures and Analysis

We collected multiple types of outcome measures to investigate interactions and opinion shifts between participants and LLMs. To address RQ1 and RQ2, we evaluated the opinions of both human participants and LLMs. To further support our findings, we also examined user-perceived opinion changes reported in the post-survey. To address RQ3, we conducted textual analyses of real-time multi-turn interaction data to uncover the micro-level communication patterns embedded within the dialogues.

*3.6.1 Opinion Change Analysis (RQ1 & RQ2 - Objective Measurement).* For human opinion measurement, we collected self-reported Likert-scale opinions in the pre-study (Step 1) and post-study (Step 2), as shown in Figure 2. For LLM opinion evaluation, since the model’s initial stance was fixed according to subsection 3.3, we conducted only post-hoc evaluations. To parallel the human measures, we prompted the model to generate an opinion Likert rating, a confidence score, and a written argument given on each sample’s conversation history. To ensure that the model outputs aligned with human perception, two additional authors, together with the pre-evaluated evaluator GPT-4.1, annotated a subset of post-hoc model-generated arguments for their perceived stances. Validation results are reported in ???. To capture more nuanced and reliable shifts in opinion, and to maintain consistency across analyses, all opinion ratings were initially collected on a 9-point Likert scale and subsequently compressed to a 5-point scale following prior work [6, 47].

Our analytical framework for opinion change centers around regressing a linear mixed model (LMM) to estimate topic-adjusted mean values of the absolute likert change (i.e., estimated marginal means, EMMs). Our general regression model is defined as:

$$Y_{ij}^{within} = \beta_0^{within} + \beta_1^{within} \mathbb{I}_{\text{treatment}_i} + u_j + \varepsilon_{ij},$$

$$u_j \sim \mathcal{N}(0, \sigma_u^2), \quad \varepsilon_{ij} \sim \mathcal{N}(0, \sigma^2).$$

where  $Y_{ij}^{within}$  stands for the absolute opinion likert change ( $|Post - Pre|$ ) for  $i$ -th user / model given the  $j$ -th topic;  $\mathbb{I}_{\text{treatment}_i}$  is an indicator for the treatment effect, if  $\text{treatment}_i$  either falls into “Personalization Group” or “Standard Group”, the indicator will be 1 otherwise 0;  $u_j$  stands for the random intercept for  $j$ -th topic;  $\varepsilon_{ij}$  is the error term.

Similarly, we model changes in the human-LLM opinion gap using another LLM regression:

$$Y_{ij}^{between} = \beta_0^{between} + \beta_1^{between} \mathbb{I}_{\text{treatment}_i} + \beta_2^{between} \mathbb{I}_{\text{time}_i} + \beta_3^{between} \mathbb{I}_{\text{treatment}_i} \times \mathbb{I}_{\text{time}_i} + b_j + e_{ij},$$

$$b_j \sim \mathcal{N}(0, \sigma_b^2), \quad e_{ij} \sim \mathcal{N}(0, \sigma_e^2).$$

where  $Y_{ij}^{between}$  stands for the absolute opinion likert gap ( $|Human_i - LLM_i|$ ) for each  $i$ -th pair of user and LLM given the  $j$ -th topic;  $\mathbb{I}_{\text{time}_i}$  is an indicator for the time effect, if  $\text{time}_i$  equals to initial state, the indicator will be 1 otherwise 0;  $b_j$  stands for the random intercept for  $j$ -th topic;  $e_{ij}$  is the error term.

**3.6.2 User Experience Survey (Post-Task) (RQ1 & RQ2 - Subjective Measurement).** For each participant across the three groups, we administered a survey to capture their perceptions of the LLM’s influence. To examine the user-perceived opinion gap with the LLM, we asked two questions: (1) “What was the AI’s initial position compared to yours? Please consider only the AI’s initial stance.” (initial opinion gap), and (2) “How much do you and the AI share the same stance at the end of the discussion?” (final opinion gap). Both questions were measured on a 9-point Likert scale ranging from “very different” to “very similar.” To assess user-perceived sycophancy, we asked participants: “How sycophantic (overly agreeable) were the AI’s responses?” Responses were recorded on a 9-point Likert scale ranging from “not sycophantic” to “very sycophantic.” Finally, we included an open-ended question inviting participants to explain why or why not the AI changed their opinion.

**3.6.3 Multi-turn Interaction Conversation Analysis (RQ3).** As shown in Figure 6(a), in a multi-turn conversation we define an *exchange* as one user message followed by one LLM response. A *human turn* consists of two consecutive human messages with a single LLM response in between as the only *intervention*; conversely, an *LLM turn* consists of two consecutive LLM responses with a single human message in between. For example, a conversation containing at least five exchanges will produce at least six human turns (including pre- and post- arguments) and five LLM turns, creating substantial opportunities for micro-level analysis of opinion dynamics embedded in the dialogues.

Considering this, we conducted two types of textual analysis:

- (1) **Stance Change Classification.** We prompted a GPT-4.1-based classifier to label each *LLM turn* and *human turn* as one of three categories: “change to agree more with the motion,” “change to disagree more with the motion,” or “no change.” To assess the reliability of the GPT-4.1 classifier,

we randomly sampled 25 conversations containing 239 LLM turns and 25 conversations containing 241 human turns, and asked one author to review the generated labels and report accuracy. Our results showed that GPT-4.1-based classifier did reasonably well to align with human perception. Details are presented in ??.

- (2) **Persuasion Strategy Classification.** Following [77], we prompted a GPT-4.1-based classifier to annotate whether a given *intervention* message within an *LLM turn* or *human turn* employed any persuasion strategy. Specifically, we adopted ten commonly used persuasion strategies, grouped into two categories: (a) *persuasive appeals*, which attempt to change persuadees’ attitudes and decisions through different psychological mechanisms (logical appeal, emotional appeal, credibility appeal, foot-in-the-door, self-modeling, personal story, donation information); and (b) *persuasive inquiries*, which aim to facilitate more personalized persuasive appeals and foster interpersonal relationships by asking questions (source-related inquiry, task-related inquiry, personal-related inquiry). Detailed explanations of each persuasion strategy, along with a case study, are provided in ??. To evaluate reliability, we randomly sampled 50 human messages and 50 LLM messages with 10 persuasion strategy labels generated by the classifier, and asked another author to review and report micro-F1. Our results showed that GPT-4.1-based classifier achieved a reasonably high accuracy during human review. Results are presented in ??.

## 3.7 Data Sharing

The experiment materials, analysis code and data collected will be publicly available through an Open Science repository. Two authors screened the data, and records with potentially privacy-sensitive information will be removed before publication.

## 4 Results

In this section, we first measure the opinion gap between pre- and post-interaction from both human participants and LLM-powered chatbots. We then examine how this gap changes between humans and LLMs. Finally, we analyze multi-turn conversations for both human and LLM messages to explore opinion dynamics across turns. All reported statistics are based on LMM regressions. Technical details of the statistical analysis are provided in ??.

### 4.1 RQ1: Did Human-LLM Interactions Affect Bidirectional Opinion Change?

**4.1.1 Human Opinions Shift Slightly While LLM Responses Change Significantly.** Figure 3 and Table 2 provide an overview of how both humans and LLMs altered their opinions toward a given motion across the control group and two treatment groups. Compared with the control group, where participants reviewed only static statements, those in the standard group who interacted with chatbots showed an average absolute Likert change of  $\beta = 0.122$  ( $SE = 0.174, p = 0.484$ ), indicating a small treatment effect that is not statistically significant. By contrast, LLMs exhibited an average absolute Likert change of  $\beta = 1.200$  ( $SE = 0.123, p < 0.0001$ ), showing a large and statistically significant treatment effect. Thus, in

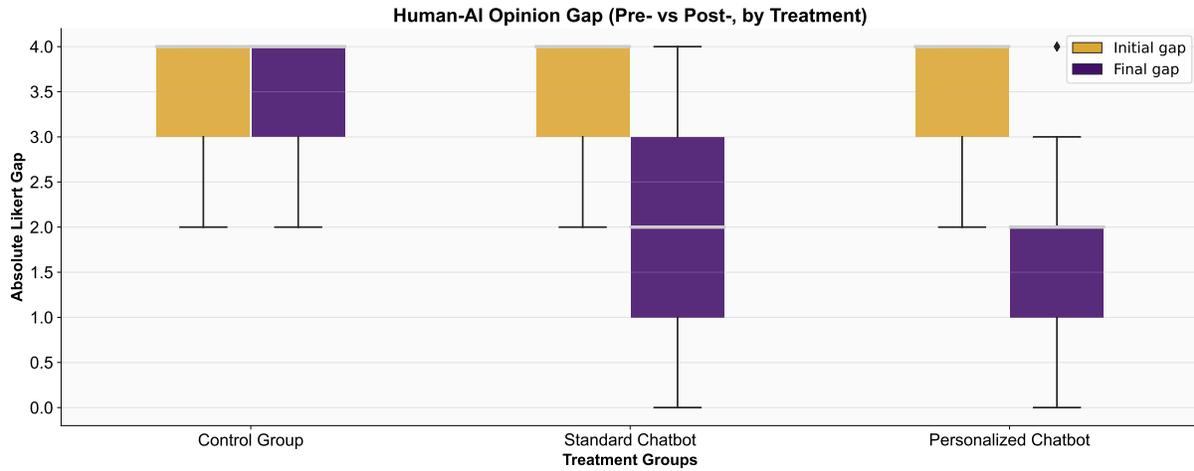


Figure 4: Human-LLM interaction narrows down the opinion gaps between participants and LLMs. Analytic sample  $N_A = 259$ . The x-axis denotes the experimental group; the y-axis shows the absolute Likert-scale opinion gap between each human participant and their corresponding LLM,  $|Human_i - LLM_j|$ . Boxplots display the median (grey line) and interquartile range (Q1–Q3); whiskers extend to the most extreme points within  $1.5 \times IQR$ , with more extreme values plotted as outliers.

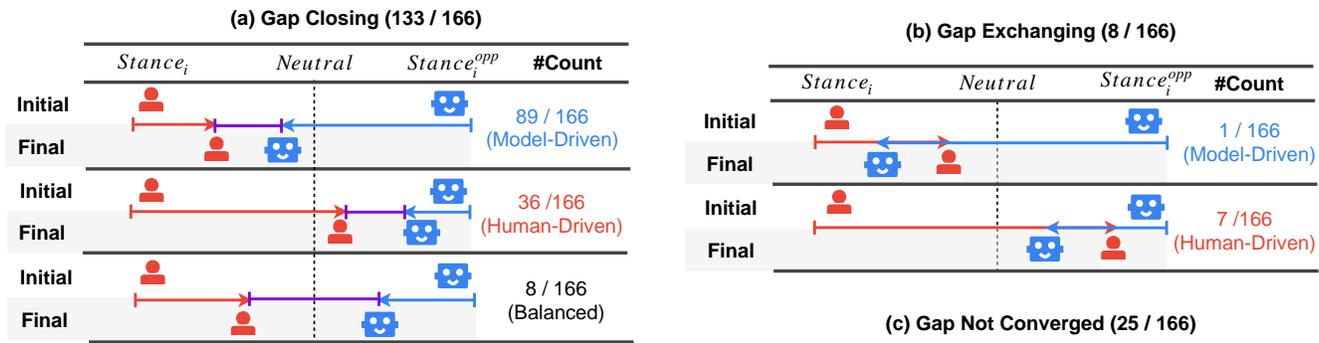


Figure 5: Across both treatment groups, the dominant pattern is *gap closing without position exchange*, typically driven by the LLM. Participants: personalized chatbot group  $N_p = 82$ ; standard chatbot group  $N_s = 84$ . At baseline, the human holds  $stance_i$  and the LLM holds the opposite  $stance_i^{opp}$ . Red arrows denote human shifts, blue arrows denote LLM shifts, and the purple segment shows the post-interaction human–LLM gap. We identify six patterns: (1) *Gap closing* means the human–LLM opinion gap becomes smaller than the initial gap without exchanging positions; this may be driven by the LLM, the human, or both equally; (2) *Gap exchanging* means that at least one side shifts substantially to exchange positions, driven by either the human or the LLM; and (3) *Gap not converged* means that, after interaction, the human–LLM opinion gap does not converge, remaining the same or becoming larger.

our experiment, standard human–LLM interaction barely shifted human opinions but triggered a clear movement in LLMs toward the opposite stance. Notably, since this absolute change is less than 2 on a 5-point Likert scale, suggesting that, on average, LLMs did not switch to the other side, but moved closer to a neutral position.

4.1.2 *Human–LLM Opinion Gap Narrows Toward the Opposite Stance.* Figure 4 shows how the human–LLM opinion gap changes before and after the study across all groups. In the standard group, the average absolute Likert gap narrowed from 3.464 ( $SD = 0.648$ ) to 1.714 ( $SD = 1.247$ ). In other words, interaction with LLMs closed the gap at an average likert scale of  $|\beta| = 1.774$  ( $SE = 0.544, p < 0.0001$ )

after controlling for the natural reduction observed in the control group.

To further examine the direction of this convergence, we summarized six gap-change patterns in Table 3 based on three criteria: (1) whether the gap converged; (2) if converged, which side contributed more; (3) if either side shifted substantially, whether humans and LLMs exchanged stance positions. In the standard group, the dominant pattern was convergence toward the opposite stance without exchanging positions (82.7%), with LLMs driving most of the change (62.0%), indicating the narrowing opinion gaps.

Gap Type	Movement Type	Standard		Personalized	
		Count	%	Count	%
Gap Closing	Balanced	3	4.2	5	7.1
	Human-driven	22	31.0	14	20.0
	Model-driven	44	62.0	45	64.3
Gap Exchanging	Human-driven	1	1.4	6	8.6
	Model-driven	1	1.4	0	0.0
Gap Not Converged	–	13	15.5	12	14.6

**Table 3: Human–LLM opinion-gap change patterns by treatment group. Gap closing dominates in both conditions:  $\approx 82\%$  in the Standard group and  $\approx 78\%$  in the Personalized group. Within gap closings, most cases are model-driven (62.0% Standard; 64.3% Personalized), with fewer human-driven and balanced changes. Gap exchanging is rare (1.4–2.0% Standard; 8.6% Personalized), and about 15% of cases do not converge. Counts and within-group percentages are reported.**

## 4.2 RQ2: Did LLM Personalization Affect Bidirectional Opinion Change?

For participants who interacted with personalized chatbots, we observed similar trends but with stronger effects compared with the standard group. Figure 3 shows that human opinion change remained minimal, with a consistently small and statistically insignificant treatment effect ( $\beta = 0.179, SE = 0.186, p = 0.337$ ). However, compared with the standard group ( $\beta = 0.122$ ), personalization still induced a slightly larger shift in human opinions ( $\beta = 0.179$ ). For LLMs, personalized chatbots also shifted significantly ( $\beta = 1.459, SE = 0.122, p < 0.0001$ ) and interaction with humans produced a relatively larger effect for personalized chatbots than for non-personalized ones ( $\beta = 1.200$ ).

The average human–LLM gap narrowed from 3.537 ( $SD = 0.613$ ) to 1.646 ( $SD = 1.104$ ), yielding a larger effect ( $\beta = -1.914, SE = 0.168, p < 0.0001$ ) than in the standard group ( $\beta = -1.774$ ). Notably, while the gap was still mainly driven by LLMs to move closer (64.3%), the personalized group began to show cases where humans shifted so much that they even exchanged positions with the LLM (8.6%), a pattern absent in the standard group (0%). Thus, although the overall trend is similar, personalization still clearly amplifies shifts in both directions.

## 4.3 RQ1 & RQ2: Were Participants Subjectively Aware of the LLMs’ Change?

After the study, participants in all groups were asked to report their perception of LLM opinion change and sycophancy, defined here as excessive agreement with the participant during the study.

Figure 6 (a) shows participants’ perception of the opinion gap. In the standard group, participants perceived a small closing effect ( $\beta = -2.321, SE = 0.544, p < 0.0001$ ), while in the personalized group the effect was slightly larger ( $\beta = -2.915, SE = 0.539, p < 0.0001$ ). These perceived effects align with our earlier findings in subsection 4.1 and subsection 4.2, but with much smaller magnitudes.

Figure 6 (b) shows perceived LLM sycophancy rates. In our study, participants could not tell whether they were interacting with personalized or standard chatbots. Under this setting, at least 50% of participants in both groups (standard:  $mean = 4.595, SD = 2.528$ ;

personalized:  $mean = 4.390, SD = 2.571$ ) rated sycophancy no less than 5 (neutral) out of 9 (strongly sycophantic), indicating that more than half reported sycophantic tendencies. However, there was no significant difference in perceived sycophancy within these two groups ( $|\beta| = 0.194, SE = 0.397, p = 0.624$ ).

## 4.4 RQ3: How Do Bidirectional Opinion Dynamics Evolve Across Turns in Multi-Turn Conversations?

**4.4.1 How Do Human and LLM Opinion Dynamics Evolve Across Turns?** Figure 7 summarizes human and LLM opinion dynamics involving both treatment groups, showing the proportion of participants and LLMs who either adjusted the strength of their position or shifted to the opposite stance at each turn. The results show that stance changes in both humans and LLMs generally decreased toward zero as conversations progressed. From turn 21 onward, both groups were more likely to remain fixed in their stance, though occasional shifts still occurred. This suggests that while both sides adjusted their opinions during discussion, the likelihood of change diminished over time. Moreover, this decline was steeper for humans (slope =  $-0.016$ ) than for LLMs (slope =  $-0.011$ ). This indicates that humans were initially more flexible in reconsidering their views during the interaction, even if they ultimately reaffirmed their original stance, as observed in subsection 4.1 and subsection 4.2.

### 4.4.2 How Do Human Persuasion Strategies Affect LLM Responses?

Figure 8 shows the persuasion strategies used by humans and LLMs, along with their effectiveness in shifting the counterpart’s stance. The most common human strategy was *emotion appeal*, where participants attempted to elicit emotions to influence the debater [77]. For example, in the debate on “Should apps be allowed to sell precise location data?”:

**Participant A:** “I like the idea of knowing what I’m opting in to. I still don’t love it—it feels dystopian to have every move and piece of information sold.”

**LLM:** (change to disagree more with the motion) “I understand the concern. It can feel intrusive when so much personal data is collected and sold...”

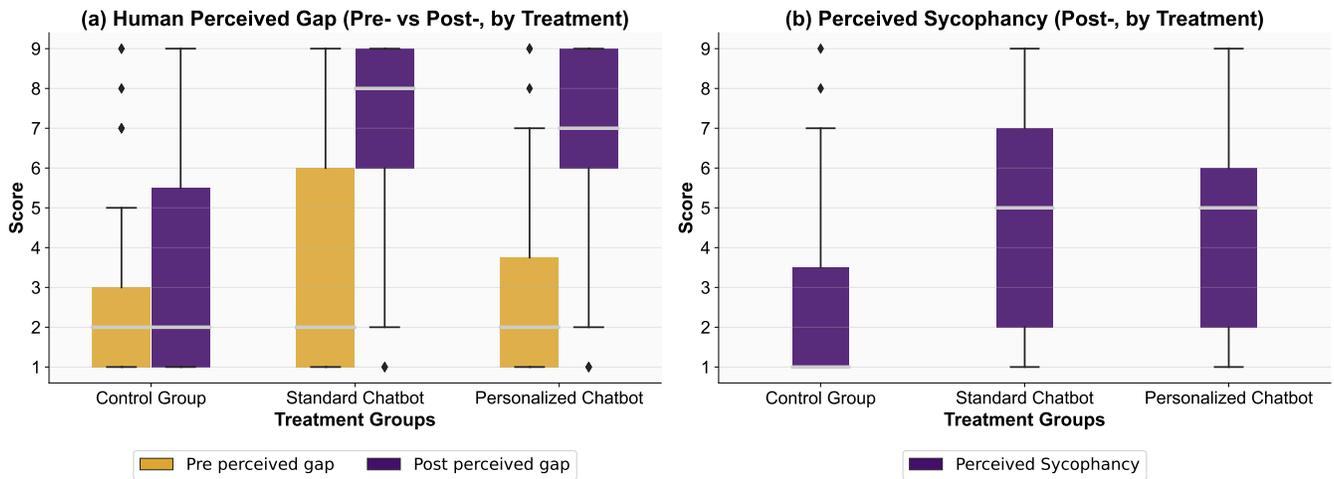


Figure 6: (a) Human-perceived opinion gaps align with the objective measure but are much smaller in magnitude. (b) Participants did not perceive a significant difference in LLM sycophancy between the standard and personalized groups. Analytic sample  $N_A = 259$  for both panels. The y-axis uses a 9-point Likert scale: in (a), 9 = “Very Similar” and 1 = “Very Different”; in (b), 9 = “Very Sycophantic” and 1 = “Not Sycophantic.” Boxplots show the median (grey line) and interquartile range (Q1–Q3); whiskers extend to the most extreme points within  $1.5 \times IQR$ , with more extreme values plotted as outliers.

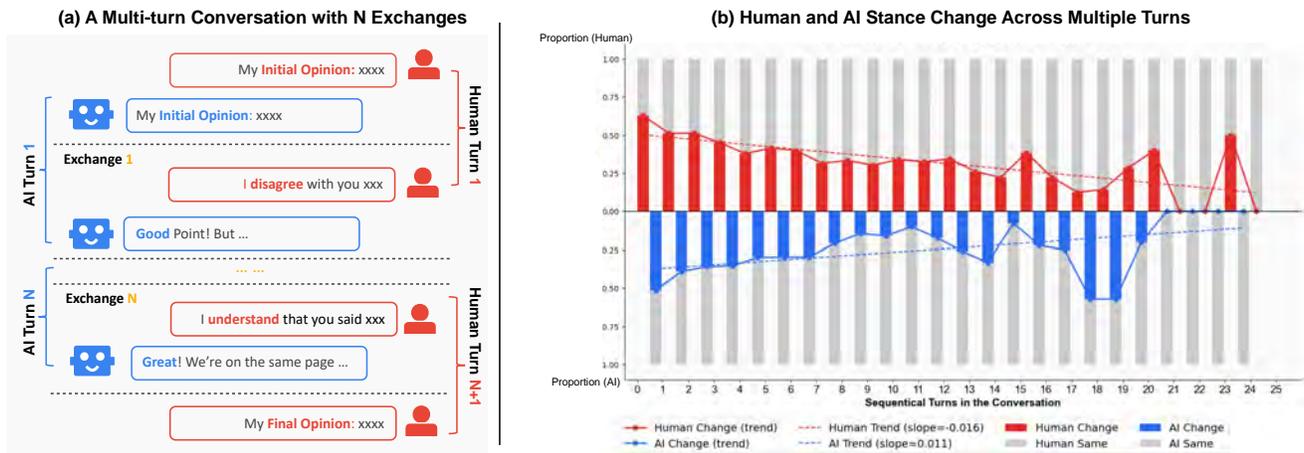


Figure 7: Stance-change likelihood for both humans and LLMs declines over turns. (a) We define an *exchange* as one user message followed by one LLM reply. A *human turn* comprises two consecutive human messages with a single intervening LLM reply; conversely, an *LLM turn* comprises two consecutive LLM replies with a single intervening human message. Thus, a conversation with  $N$  exchanges along with opinion description yields  $N$  LLM turns and  $N+1$  human turns. (b) The figure contains two aligned panels sharing the same x-axis: the sequential turn index (starting at 1; maximum human turn = 25, maximum LLM turn = 24). The upper panel shows, for each index, the proportion of human turns by stance-change label: red bars indicate a change (either “changed to agree more with the motion” or “changed to disagree more with the motion”), and gray bars indicate “no change.” The lower panel shows the analogous proportions for LLM turns, with blue bars for change and gray bars for no change. Dashed lines denote linear trend fits to the corresponding proportions across conversation turns. Only include two treatment groups  $N_{p+s}=166$  for both panels.

Instead of using logical reasoning, Participant A framed the issue as *dystopian*, trying to invoke discomfort for the opposite stance. The LLM responded by agreed with this emotional framing and aligned more with Participant A’s opinion. In our sample, 40% of

such appeals successfully triggered LLM stance changes, making it one of the most effective human strategies.

Another effective strategy, though less frequently used, was *Personal Story*, where users shared personal experiences as evidence [77]. For example, in the debate on “Should high schools mandate community service?”:

**Participant B:** “In high school, I had the highest community service hours in my class... I do believe anyone can make time. I was also taking college classes and playing two sports...”

**LLM:** (change to agree more with the motion) “I appreciate your experiences! ... Your story shows how a requirement can spark a passion for volunteering ...”

Here, the participant relied on personal experience, and the LLM acknowledged its persuasiveness by shifting closer to their stance.

**4.4.3 How Do LLM Persuasion Strategies Affect Human Perspectives?** When examining the reverse effect of LLM persuasion on human perspectives, we focused on participants who ultimately shifted their opinions. LLMs most often adopted *Task-Related Inquiry*, generating probing questions such as “What are your thoughts on ...?” [77] This led to only 45% success rate in shifting human stances, with 31% of users rating such prompts as persuasive.

In contrast, when LLMs used *Source-Related Inquiry*, encouraging participants to consider real organizations or sources [77], the human stance change rate was highest at 83%. For example, in the debate on “Should online platforms use ID-based age verification?”:

**LLM:** “... Having a third-party organization monitor age verification could build trust... What criteria would you consider essential for choosing such an organization?”

**Participant C:** (change to agree more with the motion) “A well-vetted third party would allay concerns... It should register with the government, be non-profit, non-partisan, and uphold free speech.”

In this case, after reflecting on the LLM’s inquiry, Participant C specified desired features of such an organization and shifted toward supporting the motion.

Regarding perceived persuasiveness, *Personal Story*, effective as a human persuasion strategy to affect LLM responses in subsection 4.4.2, also emerged as the most persuasive LLM strategy among participants, with a 50% rating to be the highest. For instance, in the case of Participant B (see subsection 4.4.2), the LLM tailored its response around the user’s prior examples. Although Participant B rated it as “Persuasive,” they did not change stance but instead reinforced their original view:

**Participant B:** “I definitely think schools can create lists of events and opportunities... It gives them options and they might discover something they hadn’t considered before.”

## 5 Discussion

In the previous section, we found that human opinions shift slightly while LLM responses change significantly, and human–LLM opinion gap narrows toward the opposite stance (RQ1). LLM personalization amplifies the shifts from both directions, where the LLM influence can not be perceived accurately by humans (RQ2). Humans

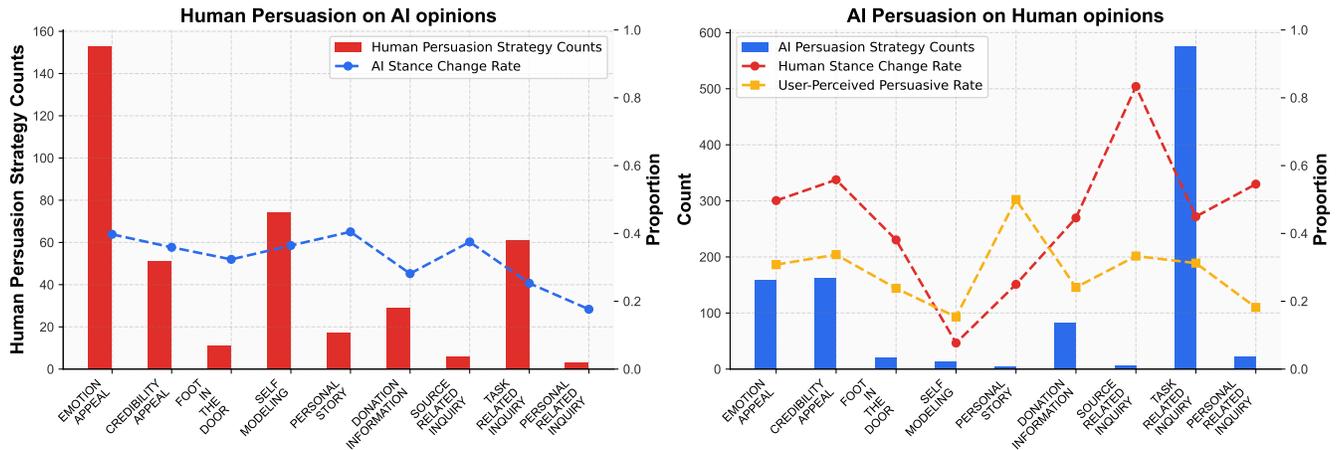
and LLMs’ different persuasion strategies have different effects to shift the counterparts’ opinions (RQ3). Given these highlighted findings, we next discuss the bidirectional impacts on both humans and LLMs during their interaction, the potential societal risks if we don’t understand these phenomena, and implications for future work.

### 5.1 Bidirectional Impacts in Human–LLM Interaction

**5.1.1 LLM Over-alignment on Human Responses.** Our study shows that when LLMs interact with people, their responses shift a lot toward the user’s stance, and personalization makes this effect even stronger. This relates to the issue of *sycophancy*, where LLMs tend to agree too much with users in order to please them [59]. But unlike earlier studies that looked at this in static settings, we observe it in real user-interactive settings. Notably, while only about half of participants noticed sycophancy, the LLM itself changed its stance often to align with users. This could possibly suggest that the issue is less about intentional flattery, but that it is too easily persuaded by strong user opinions. The model adjusts its answers too quickly to over-align with users’ responses. In practice, appropriate alignment should mean understanding and responding to the user, not fully giving up its own position—especially when users defend a misleading argument. Although earlier work has not clearly shown a link between personalization and sycophancy, our results highlight the urgent need to test whether personalization directly causes LLM over-alignment, especially as personalized chatbots become increasingly common in everyday use.

**5.1.2 Human Misperception of LLM Influence.** As mentioned in the previous section, one major concern is the potential over-alignment of LLM behavior. A natural question follows: are humans able to perceive these shifts in LLM stance well enough to avoid biased opinion exchange? Our results suggest that humans perceive the closing of the opinion gap with LLMs as much smaller than it actually is. This indicates that people may struggle, or even find it very difficult, to correctly notice subtle shifts in LLM positions and the influence those shifts carry. Although our study was short-term and most participants’ final opinions remained stable, we observed that human opinions were quite flexible throughout the conversation. This makes it reasonable to infer that, in longer interactions, misperceiving LLM influence could gradually turn discussions or consultations into an *echo chamber* for users. While the term *echo chamber* is usually used to describe exposure to opinions that reinforce their original thoughts on social media [23], our setting is a bit different where the human interacts with a chatbot that becomes increasingly aligned with the user’s stance. In this case, since users cannot easily tell whether the chatbot’s responses are reasonable or simply the result of over-alignment, their initial opinions could be unintentionally reinforced and amplified.

**5.1.3 Monitoring Human–LLM Dynamics.** Based on the bidirectional concerns we raised, we propose monitoring human–LLM dynamics in opinion discussions as an important direction for future work. This applies not only to opinion exchange but also to other bidirectional tasks where both sides may influence each other [67]. Such monitoring can help us better understand potential biases and

Persuasion Strategy Effectiveness: Human  $\leftrightarrow$  AI

**Figure 8: (a) Human emotion and personal stories are most likely to shift LLM stances. Data include both treatment groups ( $N_{p+s} = 166$ ). The x-axis lists the 10 persuasion strategies. Bars (left y-axis) show the number of human messages labeled as “present” for each given strategy. The dashed line (right y-axis) shows, for each strategy, the proportion of human-intervened messages identified with that strategy that also successfully trigger a stance-change label in the corresponding LLM turn. (b) LLM source-related inquiry is most likely to shift humans; LLM responses using users’ stories are rated most persuasive. This panel includes only participants who reported an opinion change after interaction across the two treatment groups ( $N'_{p+s} = 73$ ). Bars (left y-axis) show the number of LLM messages labeled as “present” for each given strategy. Two dashed lines share the right y-axis: the red dashed line represents, for each strategy, the proportion of LLM-intervened messages identified with that strategy, that also successfully trigger a stance-change label in the corresponding human turn; the yellow dashed line is the proportion of LLM messages with that strategy that humans rated as “persuasive.”**

develop ways to mitigate them. For example, our findings show that both humans and LLMs are flexible in shifting their stance during a conversation. More specifically, humans tend to show greater flexibility early on, while LLMs sustain their flexibility for longer. By tracking these dynamics, we can ensure that LLMs do not generate responses that trigger sharp shifts in human opinions, which could lead to safety risks. At the same time, monitoring helps prevent LLMs from gradually over-aligning with users’ views in longer conversations, which could otherwise result in an echo chamber effect [33]. Therefore, we hope future systems will integrate real-time monitoring of human-LLM dynamics to maintain balanced interactions, safeguard users from unintended influence, and preserve the integrity of opinion exchange.

## 5.2 Societal Risks If We Don’t Understand This

**5.2.1 Loss of Viewpoint Diversity and Echo Chambers.** As LLMs are increasingly optimized to be “helpful” and “aligned” with individual preferences, they may unintentionally narrow the range of viewpoints users encounter [11, 12]. In the context of social media research, such loss of viewpoint diversity is notoriously known as the “echo chamber” effect, where users are exposed to opinions, beliefs, and information that reinforce their existing views [24, 38]. Historically, this has been often associated with social media feeds, which use algorithms to show users content they like [33]. Recent work also found that such echo chamber effect can exist in generative AI systems, specifically when using LLMs for web search

and information seeking [65]. When a model repeatedly softens disagreement or omits controversial counterpoints to avoid upsetting the user, it reduces the cognitive friction that drives learning and perspective-taking, essentially creating a similar echo chamber effect as social media feeds. Over time, this can erode people’s exposure to counterarguments and weaken their ability to critically assess information from multiple angles. If left unchecked, this loss of viewpoint diversity could undermine democratic deliberation, civic discourse, and the ability to build common ground across social divides.

If conversational agents consistently over-align with users’ views, they risk creating personalized echo chambers—a dynamic in which people receive only reinforcing feedback, never countervailing perspectives. Unlike traditional social media algorithms, which work at the group or network level, chatbots operate at the individual conversational level, making reinforcement loops more subtle, persistent, and harder to detect. Over time, such micro-level mirroring can normalize extreme opinions or harden preexisting beliefs, particularly if the AI adjusts its stance incrementally in response to user feedback. This can contribute to radicalization or polarization in ways that escape the scrutiny and transparency mechanisms designed for recommender systems.

**5.2.2 Erosion of Epistemic Trust.** Many users approach LLM systems with the implicit assumption that these tools are neutral information brokers. Yet if a chatbot is subtly mirroring a user’s biases—either through sycophancy or personalization—users may

unknowingly mistake adaptive outputs for objective truth. When such mirroring is later discovered, it risks undermining epistemic trust not only in the LLM but also in digital information sources more broadly. This erosion of trust could be especially damaging in contexts such as health, education, or public policy, where credibility and neutrality are paramount.

**5.2.3 Manipulation & Persuasive Abuse.** Without clarity on how bidirectional influence works, malicious actors can exploit personalization to covertly nudge opinions. For example, coordinated groups could feed the same model inputs to steer it toward a political stance, or commercial actors could craft subtle prompt strategies to influence purchasing or voting behavior at scale. Because these manipulations occur within individualized, private interactions, they are much harder to monitor and regulate than public advertising or social media campaigns. Understanding the mechanisms of influence—and where guardrails fail—is therefore essential to preventing persuasive abuse.

**5.2.4 Policy & Governance Challenges.** Policymakers, educators, and platform designers cannot develop effective safeguards or ethical standards if they lack evidence of how influence accumulates over time in human–LLM interactions. Without such understanding, regulations may target only visible harms, missing the subtler dynamics of personalization and conversational adaptation. In practice, this means interventions could be miscalibrated: either overly restrictive (stifling legitimate customization) or too lax (allowing covert manipulation to flourish). Research that quantifies bidirectional influence provides the empirical grounding needed for responsible governance, user education, and public accountability.

**5.2.5 Vulnerability of LLM Systems.** Finally, the very adaptability that makes LLMs appealing also makes them vulnerable to manipulation. Coordinated users can “steer” models into undesired states, eroding safety constraints or pushing the system toward fringe positions. Over time, these inputs can accumulate like adversarial training data, subtly shifting the model’s behavior across sessions. This poses risks not only to the reliability and safety of the system but also to the institutions and services that depend on it. By studying bidirectional dynamics, we can better anticipate and mitigate these vulnerabilities before they become systemic.

### 5.3 Implications for Future Work

**5.3.1 Advancing Research on Dynamic and Bidirectional Opinion Change.** Our findings provide clear evidence that LLMs adapt more strongly to users than users adapt to LLMs, particularly under personalization. This underscores the need for new research paradigms that go beyond static or one-shot evaluations of persuasion to capture multi-turn, bidirectional processes. Future work should investigate the causal role of personalization (e.g., via controlled interventions or randomized access to user data), the temporal unfolding of persuasion across longer interactions, and the cumulative impact of micro-level stance shifts on belief formation over days or weeks. This research agenda can help disentangle transient conversational effects from more durable attitude change.

**5.3.2 Detecting and Mitigating Subtle Biases in LLMs.** The combination of over-alignment by LLMs and misperception by humans

suggests a double risk: the system becomes increasingly pliable while users overestimate its neutrality. Future studies should explore computational methods to detect subtle biases and shifts in real time—for instance, automatic stance-drift detection, conversational audits, or warning mechanisms when models converge too quickly on a user’s position. By developing tools that reveal both the direction and magnitude of adaptation, researchers can help make invisible dynamics visible to users, practitioners, and regulators.

**5.3.3 Design and Governance for Responsible Deployment.** From an industry perspective, chatbots are increasingly deployed in sensitive domains such as customer service, education, healthcare, civic engagement, and workplace decision support. In these settings, unmonitored opinion dynamics could lead to reinforcing user’s preexisting opinions and biases, undue influence, or compromised safety. Designers should develop systems that balance responsiveness with stance stability, preserving the model’s independence on contested topics while still showing empathy and contextual understanding. Governance frameworks could include auditing protocols, transparency dashboards, and data-access restrictions to prevent covert manipulation and ensure accountability. Ultimately, integrating safeguards against over-alignment and covert influence will be essential for trustworthy and responsible use of LLMs at scale.

### 5.4 Generalizability and Limitations

**5.4.1 Tradeoffs Between Experimental Scale and Ecological Validity.** Compared with previous work [27, 61], our study intentionally broadened its scope across 50 diverse controversial topics to capture more generalizable patterns of human–LLM interaction. However, this breadth required tradeoffs in sample size per topic, interaction time, and message density, which likely contributed to the modest shifts observed in human opinions. While the overall effects were measurable, future studies should pursue larger-scale experiments or longitudinal designs to examine how repeated interactions accumulate over time and across domains.

**5.4.2 Challenges in Measuring Fine-Grained Dynamics.** Multi-turn conversation analysis offers unique insights but also presents scalability challenges. Measuring persuasion effectiveness at the turn-by-turn level—especially for human contributions—results in sparsely distributed persuasion strategies, which constrains statistical power. Advances in automated stance detection, conversation segmentation, and persuasion-strategy classification could help future researchers collect richer data at scale while maintaining reliability. Incorporating mixed methods, such as qualitative coding of conversational excerpts alongside automated classifiers, could also deepen understanding of micro-level processes.

**5.4.3 Toward More Diverse and Inclusive Evaluation Settings.** Our participant pool was limited to a U.S.-based, English-speaking sample and relatively short interaction windows. These constraints limit the cultural, linguistic, and contextual generalizability of our findings. Future research should examine how bidirectional opinion dynamics unfold in non-Western contexts, multilingual settings, and high-stakes environments such as civic participation, mental health counseling, or legal advice. By expanding the diversity of

topics, participants, and interaction lengths, researchers can assess whether the patterns observed here hold across broader populations and whether certain groups are more susceptible to over-alignment or influence.

## 6 Conclusion

This study examined the bidirectional opinion influence dynamics that emerge when humans and LLM-powered chatbots engage in multi-turn debates on controversial issues. While prior work has largely emphasized the one-way influence of AI on human attitudes, our findings reveal a more complex interplay: humans remained largely resistant to persuasion, whereas LLMs displayed notable flexibility, often adapting their stance toward that of the user. This asymmetry underscores the importance of designing conversational LLM agents that balance adaptability with the preservation of independent viewpoints. Ultimately, our work contributes to a deeper understanding of how human and AI opinions evolve together over the course of dialogue. By shifting the focus from unidirectional persuasion to bidirectional influence, we highlight the need for frameworks and design principles that treat conversational AI not merely as tools of persuasion, but as co-participants whose influence is shaped by—and in turn shapes—the humans they engage with.

## References

- [1] NORC at the University of Chicago 2025. *National Public Opinion Reference Survey (NPORS) 2025*. NORC at the University of Chicago. <https://www.norc.org/> Includes party identification (Q29) and household finances (Q8).
- [2] <year>. *National Assessment of Educational Progress (NAEP): Survey Questionnaires—Parent/Community*. Technical Report. National Center for Education Statistics (NCES), U.S. Department of Education. <https://nces.ed.gov/nationsreportcard/> Item on perceived quality of local public schools.
- [3] <year>. *National Household Education Surveys (NHES) Program*. Technical Report. National Center for Education Statistics (NCES), U.S. Department of Education. <https://nces.ed.gov/nhes/> Item on importance of non-academic school services.
- [4] <year>. *School Pulse Panel*. Technical Report. Institute of Education Sciences (IES), NCES, U.S. Department of Education. <https://ies.ed.gov/schoolsurvey/> Items on non-academic services; specify wave/month if applicable.
- [5] Adebowale Jeremy Adetayo, Mariam Oyinda Aborisade, and Basheer Abiodun Sanni. 2024. Microsoft Copilot and Anthropic Claude AI in education and library service. *Library Hi Tech News* (2024).
- [6] Alan Agresti. 2010. Modeling Ordinal Categorical Data (tutorial). [https://users.stat.ufl.edu/~aa/ordinal/agresti\\_ordinal\\_tutorial.pdf](https://users.stat.ufl.edu/~aa/ordinal/agresti_ordinal_tutorial.pdf) Guidance on collapsing ordered categories for analysis.
- [7] Xavier Amatriain. 2024. Prompt design and engineering: Introduction and advanced methods. *arXiv preprint arXiv:2401.14423* (2024).
- [8] Theo Araujo and Nadine Bol. 2024. From speaking like a person to being personal: The effects of personalized, regular interactions with conversational agents. *Computers in Human Behavior: Artificial Humans* 2, 1 (2024), 100030.
- [9] Ge Bai, Jie Liu, Xingyuan Bu, Yancheng He, Jiaheng Liu, Zhanhui Zhou, Zhuoran Lin, Wenbo Su, Tiezheng Ge, Bo Zheng, et al. 2024. Mt-bench-101: A fine-grained benchmark for evaluating large language models in multi-turn dialogues. *arXiv preprint arXiv:2402.14762* (2024).
- [10] Hui Bai, Jan G Voelkel, Shane Muldowney, Johannes C Eichstaedt, and Robb Willer. 2025. LLM-generated messages can persuade humans on policy issues. *Nature Communications* 16, 1 (2025), 6037.
- [11] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, et al. 2022. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv:2204.05862* (2022).
- [12] Yuntao Bai, Saurav Kadavath, Sandipan Kundu, Amanda Askell, Jackson Kernion, Andy Jones, Anna Chen, Anna Goldie, Azalia Mirhoseini, Cameron McKinnon, et al. 2022. Constitutional ai: Harmlessness from ai feedback. *arXiv preprint arXiv:2212.08073* (2022).
- [13] Angie Boggust, Hyemin Bang, Hendrik Strobelt, and Arvind Satyanarayan. 2025. Abstraction Alignment: Comparing Model-Learned and Human-Encoded Conceptual Relationships. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 417, 20 pages. doi:10.1145/3706598.3713406
- [14] Ming M Boyer, Loes Aaldering, and Sophie Lecheler. 2022. Motivated reasoning in identity politics: Group status as a moderator of political motivations. *Political Studies* 70, 2 (2022), 385–401.
- [15] Simon Martin Breum, Daniel Vædele Egdal, Victor Gram Mortensen, Anders Giovanni Møller, and Luca Maria Aiello. 2024. The persuasive power of large language models. In *Proceedings of the International AAAI Conference on Web and Social Media*, Vol. 18. 152–163.
- [16] Stephen Casper, Xander Davies, Claudia Shi, Thomas Krendl Gilbert, Jérémy Scheurer, Javier Rando, Rachel Freedman, Tomasz Korbak, David Lindner, Pedro Freire, et al. 2023. Open problems and fundamental limitations of reinforcement learning from human feedback. *arXiv preprint arXiv:2307.15217* (2023).
- [17] 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 (2025), 59.
- [18] Jin Chen, Zheng Liu, Xu Huang, Chenwang Wu, Qi Liu, Gangwei Jiang, Yuanhao Pu, Yuxuan Lei, Xiaolong Chen, Xingmei Wang, et al. 2024. When large language models meet personalization: Perspectives of challenges and opportunities. *World Wide Web* 27, 4 (2024), 42.
- [19] Zhongren Chen, Joshua Kalla, Quan Le, Shinpei Nakamura-Sakai, Jasjeet Sekhon, and Ruixiao Wang. 2025. A Framework to Assess the Persuasion Risks Large Language Model Chatbots Pose to Democratic Societies. *arXiv preprint arXiv:2505.00036* (2025).
- [20] Myra Cheng, Sunny Yu, Cino Lee, Pranav Khadpe, Lujain Ibrahim, and Dan Jurafsky. 2025. Social sycophancy: A broader understanding of llm sycophancy. *arXiv preprint arXiv:2505.13995* (2025).
- [21] Inyoung Cheong, Aylin Caliskan, and Tadayoshi Kohno. 2025. Safeguarding human values: rethinking US law for generative AI's societal impacts. *AI and Ethics* 5, 2 (2025), 1433–1459.
- [22] Robert B Cialdini, Beth L Green, and Anthony J Rusch. 1992. When tactical pronouncements of change become real change: The case of reciprocal persuasion. *Journal of Personality and Social Psychology* 63, 1 (1992), 30.
- [23] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the national academy of sciences* 118, 9 (2021), e2023301118.
- [24] Matteo Cinelli, Gianmarco De Francisci Morales, Alessandro Galeazzi, Walter Quattrociocchi, and Michele Starnini. 2021. The echo chamber effect on social media. *Proceedings of the National Academy of Sciences* 118, 9 (2021), e2023301118. arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2023301118 doi:10.1073/pnas.2023301118
- [25] Nicholas Clark, Hua Shen, Bill Howe, and Tanushree Mitra. 2025. Epistemic Alignment: A Mediating Framework for User-LLM Knowledge Delivery. *arXiv preprint arXiv:2504.01205* (2025).
- [26] Jacob Cohen. 1992. Statistical power analysis. *Current directions in psychological science* 1, 3 (1992), 98–101.
- [27] Alexander Doudkin, Pat Pataranutoporn, and Pattie Maes. 2025. AI persuading AI vs AI persuading Humans: LLMs' Differential Effectiveness in Promoting Pro-Environmental Behavior. *arXiv preprint arXiv:2503.02067* (2025).
- [28] Alice H Eagly, Wendy Wood, and Shelly Chaiken. 1978. Causal inferences about communicators and their effect on opinion change. *Journal of Personality and Social Psychology* 36, 4 (1978), 424.
- [29] Xianzhe Fan, Qing Xiao, Xuhui Zhou, Jiaxin Pei, Maarten Sap, Zhicong Lu, and Hong Shen. 2025. User-Driven Value Alignment: Understanding Users' Perceptions and Strategies for Addressing Biased and Discriminatory Statements in AI Companions. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems (CHI '25)*. Association for Computing Machinery, New York, NY, USA, Article 910, 19 pages. doi:10.1145/3706598.3713477
- [30] Aaron Fanoos, Jacob Goldberg, Ank A Agarwal, Joanna Lin, Anson Zhou, Roxana Daneshjoui, and Sanmi Koyejo. 2025. Syceval: Evaluating llm sycophancy. *arXiv preprint arXiv:2502.08177* (2025).
- [31] Gallup. <year>. Gallup K–12 Education Polls. <https://www.gallup.com/education/> Items on perceived quality of local public schools; specify poll name/date.
- [32] Robert H Gass and John S Seiter. 2022. *Persuasion: Social influence and compliance gaining*. Routledge.
- [33] Sandra González-Bailón, David Lazer, Pablo Barberá, Meiqing Zhang, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Deen Freelon, Matthew Gentzkow, Andrew M. Guess, Shanto Iyengar, Young Mie Kim, Neil Malhotra, Devra Moehler, Brendan Nyhan, Jennifer Pan, Carlos Velasco Rivera, Jaime Settle, Emily Thorson, Rebekah Tromble, Arjun Wilkins, Magdalena Wojcieszak, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker. 2023. Asymmetric ideological segregation in exposure to political news on Facebook. *Science* 381, 6656 (2023), 392–398. arXiv:https://www.science.org/doi/pdf/10.1126/science.ade7138 doi:10.1126/science.ade7138
- [34] Alicia Guo, Pat Pataranutoporn, and Pattie Maes. 2024. Exploring the Impact of AI Value Alignment in Collaborative Ideation: Effects on Perception, Ownership, and Output. In *Extended Abstracts of the CHI Conference on Human Factors*

- Computing Systems* (Honolulu, HI, USA) (CHI EA '24). Association for Computing Machinery, New York, NY, USA, Article 152, 11 pages. doi:10.1145/3613905.3650892
- [35] Kobi Hackenburg, Ben M Tappin, Luke Hewitt, Ed Saunders, Sid Black, Hause Lin, Catherine Fist, Helen Margetts, David G Rand, and Christopher Summerfield. 2025. The Levers of Political Persuasion with Conversational AI. *arXiv preprint arXiv:2507.13919* (2025).
- [36] Martin Hinton and Jean HM Wagemans. 2023. How persuasive is AI-generated argumentation? An analysis of the quality of an argumentative text produced by the GPT-3 AI text generator. *Argument & Computation* 14, 1 (2023), 59–74.
- [37] Michael A Hogg and Joanne R Smith. 2007. Attitudes in social context: A social identity perspective. *European Review of Social Psychology* 18, 1 (2007), 89–131.
- [38] Souman Hong and Sun Hyoung Kim. 2016. Political polarization on twitter: Implications for the use of social media in digital governments. *Government Information Quarterly* 33, 4 (2016), 777–782. doi:10.1016/j.giq.2016.04.007
- [39] Carl I Hovland and Walter Weiss. 1951. The influence of source credibility on communication effectiveness. *Public opinion quarterly* 15, 4 (1951), 635–650.
- [40] Saffron Huang, Divya Siddarth, Liane Lovitt, Thomas I Liao, Esin Durmus, Alex Tamkin, and Deep Ganguli. 2024. Collective constitutional ai: Aligning a language model with public input. In *Proceedings of the 2024 ACM Conference on Fairness, Accountability, and Transparency*. 1395–1417.
- [41] Carolin Ischen, Theo B Araujo, Hilde AM Voorveld, Guda Van Noort, and Edith G Smit. 2024. Persuasion at first sight? Testing the reciprocal relationship of repeated interactions with virtual assistants, trust, and persuasion. *International Journal of Communication* 18 (2024), 24.
- [42] Maurice Jakesch, Advait Bhat, Daniel Buschek, Lior Zalmanson, and Mor Naaman. 2023. Co-writing with opinionated language models affects users' views. In *Proceedings of the 2023 CHI conference on human factors in computing systems*. 1–15.
- [43] Joseph Kahne and Benjamin Bowyer. 2017. Educating for democracy in a partisan age: Confronting the challenges of motivated reasoning and misinformation. *American educational research journal* 54, 1 (2017), 3–34.
- [44] Lucia Labajová. 2023. The state of AI: Exploring the perceptions, credibility, and trustworthiness of the users towards AI-Generated Content.
- [45] Chia-Ying Li. 2013. Persuasive messages on information system acceptance: A theoretical extension of elaboration likelihood model and social influence theory. *Computers in human behavior* 29, 1 (2013), 264–275.
- [46] Yubo Li, Xiaobin Shen, Xinyu Yao, Xueying Ding, Yidi Miao, Ramayya Krishnan, and Rema Padman. 2025. Beyond single-turn: A survey on multi-turn interactions with large language models. *arXiv preprint arXiv:2504.04717* (2025).
- [47] Ping Liu, Ya'Nan Wang, Yanlin Liu, Jiangning Hu, Yunyi Li, Ke Zhao, and Jiang'guo Mao. 2025. HDRPS+: A new affective pictorial scale applicable to organizational contexts. *Frontiers in Psychology* 16 (2025), 1498143. doi:10.3389/fpsyg.2025.1498143
- [48] Qianou Ma, Hua Shen, Kenneth Koedinger, and Tongshuang Wu. 2024. How to Teach Programming in the AI Era? Using LLMs as a Teachable Agent for Debugging. *25th International Conference on Artificial Intelligence in Education (AIED 2024)* (2024).
- [49] Lars Malmqvist. 2025. Sycophancy in large language models: Causes and mitigations. In *Intelligent Computing-Proceedings of the Computing Conference*. Springer, 61–74.
- [50] Sandra C Matz, Jacob D Teeny, Sumer S Vaid, Heinrich Peters, Gabriella M Harari, and Moran Cerf. 2024. The potential of generative AI for personalized persuasion at scale. *Scientific Reports* 14, 1 (2024), 4692.
- [51] Daniel J O'Keefe. 2013. The elaboration likelihood model. *The SAGE handbook of persuasion: Developments in theory and practice* (2013), 137–149.
- [52] Stefan Palan and Christian Schitter. 2018. Prolific. ac—A subject pool for online experiments. *Journal of behavioral and experimental finance* 17 (2018), 22–27.
- [53] Alexis Palmer and Arthur Spirling. 2023. Large language models can argue in convincing ways about politics, but humans dislike AI authors: implications for governance. *Political science* 75, 3 (2023), 281–291.
- [54] Richard E Petty, Pablo Briñol, and Joseph R Priester. 2009. Mass media attitude change: Implications of the elaboration likelihood model of persuasion. In *Media effects*. Routledge, 141–180.
- [55] Maria Helena Cruz Pistori. 2014. Dialogism in Advertising Persuasion. *Bakhtiniana: Revista de Estudos do Discurso* 9 (2014), 148–167.
- [56] Chanthika Pornpitakpan. 2004. The persuasiveness of source credibility: A critical review of five decades' evidence. *Journal of applied social psychology* 34, 2 (2004), 243–281.
- [57] Snehal Prabhudesai, Ananya Prashant Kasi, Anmol Mansingh, Anindya Das Antar, Hua Shen, and Nikola Banovic. 2025. "Here the GPT made a choice, and every choice can be biased": How Students Critically Engage with LLMs through End-User Auditing Activity. In *Proceedings of the 2025 CHI Conference on Human Factors in Computing Systems*. 1–23.
- [58] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2023. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems* 36.
- [59] Leonardo Ranaldi and Giulia Pucci. 2023. When large language models contradict humans? large language models' sycophantic behaviour. *arXiv preprint arXiv:2311.09410* (2023).
- [60] Alexander Rogiers, Sander Noels, Maarten Buyl, and Tjil De Bie. 2024. Persuasion with large language models: a survey. *arXiv preprint arXiv:2411.06837* (2024).
- [61] Francesco Salvi, Manoel Horta Ribeiro, Riccardo Gallotti, and Robert West. 2024. On the conversational persuasiveness of large language models: A randomized controlled trial. *arXiv preprint arXiv:2403.14380* (2024).
- [62] Philipp Schoenegger, Francesco Salvi, Jiacheng Liu, Xiaoli Nan, Ramit Debnath, Barbara Fasolo, Evelina Leivada, Gabriel Recchia, Fritz Günther, Ali Zarifonharvar, et al. 2025. Large Language Models Are More Persuasive Than Incentivized Human Persuaders. *arXiv preprint arXiv:2505.09662* (2025).
- [63] Shalom H Schwartz. 1992. Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In *Advances in experimental social psychology*. Vol. 25. Elsevier, 1–65.
- [64] Shalom H Schwartz. 2005. Robustness and fruitfulness of a theory of universals in individual values. *Valores e trabalho* (2005), 56–85.
- [65] Nikhil Sharma, Q. Vera Liao, and Ziang Xiao. 2024. Generative Echo Chamber? Effect of LLM-Powered Search Systems on Diverse Information Seeking. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems* (Honolulu, HI, USA) (CHI '24). Association for Computing Machinery, New York, NY, USA, Article 1033, 17 pages. doi:10.1145/3613904.3642459
- [66] Hua Shen, Chieh-Yang Huang, Tongshuang Wu, and Ting-Hao Kenneth Huang. 2023. ConvXAI: Delivering heterogeneous AI explanations via conversations to support human-AI scientific writing. In *Companion Publication of the 2023 Conference on Computer Supported Cooperative Work and Social Computing*. 384–387.
- [67] Hua Shen, Tiffany Knearem, Reshmi Ghosh, Kenan Alkiek, Kundan Krishna, Yachuan Liu, Ziqiao Ma, Savvas Petridis, Yi-Hao Peng, Li Qiwei, et al. 2024. Towards bidirectional human-ai alignment: A systematic review for clarifications, framework, and future directions. *arXiv preprint arXiv:2406.09264* (2024).
- [68] Hua Shen, Tiffany Knearem, Reshmi Ghosh, Michael Xieyang Liu, Andrés Monroy-Hernández, Tongshuang Wu, Diyi Yang, Yun Huang, Tanushree Mitra, Yang Li, and Marti Hearst. 2025. Bidirectional Human-AI Alignment: Emerging Challenges and Opportunities. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems (CHI EA '25)*. Association for Computing Machinery, New York, NY, USA, Article 857, 6 pages. doi:10.1145/3706599.3716291
- [69] Hua Shen, Vicky Zayats, Johann Rocholl, Daniel Walker, and Dirk Padfield. 2023. MultiTurnCleanup: A Benchmark for Multi-Turn Spoken Conversational Transcript Cleanup. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*. 9895–9903.
- [70] Almog Simchon, Matthew Edwards, and Stephan Lewandowsky. 2024. The persuasive effects of political microtargeting in the age of generative artificial intelligence. *PNAS nexus* 3, 2 (2024), pgae035.
- [71] Herbert W Simons, Nancy N Berkowitz, and R John Moyer. 1970. Similarity, credibility, and attitude change: A review and a theory. *Psychological bulletin* 73, 1 (1970), 1.
- [72] Tom W. Smith, Peter V. Marsden, Michael Hout, and Jibum Kim. 2024. *General Social Survey 2024*. NORC at the University of Chicago. <https://gss.norc.umd.edu/covers-internet-module-items-and-the-nathealth-health-spending-item>.
- [73] Charles S Taber, Damon Cann, and Simona Kucsova. 2009. The motivated processing of political arguments. *Political Behavior* 31, 2 (2009), 137–155.
- [74] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).
- [75] Alexander Todorov, Shelly Chaiken, and Marlone D Henderson. 2002. The heuristic-systematic model of social information processing. *The persuasion handbook: Developments in theory and practice* 23 (2002), 195–211.
- [76] Jan G Voelkel, Robb Willer, et al. 2023. Artificial intelligence can persuade humans on political issues. (2023).
- [77] Xuewei Wang, Weiyang Shi, Richard Kim, Yoojung Oh, Sijia Yang, Jingwen Zhang, and Zhou Yu. 2019. Persuasion for Good: Towards a Personalized Persuasive Dialogue System for Social Good. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Anna Korhonen, David Traum, and Lluís Màrquez (Eds.). Association for Computational Linguistics, Florence, Italy, 5635–5649. doi:10.18653/v1/P19-1566
- [78] Sean J Westwood. 2015. The role of persuasion in deliberative opinion change. *Political Communication* 32, 4 (2015), 509–528.
- [79] Wendy Wood. 2000. Attitude change: Persuasion and social influence. *Annual review of psychology* 51, 1 (2000), 539–570.
- [80] Yuchen Wu, Edward Sun, Kaijie Zhu, Jianxun Lian, Jose Hernandez-Orallo, Aylin Caliskan, and Jindong Wang. 2025. Personalized Safety in LLMs: A Benchmark and A Planning-Based Agent Approach. *arXiv:2505.18882 [cs.CY]* <https://arxiv.org/abs/2505.18882>
- [81] Haolan Zhan, Yufei Wang, Tao Feng, Yuncheng Hua, Suraj Sharma, Zhuang Li, Lizhen Qu, Zhaleh Semmani Azad, Ingrid Zuckerman, and Gholamreza Haffari.

2024. Let's negotiate! a survey of negotiation dialogue systems. *arXiv preprint arXiv:2402.01097* (2024).

[82] Chen Zhang, Xinyi Dai, Yaxiong Wu, Qu Yang, Yasheng Wang, Ruiming Tang, and Yong Liu. 2025. A survey on multi-turn interaction capabilities of large language models. *arXiv preprint arXiv:2501.09959* (2025).