Belief-Sim: Towards Belief-Driven Simulation of Demographic Misinformation Susceptibility

Anonymous ACL submission

Abstract

Misinformation is a growing societal threat, and susceptibility to misinformative claims varies across demographic groups due to differences in underlying beliefs. As Large Language Models (LLMs) are increasingly used to simulate human behaviors, we investigate whether they can simulate demographic misinformation susceptibility, treating *beliefs* as a primary driving factor. We introduce BELIEFSIM, a simulation framework that constructs demographic belief profiles using psychology-informed taxonomies and survey priors. We study prompt-based conditioning and post-training adaptation, and conduct a multi-fold evaluation using: (i) susceptibility accuracy and (ii) counterfactual demographic sensitivity. Across both datasets and modeling strategies, we show that beliefs provide a strong prior for simulating misinformation susceptibility, with accuracy up to 92%.

1 Introduction

Misinformation is a critical challenge in today's information ecosystem, with impacts that vary substantially across demographic groups (Khachaturov et al., 2025; Timm et al., 2025). These differences reflect not only differential targeting (Ribeiro et al., 2019), but also variation in how content is perceived (Guess et al., 2019a). Prior work illustrates these complexities: younger adults were more susceptible to believing COVID-related misinformation in the UK and Brazil (Vijaykumar et al., 2021), while older adults were more likely to share false articles on social media (Guess et al., 2019b; Brashier and Schacter, 2020). Such mixed findings suggest that demographics alone are insufficient to explain susceptibility, i.e., how likely someone is to believe a misinformative claim. Prior studies show that beliefs may play a stronger role (Sultan et al., 2024): people are more likely to accept false information when it aligns or coincides with beliefs they already hold, a phenomenon known as "belief consistency" (Flynn et al., 2017; Taber and Lodge, 2006; Roozenbeek et al., 2020). Therefore, understanding how demographic differences in misinformation susceptibility relate to underlying beliefs is essential for designing effective interventions.

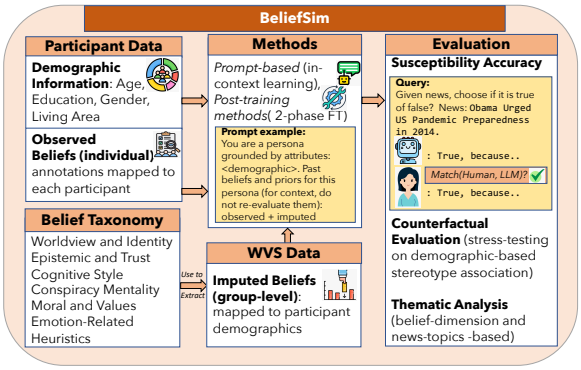

Figure 1: BeliefSim Framework: (1) Participant Data, Observed and Imputed Beliefs (based on Belief Taxonomy) are collected from surveys, (2) Methods consist of prompt-conditioning and post-training adaptation and (3) Evaluation using Susceptibility Accuracy, Counterfactual and Thematic Analysis.

LLMs are increasingly used to simulate social processes (Park et al., 2023; Zhou et al., 2023; Borah et al., 2025a), including behaviors such as biased endorsement and echo chambers (Acerbi and Stubbersfield, 2023; Nehring et al., 2024; Borah et al., 2025b; Sharma et al., 2024). Closest prior work on demographic misinformation susceptibility simulation relies on coarse personas (Borah et al., 2025b), without modeling belief structures. Conversely, most belief-based studies primarily align LLMs with broad political ideologies or moral foundations (Santurkar et al., 2023; Argyle et al., 2023), or simulate general personality traits (Pratelli and Petrocchi, 2025).

Motivated by these gaps, we ask three research questions: (1) Do beliefs improve demographic-aware simulations of misinformation susceptibility? (2) How can such simulations be rigorously

Dimension	Description	Example from WVS
Worldview & Identity	Group-based and political identities shape how individuals interpret and endorse information.	How proud are you to be from your country?
Epistemic Trust	Beliefs about what sources are credible and how knowledge should be evaluated (e.g., trust in science, media, institutions).	Would you say that the world is better off, or worse off, because of science?
Cognitive Style	Preferred thinking approach, ranging from analytical, accuracy-driven reasoning to intuitive impression-based judgments.	Nowadays, does one often have trouble deciding which moral rules to follow?
Conspiracy Mentality	General tendency to perceive hidden plots or malevolent intent behind significant social or political events.	How many people who are in the media do you believe are involved in corruption?
Morals & Values	Core moral values and cultural worldviews influence judgments of right, wrong, fairness, and societal norms.	Is it justifiable - claiming government benefits to which you are not entitled?
Emotion-Related	Emotion-driven beliefs related to fear, perceived risk, or threat sensitivity that heighten susceptibility to alarming misinformation.	How satisfied are you with your life as a whole these days?
Heuristics	Reliance on mental shortcuts (e.g., familiarity implies truth) that increase vulnerability to repeated or simplified claims.	Please indicate how much you use it: Daily newspaper, TV News, Radio News.

Table 1: Belief taxonomy for modeling demographic-aware misinformation susceptibility in LLMs. We provide representative WVS items as examples of how survey questions align with each belief dimension.

evaluated using utility and counterfactual demographic analyses? and (3) What modeling strategies are best suited for LLM-based simulation?

To answer these questions, we propose **BeliefSim**, a belief-driven simulation framework for demographic misinformation susceptibility, illustrated in Fig 1. We summarize our contributions as follows: **(1)** We construct a belief taxonomy and build a simulation dataset by aggregating data from existing surveys and studies; **(2)** We conduct empirical analyses to identify factors across belief and demographic dimensions that are most important for the simulation; **(3)** We explore simulation techniques including prompt-based and post-training approaches using several LLMs; and **(4)** We perform a counterfactual study to understand when and how demographic information may provide useful priors vs stereotype-like sensitivity. Finally, we outline actionable steps that can help design future intervention methods.

2 Related Work

Psychological research has established that susceptibility to misinformation is heavily driven by *belief consistency*. While demographic factors like age, gender, and education provide significant predictive signals, recent large-scale meta-analyses suggest these effects can be context- and belief-dependent Sultan et al. (2024). However, most studies that establish benchmarks with large-scale human annotations face an inherent scalability bottleneck (Maertens et al., 2024; Borah et al., 2025b). High-quality human data collection is resource-intensive and slow, often limiting research to static snapshots of specific news cycles (e.g. COVID-19) or narrow domains.

The emerging capability of LLMs to simulate human behavior offers a potential solution. Early work showed that LLMs could act as "silicon subjects" via demographic prompting (Argyle et al., 2023). Studies also investigated if LLMs reproduce classic social science findings and moral values (Park et al., 2023; Borah et al., 2025a; Nair and Wang, 2025). However, fidelity often degrades when simulations depend only on demographic labels (Giorgi et al., 2024), leading to stereotypical associations (Borah and Mihalcea, 2024).

To move beyond demographics, recent research has focused on conditioning models with richer context such as beliefs (Moon et al., 2024; Namikoshi et al., 2024). Most relevant to our work, Chuang et al. (2024) introduced "Human Belief Networks", showing that seeding an agent with even a single belief (e.g., a stance on welfare) was more predictive of downstream responses than the agent's entire demographic profile. However, prior belief-focused studies have largely not examined their use in simulating misinformation susceptibility. Our work bridges this gap by focusing on this specific domain, where the interplay between demographics and beliefs is highly complex.

3 Belief Taxonomy

Beliefs are an important tool for modeling demographic simulations for misinformation. Prior work has shown that belief dimensions, such as conspir-

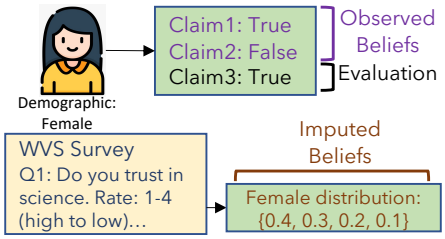

Figure 2: Simulation Data Example.

acy beliefs or trust in science, can help predict susceptibility to misinformation (Ecker et al., 2022; Munusamy et al., 2024).

Grounded in psychological and cognitive science research, we compiled a belief taxonomy shown in Table 1, that includes seven core dimensions most associated with misinformation susceptibility: (1) Worldview and Identity Beliefs (Kahan, 2017; Van Bavel et al., 2021), (2) Epistemic Trust Beliefs (De Coninck et al., 2021; Lewandowsky et al., 2023), (3) Cognitive style (Pennycook and Rand, 2019; Ecker et al., 2022), (4) Conspiracy mentality (Douglas et al., 2019; De Coninck et al., 2021), (5) Moral and Value Beliefs (D'Errico et al., 2022; Yang et al., 2024), (6) Emotion-Related Beliefs (Brady et al., 2017; McLoughlin et al., 2024) and (7) Heuristic Beliefs (Lin et al., 2016; Fazio, 2020). While the above dimensions have traditionally been studied in isolation and within human populations, we unify them into a structured taxonomy designed for computational modeling (Further details are provided in Appendix A. This enables systematic belief simulation by providing a framework for organizing belief data and evaluating it.

4 Simulation Data

For our study, we consider four demographic axes: Gender (female/male), Age (younger: <=35 years/older: >=60 years), Living Area (Rural/Urban), and Education (completed high school/not completed high school).[1] These demographics are commonly associated with systematic differences in media exposure, institutional trust, etc., all of which influence vulnerability to misinformation (Allcott and Gentzkow, 2017; Guess et al., 2019b; Allcott et al., 2019; Anspach and Carlson, 2024). Focusing on these axes enables controlled analysis of demographic context for simulating susceptibility signals in LLMs.

[1]For ease of analysis, we focus on binary groupings, leaving finer-grained analysis to future work.

4.1 Evaluation Data

For susceptibility evaluation, we use two ground-truth datasets containing human judgments of whether they believe a given claim: (1) PANDORA Dataset from Borah et al. (2025b), containing annotations from 318 participants. Each participant provided judgments on 3 distinct claims, along with demographic information including age, gender, living area, and education. These claims are collected from RumorEval (Gorrell et al., 2019), which consists of true or false rumors covering eight major news events and natural disaster events; (2) MIST dataset from Maertens et al. (2024). From MIST, we use only Study 1 (MIST-1) which includes 409 participants, each providing judgment on the same 100 claims, and demographic information including age, gender, and education. From the two datasets combined, we obtain 13.8K claims for evaluation.

4.2 Belief Data

We collect two complementary belief signals, consisting of individually observed belief judgments and group-level (demographic) belief distributions:

(1) Observed Data - claims that were directly judged by participants in PANDORA and MIST-1 datasets. These responses capture individual-level belief judgments, reflecting each participant's personal stance rather than demographic group averages. We use two claim judgments as observed beliefs for each participant (keeping it separate from the evaluation data). This leads to 27.6 claim judgments as observed beliefs.

(2) Imputed Data - inferred from the World Values Survey Wave 7[2] distributions. Imputed data represent demographic belief priors (group–level), inferred from WVS, conditioned solely on demographic attributes. We map these imputed belief items to our belief taxonomy using exploratory factor analysis. This yields 126 imputed belief questions and corresponding demographic distributions. Table 1 shows examples, and Appendix A.2 contains mapping details and all questions.

Fig. 2 provides an example of simulation data. We restrict all our analysis to U.S.-based participants and English headlines, consistent with the scope of the available datasets, while still leveraging variation in living area, age, gender, and education.

[2]https://www.worldvaluessurvey.org/WVSDocumentationWV7.jsp

5 Prompt-based Conditioning for Susceptibility Simulation

We first use prompt-based methods, where each prompt conditions the model on participant demographics and belief signals (both observed and imputed) from the surveys. Note that we simulate each demographic axis separately, as it enables controlled analysis of demographic-level differences.

Method. For WVS questions, answers are either Likert scales (Bertram, 2007) or Yes/No questions. Therefore, we input modal (most frequent) responses in the prompt as beliefs. For example, consider the WVS question "How important is religion in your life?", rated on a 1–10 scale. If a demographic group most frequently responds with a value of 3, we use this modal response in the prompt. Given demographic group <d>, and belief annotations <b>, we use: `You are a persona grounded by attributes: <demographic group>. Past beliefs and priors for this persona (for context, do not re-evaluate them): <beliefs>. When judging a claim, stay consistent with this persona's prior beliefs where reasonable.`

We perform experiments under four primary input conditions for prompt-based conditioning: (1) zero-shot (without any demographic or belief information), (2) demographic information only, (4) belief information only, and (4) both demographics+beliefs. We additionally ablate over belief dimensions (the seven in our taxonomy) and belief sources (observed vs. imputed), as well as over demographic attributes (analyzing groups separately). These experiments yield several insights across 12 total settings. We perform these experiments using three LLMs: `llama-3-8b-instr.`, `qwen-2.5-14b-instr.` and `mistral-7b-instr.-v02`, and average our results across three separate runs.

Evaluation. We measure susceptibility accuracy: whether the LLM correctly predicts if an individual would judge a claim as true or false.

5.1 Results

Fig 3 shows the susceptibility accuracies across the belief and demographic configurations for both datasets. For cases that include `Imputed` beliefs, we only report `best` results by selecting the belief dimension that achieves the highest accuracy. Findings show that incorporating belief information consistently improves performance over the zero-shot and demographic-only baselines. We next analyze the impact of specific belief types, demographic ablations, datasets, and model performances. Appendix C contains detailed results for all settings.

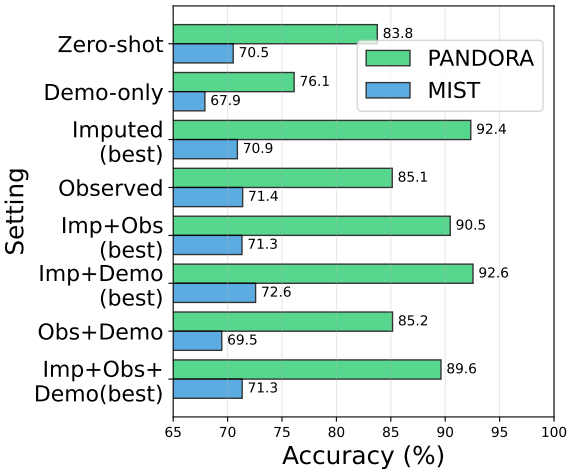

Figure 3: Susceptibility Accuracies by settings and datasets, averaged across demographic groups and models. Imputed + Demo(graphic) (best) performs the best, with all belief included settings doing better than zero-shot and demo-only.

Demographic vs Belief information. Overall, beliefs drive most of the positive gains. Demo-only degrades zero-shot performances for both datasets, while beliefs-only (imputed) show larger gains (+1-16) points over demo-only and zero-shot settings. Imputed + Demo achieves the highest performance. This shows that what people *believe* may be more predictive than *who they are*. This is also in line with findings from Borah et al. (2025a), which show that belief-based signals are stronger than demographic signals in LLMs.

Imputed vs Observed Beliefs. Across settings, imputed beliefs consistently outperform observed beliefs, more so for PANDORA. A likely reason is that observed beliefs, being tied to specific participants, can be sparse, or mismatched to the new claim, while imputed demographic beliefs from WVS may act as *smoother population priors*. Furthermore, combining imputed+observed helps but is still lower than imputed-only. This may be due to conflict or being on different scales, and the individualistic observed signal can dilute the cleaner imputed prior. Average accuracy trends are: $imputed > imputed + observed > observed$.

Imputed Belief Ablations. Adding one belief dimension at a time to the prompt consistently yields higher performance than including all belief dimensions together. Across datasets, emotion-related

and moral-values are the strongest. These also vary by demographics: emotion-related is the strongest in Gender; moral-values is the strongest for Rural/Urban and Education; heuristics/cognitive beliefs are the strongest in Education (Detailed results are in Appendix C.2).

Demographic Ablations. Adding demographic information on top of beliefs yields small gains (in the range 0.2-1.2%), especially for imputed beliefs. Thus suggesting that demographics serve as an additional context signal, helping with relevant belief priors. On the effect of belief across demographic groups, accuracies for age, education and living area are improved with belief addition, specifically using the imputed beliefs. Gender shows the largest sensitivity to which beliefs are included; prior work also show that gender effects are often smaller/mixed in misinformation contexts (Sultan et al., 2024). This underscores the need for caution in selecting belief evidence for gender-specific cases.

DATASET	LLAMA	QWEN	MISTRAL
PANDORA	85.90	88.84	87.51
MIST	69.77	73.13	72.71

Table 2: Model Performance comparison across datasets, accuracy trends: qwen > mistral > llama.

Model and Dataset comparison. Across models, `Qwen` consistently shows higher accuracies, followed by `mistral` and `llama`, as shown in Table 2.

Across datasets, the PANDORA Dataset has higher accuracy scores than the MIST dataset across all setups. PANDORA yields a high zero-shot baseline (83.8%) and gains the highest from belief integration in the (best imputed + demographic) setting (+8.8). In contrast, MIST has lower zero-shot performance (70.5%), with the highest gain in (best imputed + demographic) setting (+2). Therefore, the impact of adding beliefs also varies with dataset scale: Prolific contains much fewer examples than MIST (13000 vs 318). Nevertheless, adding belief priors is beneficial across datasets.

Thematic Analysis. We cluster news claims into latent topics and test whether LLM susceptibility accuracy differs across demographics by topic. We conduct our analysis on the MIST dataset, as it is much larger than PANDORA. We find that demographic differences in susceptibility are topic-dependent: clusters corresponding to lexically obvious or opinion-based claims exhibit minimal demographic variation, whereas more ambiguous or diverse topics like science/health claims and government show larger differences across age (older groups scoring higher) and education groups (higher-educated groups scoring higher). Gender-based differences are consistently small across all topics. These findings suggest that demographic susceptibility is also contextual and topic-sensitive, motivating further evaluation in future studies (More details are in Appendix C.3).

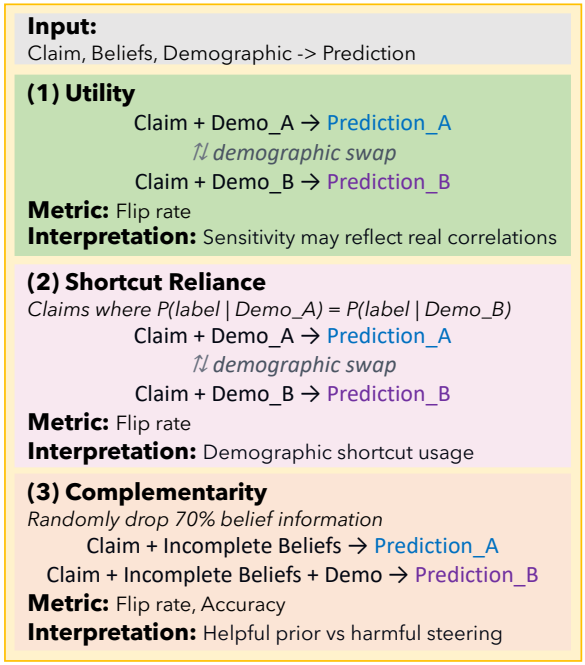

Figure 4: Demographic-based Counterfactual Eval Framework: (1) Utility of demographics, (2) Shortcut Reliance, and (3) Complementarity. Note that we perform a demographic swap within the same demographic group, i.e., swap male with female, rural with urban, etc. keeping the claim constant.

5.2 Demographic-based Counterfactual Evaluation

Demographics can be a helpful context signal in addition to belief information and legitimately correlate with misinformation susceptibility in real data, but they can also induce spurious signals that appear accurate but could be stereotypical (Wan et al., 2025; Geirhos et al., 2020). To investigate this trade-off, we conduct counterfactual evaluations to assess whether demographic attributes provide informative priors or introduce stereotype-sensitive associations.

Method. We conduct three complementary analyses on the PANDORA+MIST-1 data to investigate effects of demographics in suceptibility prediction (Fig 4):

(1) Utility: Do demographics improve susceptibil-

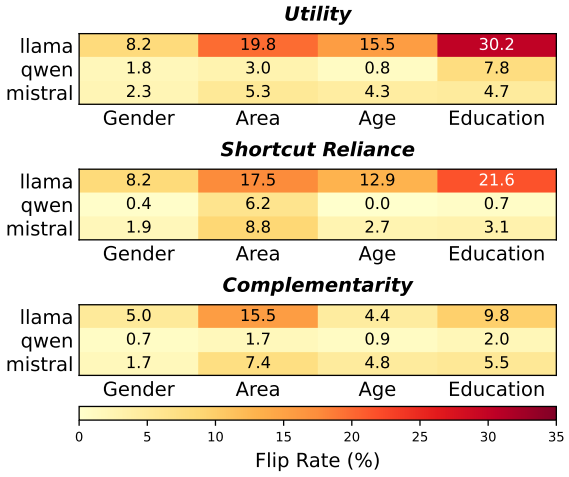

Figure 5: Flip-Rates for Counterfactual Evaluation: qwen and mistral models have lower flip-rates in comparison to llama.

ity prediction beyond claim content? We first evaluate utility on the original data distribution using the demographic-only setting. Next, we create a counterfactual demographic swap (keeping the query claim fixed), re-prompt the model, and measure the flip rate between predictions. Higher flip rates indicate stronger demographic sensitivity, which may reflect real-world demographic–label correlations rather than purely spurious associations.

(2) Shortcut reliance: Does the model rely on demographic cues even when they are non-informative by construction? Here, we design a controlled subset where the distribution of human labels in each claim is matched across demographic groups (e.g., Female and Male have equal counts of 'true' and 'misinfo' for that claim). Therefore, this makes the demographic attribute non-predictive. A high flip rate here indicates demographic short-cutting and possible stereotyping.

(3) Complementarity: When belief evidence is incomplete, do demographics act as a helpful weak prior or a harmful shortcut? We similate missing evidence by randomly dropping a large fraction of WVS belief statements (70%). We then compare model predictions under two conditions: (1) belief only and (2) belief + demographic. High flip rates between these conditions indicate that demographic cues influence predictions under uncertainty. Additionally, we compare susceptibility prediction accuracy in both conditions to assess whether adding demographics improves performance.

Results. Fig 5 shows the flip-rates per model and demographic groups averaged across datasets. Overall, qwen consistently achieves lower flip rates

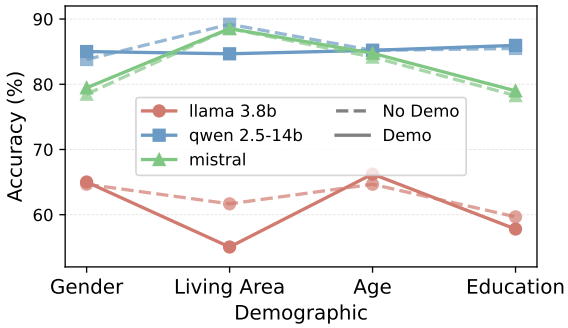

Figure 6: Complementarity Analysis Results: when partial belief information is added, adding demographics has mixed effects, showing very weak gains.

(averaging 2.16%), suggesting lower dependence on demographic cues. Mistral also exhibits lower flip rates (averaging 4.3%). Contrastingly, llama shows substantially higher flip rates (averaging 14.31%), thus higher demographic sensitivity. Across models, the largest flip rates concentrate in Education and Living Area, especially in the Shortcut Reliance setting, where demographics are designed to be non-predictive. This could also relate to stereotype-driven reliance. For complementarity analysis, Fig 6 shows that demographics are only weakly complementary: mistral shows small, consistent gains across settings. qwen improves slightly for education/gender, while drops for living area. llama also shows unstable patterns.

Overall, these findings suggest that demographic sensitivity is highly model-dependent. qwen and mistral are comparatively stable under counterfactual demographic changes, while llama shows substantially higher reliance on demographics, even when demographics are explicitly made non-predictive (Shortcut Reliance task). At the same time, demographics provide limited practical benefit when partial beliefs are present, showing they are a weak prior at best. These findings motivate moving beyond prompt-only conditioning, as we find that demographic cues in the prompt can introduce unstable shortcuts. Therefore, belief signal must be integrated in a way that is robust to conflicting sources (individual and group-based beliefs).

6 Post-training Adaptation for Susceptibility Simulation

Prompt-based conditioning experiments show the benefits of using beliefs for demographic misinformation susceptibility simulation. However, mixing imputed and observed beliefs in the prompt is associated with lower performance. This is intuitive, as

observed (ground-truth) beliefs are tied to individual participants and imputed beliefs are only linked to demographic groups and not individuals, creating inconsistencies. Furthermore, training directly on demographics and susceptibility labels risks leakage, where the model may learn label shortcuts as observed in counterfactual evaluation (Wan et al., 2025; Geirhos et al., 2020). To address this, we propose ***BAFT*** (Belief-Adapter Fine-tuning), which decouples imputed beliefs from observed data, while still learning transferable belief representations through a separate adapter.

Method. To decouple group-based imputed beliefs from individualistic observed beliefs, BAFT consists of a two-phase design: (1) belief modeling, and (2) susceptibility fine-tuning.

Phase 1: Belief Modeling. We train a belief adapter by freezing the base LLM and training a belief head, implemented as a linear projection followed by softmax, to predict demographic-conditioned response distributions for WVS belief questions. Concretely, given the encoder representation $h_\phi(q, d)$ for question q and demographic context d, we model $p_\theta(y \mid q, d) = \text{softmax}(W h_\phi(q, d) + b)$. This captures population-level variability in belief responses. Importantly, this stage uses only survey supervision (no misinformation labels), reducing the risk of label leakage or reward hacking.

Phase 2: Susceptibility Fine-tuning. In this phase, we freeze the belief adapter and train a lightweight susceptibility head that combines the base LLM's semantic representation of the prompt with the belief representation produced by the adapter. Given a persona-conditioned input (demographics and observed beliefs) and target claim x, the frozen adapter outputs a belief embedding z_{bel}, which we concatenate with the claim representation $h_\phi(x)$ and predict susceptibility via a binary classifier:

$$p_\psi(y \mid x, d) = \text{softmax}(U\,[\,h_\phi(x)\,;\, z_{\text{bel}}(d)\,] + c)\,, \quad (1)$$

where $y \in \{\text{true}, \text{misinformation}\}$. We train only the susceptibility head parameters $\psi = \{U, c\}$ with cross-entropy loss, keeping the base model and belief adapter frozen. This is to ensure susceptibility learning relies on transferable belief structure rather than encoding demographic shortcuts.

Data. For Phase 1, we construct demographic belief priors from empirical WVS response distributions. For each question q and demographic group d, we estimate a categorical distribution $P(r \mid q, d)$ over response options $r \in \{1, \ldots, K_q\}$. These distributions capture meaningful demographic variability across WVS belief questions. We restrict our experiments to six demographic groups, as MIST does not include living-area annotations, and the PANDORA data is too small to support standalone data for fine-tuning. This results in 126 WVS response distributions per demographic group, yielding $6 \times 125 = 750$ examples in total. We apply textual augmentation via back-translation using Google Translate,[3] producing 1250 examples. This is sufficient as the belief adapter is a lightweight linear head requiring fewer examples to reliably capture demographic belief patterns. Training is performed using an 80/20 train/val split.

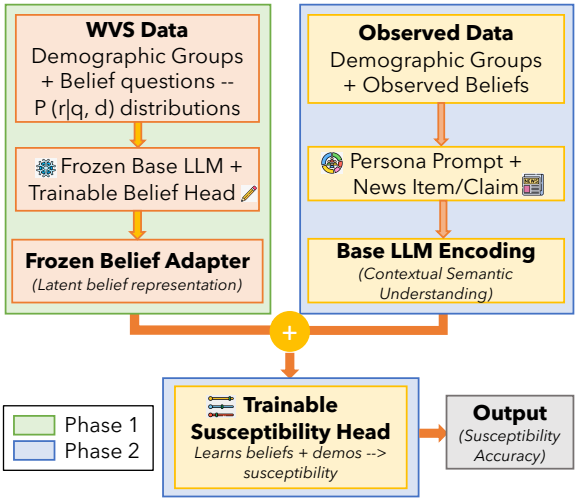

Figure 7: Belief Adapter Fine-Tuning framework. Green-shaded components correspond to Phase 1 (Belief Modeling), and blue-shaded components correspond to Phase 2 (Susceptibility Fine-Tuning).

For Phase 2, we use (PANDORA+MIST-1) data, keeping each demographic group separate per data point (as in prompt conditioning). The target is the prediction on a claim, giving susceptibility scores. We use a 80/20 train/val split, yielding approximately 33.1k training examples and 5.1k validation examples (after removal of overlapping examples). For evaluation, we also used the (MIST Study-2b) (to investigate cross-study generalization), consisting of 7k evaluation points.

We compare BAFT against two strong baselines: (1) standard LoRA fine-tuning using PANDORA+MIST-1 (Baseline (1-phase)), and (2) a two-phase variant of our pipeline that replaces the head-based adaptation with LoRA fine-tuning (LoRA-FT (2-phase)).

[3] https://translate.google.com/

Model	PANDORA+MIST-1		MIST-2	
	Acc	F1	Acc	F1
Baseline (1-phase)				
llama	0.750	0.751	0.680	0.673
qwen	0.760	0.761	0.720	0.712
mistral	0.740	0.761	0.650	0.653
LoRA-FT (2-phase)				
llama	0.809	0.809	0.876	0.861
qwen	0.800	0.800	0.893	0.892
mistral	0.800	0.790	0.887	0.867
BAFT (2-phase)				
llama	0.793	0.793	0.884	0.896
qwen	0.792	0.791	0.924	0.906
mistral	0.792	0.792	0.884	0.823

Table 3: Fine-tuning performances. 2-phase belief training improves performance across both datasets.

Evaluation. Phase 1 reflects population-level belief fidelity, so we use distribution-level metric over Likert-scale responses, Kullback–Leibler (KL) divergence (Kullback and Leibler, 1951). This measures how well predicted belief distributions align with each demographic (lower is better for both). Phase 2 reflects task utility. So, we compute susceptibility accuracy for this phase.

6.1 Results

Phase 1. `Qwen` has the lowest KL-divergence (0.051), followed by `mistral` (0.087) and `llama` (0.287) for BAFT. With LoRA-ft, the trends remain similar (complete results are in Appendix D).

Phase 2. Table 3 shows that cross-study (MIST-2) generalization improves drastically once we move from 1-phase to 2-phase. The 1-phase baseline model performs reasonably well on PANDORA+MIST-1, however, it drops sharply on MIST-2. Comparing the 2-phase variants, BAFT remains the strongest on MIST-2 (up to 92.4%). LoRA-ft slightly improves on PANDORA+MIST-1 but gives a weaker transfer, may be due to mild overfitting when more parameters are updated. Across models, `qwen` consistently performs the best for 2-phase settings. Additionally, we perform shortcut reliance experiments, finding that BAFT leads to 0% flip-rates even for `llama` models (which earlier had the highest flip-rate: Appendix D.3).

7 Lessons Learned and Actionable Steps

Our findings show that belief priors are central to simulating demographic misinformation susceptibility with LLMs. We introduce Belief-Sim and BAFT, demonstrating that decoupling belief modeling from susceptibility prediction enables higher accuracy while reducing spurious demographic sensitivity. These findings offer actionable information for designing targeted demographic interventions.

Beliefs are crucial to demographic susceptibility simulation. Across both simulation techniques, adding belief priors (especially imputed) leads to high accuracy gains. Future studies can investigate additional survey sources (e.g., Pew and other cross-national surveys) to broaden coverage and focus on diverse demographic groups and contexts. Furthermore, studying the effective decoupling of different types of beliefs is important.

Simple head tuning is effective for generalization. Our two-phase BAFT approach generalizes better than full fine-tuning, suggesting that lightweight, modular adaptation can be both effective and cheaper for simulation. Future work can further improve robustness by exploring alternative adapter designs and fusion mechanisms for combining belief and text representations.

Counterfactual evaluation is important. Our counterfactual experiments quantify each model's reliance on demographic shortcuts. Such measures are necessary as sensitivity can be either a real or a spurious cue. Future work can build on this by developing fine-grained stress tests and evaluating which demographic cues drive model flips.

8 Conclusion

This paper introduced Belief-Sim, a belief-driven framework for simulating demographic misinformation susceptibility using LLMs. Across two datasets and through prompt-based conditioning and post-training adaptation, we show that belief priors are a key driver for demographic-level susceptibility patterns. In contrast, demographics alone are an unreliable signal that can induce shortcut reliances. Our results highlight the importance of decoupling imputed and observed belief sources and we provide practical evaluation tools for it, such as prediction accuracy and counterfactual sensitivity. Based on these findings, we outline directions for future work and release our open-source framework, BeliefSim.[4]

[4] https://anonymous.4open.science/r/belief-sim

Limitations

Coverage and Intersectionality

Our demographic modeling is intentionally largely single-axis: we evaluate on a small set of demographic attributes (8 groups) and treat them independently. Our goal was to isolate the effects of simulation across demographic groups and simplify counterfactual analysis as well as balance labels across groups. However, intersectional groups (e.g., older x low-education, female x rural) can have interesting implications for simulation purposes, and future work can further investigate Belief-Sim for intersectional belief profiles. In addition, while we mostly adhere to binary demographic axes, each axis can be further analyzed in a finer-grained approach. Finally, our study is based on US participants only, due to the availability of current data sources. We acknowledge that these may have WEIRD implications (Mihalcea et al., 2025), and future work should focus on investigating wider, cross-cultural misinformation susceptibility.

Counterfactual Flips do not always mean stereotyping

Our flip-rate analyses quantify sensitivity to demographic perturbations, but flips do not uniquely measure stereotype-driven shortcutting. Furthermore, low flip rates do not guarantee fairness - models may still encode demographic effects indirectly through correlations or beliefs. Although our experiments help isolate spurious dependence, our experiments may not capture real-world causal pathways as they require much extensive analysis. Future work can investigate extending our counterfactual flips experiments with stronger causal/robustness evaluations to better understand causal demographic effects and/or measure stereotype-driven shortcutting.

Experimental Limitations

Access to closed-source LLMs is constrained by transparency and cost. Therefore, we limit our experiments to three open-source models (to support controlled counterfactual swaps, ablations, and multiple runs). As a result, our conclusions may not fully generalize to the strongest closed-source models. Future work can extend our framework by benchmarking a broader set of closed-source models under matched prompting and budgeted sampling.

References

Alberto Acerbi and Joseph M Stubbersfield. 2023. Large language models show human-like content biases in transmission chain experiments. *Proceedings of the National Academy of Sciences*, 120(44):e2313790120.

Hunt Allcott and Matthew Gentzkow. 2017. Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–236.

Hunt Allcott, Matthew Gentzkow, and Chuan Yu. 2019. Trends in the diffusion of misinformation on social media. *Research & politics*, 6(2):2053168019848554.

Nicolas M Anspach and Taylor N Carlson. 2024. Not who you think? exposure and vulnerability to misinformation. *New Media & Society*, 26(8):4847–4866.

Lisa P. Argyle, Ethan C. Busby, Nancy Fulda, Joshua R. Gubler, Christopher Rytting, and David Wingate. 2023. Out of one, many: Using language models to simulate human samples. *Political Analysis*, 31(3):337–351.

Dane Bertram. 2007. Likert scales. *Retrieved November*, 2(10):1–10.

Angana Borah, Marwa Houalla, and Rada Mihalcea. 2025a. Mind the (belief) gap: Group identity in the world of LLMs. In *Findings of the Association for Computational Linguistics: ACL 2025*, pages 18441–18463, Vienna, Austria. Association for Computational Linguistics.

Angana Borah and Rada Mihalcea. 2024. Towards implicit bias detection and mitigation in multi-agent LLM interactions. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 9306–9326, Miami, Florida, USA. Association for Computational Linguistics.

Angana Borah, Rada Mihalcea, and Verónica Pérez-Rosas. 2025b. Persuasion at play: Understanding misinformation dynamics in demographic-aware human-llm interactions. *arXiv preprint arXiv:2503.02038*.

William J Brady, Julian A Wills, John T Jost, Joshua A Tucker, and Jay J Van Bavel. 2017. Emotion shapes the diffusion of moralized content in social networks. *Proceedings of the National Academy of Sciences*, 114(28):7313–7318.

Nadia M Brashier and Daniel L Schacter. 2020. Aging in an era of fake news. *Current directions in psychological science*, 29(3):316–323.

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

David De Coninck, Thomas Frissen, Koen Matthijs, Leen d'Haenens, Grégoire Lits, Olivier Champagne-Poirier, Marie-Eve Carignan, Marc D David, Nathalie Pignard-Cheynel, Sébastien Salerno, et al. 2021. Beliefs in conspiracy theories and misinformation about covid-19: Comparative perspectives on the role of anxiety, depression and exposure to and trust in information sources. *Frontiers in psychology*, 12:646394.

Karen M Douglas, Joseph E Uscinski, Robbie M Sutton, Aleksandra Cichocka, Turkay Nefes, Chee Siang Ang, and Farzin Deravi. 2019. Understanding conspiracy theories. *Political psychology*, 40:3–35.

Francesca D'Errico, Giuseppe Corbelli, Concetta Papapicco, and Marinella Paciello. 2022. How personal values count in misleading news sharing with moral content. *Behavioral Sciences*, 12(9):302.

Ullrich KH Ecker, Stephan Lewandowsky, John Cook, Philipp Schmid, Lisa K Fazio, Nadia Brashier, Panayiota Kendeou, Emily K Vraga, and Michelle A Amazeen. 2022. The psychological drivers of misinformation belief and its resistance to correction. *Nature Reviews Psychology*, 1(1):13–29.

Jonathan St BT Evans and Keith E Stanovich. 2013. Dual-process theories of higher cognition: Advancing the debate. *Perspectives on psychological science*, 8(3):223–241.

Lisa K Fazio. 2020. Repetition increases perceived truth even for known falsehoods. *Collabra: Psychology*, 6(1).

Daniel J Flynn, Brendan Nyhan, and Jason Reifler. 2017. The nature and origins of misperceptions: Understanding false and unsupported beliefs about politics. *Political psychology*, 38:127–150.

Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.

Salvatore Giorgi, Tingting Liu, Ankit Aich, Kelsey Jane Isman, Garrick Sherman, Zachary Fried, João Sedoc, Lyle Ungar, and Brenda Curtis. 2024. Modeling human subjectivity in LLMs using explicit and implicit human factors in personas. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7174–7188, Miami, Florida, USA. Association for Computational Linguistics.

Genevieve Gorrell, Elena Kochkina, Maria Liakata, Ahmet Aker, Arkaitz Zubiaga, Kalina Bontcheva, and Leon Derczynski. 2019. SemEval-2019 task 7: RumourEval, determining rumour veracity and support for rumours. In *Proceedings of the 13th International Workshop on Semantic Evaluation*, pages 845–854, Minneapolis, Minnesota, USA. Association for Computational Linguistics.

Andrew Guess, Jonathan Nagler, and Joshua Tucker. 2019a. Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science Advances*, 5(1):eaau4586.

Andrew Guess, Jonathan Nagler, and Joshua Tucker. 2019b. Less than you think: Prevalence and predictors of fake news dissemination on facebook. *Science advances*, 5(1):eaau4586.

Dan M Kahan. 2017. Misconceptions, misinformation, and the logic of identity-protective cognition.

Carolin Kaiser, Jakob Kaiser, Vladimir Manewitsch, Lea Rau, and Rene Schallner. 2025. Simulating human opinions with large language models: Opportunities and challenges for personalized survey data modeling. In *Adjunct Proceedings of the 33rd ACM Conference on User Modeling, Adaptation and Personalization*, pages 82–86.

David Khachaturov, Roxanne Schnyder, and Robert Mullins. 2025. Governments should mandate tiered anonymity on social-media platforms to counter deepfakes and LLM-driven mass misinformation. *arXiv preprint arXiv:2506.12814*.

Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1):79–86.

Stephan Lewandowsky, Ullrich KH Ecker, John Cook, Sander Van Der Linden, Jon Roozenbeek, and Naomi Oreskes. 2023. Misinformation and the epistemic integrity of democracy. *Current opinion in psychology*, 54:101711.

Xialing Lin, Patric R Spence, and Kenneth A Lachlan. 2016. Social media and credibility indicators: The effect of influence cues. *Computers in human behavior*, 63:264–271.

Rakoen Maertens, Friedrich M Götz, Hudson F Golino, Jon Roozenbeek, Claudia R Schneider, Yara Kyrychenko, John R Kerr, Stefan Stieger, William P McClanahan, Karly Drabot, et al. 2024. The misinformation susceptibility test (mist): A psychometrically validated measure of news veracity discernment. *Behavior Research Methods*, 56(3):1863–1899.

Ioana E Marinescu, Patrick N Lawlor, and Konrad P Kording. 2018. Quasi-experimental causality in neuroscience and behavioural research. *Nature human behaviour*, 2(12):891–898.

Killian L McLoughlin, William J Brady, Aden Goolsbee, Ben Kaiser, Kate Klonick, and MJ Crockett. 2024. Misinformation exploits outrage to spread online. *Science*, 386(6725):991–996.

Rada Mihalcea, Oana Ignat, Longju Bai, Angana Borah, Luis Chiruzzo, Zhijing Jin, Claude Kwizera, Joan Nwatu, Soujanya Poria, and Thamar Solorio. 2025. Why ai is weird and shouldn't be this way: Towards ai for everyone, with everyone, by everyone. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 28657–28670.

Suhong Moon, Marwa Abdulhai, Minwoo Kang, Joseph Suh, Widyadewi Soedarmadji, Eran Kohen Behar, and David M. Chan. 2024. Virtual personas for language models via an anthology of backstories. *Preprint*, arXiv:2407.06576.

Shalini Munusamy, Kalaivanan Syasyila, Azahah Abu Hassan Shaari, Muhammad Adnan Pitchan, Mohammad Rahim Kamaluddin, and Ratna Jatnika. 2024. Psychological factors contributing to the creation and dissemination of fake news among social media users: a systematic review. *BMC psychology*, 12(1):673.

Inderjeet Nair and Lu Wang. 2025. Do language models think consistently? a study of value preferences across varying response lengths. *arXiv preprint arXiv:2506.02481*.

Keiichi Namikoshi, Alex Filipowicz, David A Shamma, Rumen Iliev, Candice L Hogan, and Nikos Arechiga. 2024. Using llms to model the beliefs and preferences of targeted populations. *arXiv preprint arXiv:2403.20252*.

Jan Nehring, Aleksandra Gabryszak, Pascal Jürgens, Aljoscha Burchardt, Stefan Schaffer, Matthias Spielkamp, and Birgit Stark. 2024. Large language models are echo chambers. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10117–10123, Torino, Italia. ELRA and ICCL.

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.

Gordon Pennycook and David G Rand. 2019. Lazy, not biased: Susceptibility to partisan fake news is better explained by lack of reasoning than by motivated reasoning. *Cognition*, 188:39–50.

Manuel Pratelli and Marinella Petrocchi. 2025. *Evaluating the Simulation of Human Personality-Driven Susceptibility to Misinformation with LLMs*. IOS Press.

Filipe N. Ribeiro, Koustuv Saha, Mahmoudreza Babaei, Lucas Henrique, Johnnatan Messias, Fabricio Benevenuto, Oana Goga, Krishna P. Gummadi, and Elissa M. Redmiles. 2019. On microtargeting socially divisive ads: A case study of russia-linked ad campaigns on facebook. In *Proceedings of the Conference on Fairness, Accountability, and Transparency*, FAT* '19, page 140–149, New York, NY, USA. Association for Computing Machinery.

Jon Roozenbeek, Claudia R Schneider, Sarah Dryhurst, John Kerr, Alexandra LJ Freeman, Gabriel Recchia, Anne Marthe Van Der Bles, and Sander Van Der Linden. 2020. Susceptibility to misinformation about covid-19 around the world. *Royal Society open science*, 7(10):201199.

Shibani Santurkar, Esin Durmus, Faisal Ladhak, Cinoo Lee, Percy Liang, and Tatsunori Hashimoto. 2023. Whose opinions do language models reflect? *Preprint*, arXiv:2303.17548.

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*, pages 1–17.

Mubashir Sultan, Alan N Tump, Nina Ehmann, Philipp Lorenz-Spreen, Ralph Hertwig, Anton Gollwitzer, and Ralf HJM Kurvers. 2024. Susceptibility to online misinformation: A systematic meta-analysis of demographic and psychological factors. *Proceedings of the National Academy of Sciences*, 121(47):e2409329121.

Charles S Taber and Milton Lodge. 2006. Motivated skepticism in the evaluation of political beliefs. *American journal of political science*, 50(3):755–769.

Jasper Timm, Chetan Talele, and Jacob Haimes. 2025. Tailored truths: Optimizing llm persuasion with personalization and fabricated statistics. *arXiv preprint arXiv:2501.17273*.

Jay J Van Bavel, Diego A Reinero, Victoria Spring, Elizabeth A Harris, and Annie Duke. 2021. Speaking my truth: Why personal experiences can bridge divides but mislead. *Proceedings of the National Academy of Sciences*, 118(8):e2100280118.

Santosh Vijaykumar, Yan Jin, Daniel Rogerson, Xuerong Lu, Swati Sharma, Anna Maughan, Bianca Fadel, Mariella Silva de Oliveira Costa, Claudia Pagliari, and Daniel Morris. 2021. How shades of truth and age affect responses to covid-19 (mis) information: randomized survey experiment among whatsapp users in uk and brazil. *Humanities and Social Sciences Communications*, 8(1).

Herun Wan, Jiaying Wu, Minnan Luo, Zhi Zeng, and Zhixiong Su. 2025. Truth over tricks: Measuring and mitigating shortcut learning in misinformation detection. *Preprint*, arXiv:2506.02350.

Aimei Yang, Alvin Zhou, Jieun Shin, Ke Huang-Isherwood, Wenlin Liu, Chuqing Dong, Eugene Lee, and Jingyi Sun. 2024. Sharing is caring? how moral foundation frames drive the sharing of corrective messages and misinformation about covid-19 vaccines. *Journal of Computational Social Science*, 7(3):2701–2733.

Xuhui Zhou, Hao Zhu, Leena Mathur, Ruohong Zhang, Haofei Yu, Zhengyang Qi, Louis-Philippe Morency, Yonatan Bisk, Daniel Fried, Graham Neubig, et al. 2023. Sotopia: Interactive evaluation for social intelligence in language agents. *arXiv preprint arXiv:2310.11667*.

A Taxonomy Dimensions and Imputed Beliefs

Beliefs are an important tool for modeling demographic simulations in the context of misinformation. Psychological work shows that several belief dimensions, such as conspiracy beliefs, political ideology, trust in science, etc. can help us predict who is most susceptible to misinformation (Ecker et al., 2022; Munusamy et al., 2024). Rather than using only demographic information as a flat category, adding belief profiles provides a more fine-grained representation of how different groups may perceive information. Recent LLM studies also argue for modeling population-level beliefs and preferences to simulate targeted groups, using belief-like representations to approximate human responses at scale (Namikoshi et al., 2024; Kaiser et al., 2025).

A.1 Taxonomy formation

We create a belief taxonomy of beliefs consisting of seven core dimensions that are associated with misinformation susceptibility grounded in psychological and cognitive science research:
(1) Worldview and Identity Beliefs: individuals interpret information through the lens of social identity and worldview (Kahan, 2017; Van Bavel et al., 2021), **(2) Epistemic Trust Beliefs**: individuals differ systematically in epistemic trust toward institutions and experts (De Coninck et al., 2021; Lewandowsky et al., 2023), **(3) Cognitive style**: individuals vary in cognition, such as reliance on analytic versus intuitive reasoning (Pennycook and Rand, 2019; Ecker et al., 2022), **(4) Conspiracy mentality**: individuals differ in a generalized predisposition to conspiracies (Douglas et al., 2019; De Coninck et al., 2021), **(5) Moral and Value Beliefs**: individuals prioritize different moral values (D'Errico et al., 2022; Yang et al., 2024), **(6) Emotion-Related Beliefs**: individuals vary in emotional responsiveness (heightened emotional arousal, such as anger or fear may amplify belief in misinformation) (Brady et al., 2017; McLoughlin et al., 2024) and **(7) Heuristic Beliefs**: individuals rely to different degrees on shortcuts such as repetition, familiarity, and social endorsement (Lin et al., 2016; Fazio, 2020).

While the above dimensions have traditionally been studied in isolation and within human populations, we unify them into a structured taxonomy designed for computational modeling. This taxonomy enables systematic belief simulation in models by providing a framework for collecting and organizing belief data and providing interpretable axes along which belief priors can be instantiated in models.

A.2 Mapping to Belief Taxonomy.

We map imputed data to the above belief dimensions. We apply exploratory factor analysis on WVS responses to identify latent belief groupings, examining which items naturally cluster together (e.g., trust-related items to Epistemic & Trust Beliefs). Post that, we conduct a manual review to ensure proper alignment with our taxonomy and validate the final question-to-dimension mapping. Overall, we obtain 126 imputed belief questions across dimensions. Table 9 contains the 126 WVS questions mapped to the 7 belief taxonomy dimensions.

B Datasets

B.1 Evaluation Data

For susceptibility evaluation, we use two ground-truth datasets containing human judgments of whether they believe a given claim: (1) PANDORA Dataset from Borah et al. (2025b), containing annotations from 318 participants. Each participant provided judgments on 3 distinct claims, along with demographic information including age, gender, living area, and education. These claims are collected from RumorEval (Gorrell et al., 2019), which consists of true or false rumors covering eight major news events and natural disaster events; (2) MIST dataset from Maertens et al. (2024). From MIST, we use only Study 1 (MIST-1) which includes 409 participants, each providing judgment on the same 100 claims, and demographic information including age, gender, and education. From the two datasets combined, we obtain 13.8K claims for evaluation.

B.2 Belief Data

We collect two complementary belief signals, consisting of individually observed belief judgments and group-level (demographic) belief distributions: **(1) Observed Data** - claims that were directly judged by participants in PANDORA and MIST-1 datasets. These responses capture individual-level belief judgments, reflecting each participant's personal stance rather than demographic group averages. We use two claim judgments as observed be-

Dataset Examples

MIST:
Participant ID: 1
Participant Demographic Group: d1
Claim: The Government Is Knowingly Spreading Disease Through the Airwaves and Food Supply
Participant Choice: fake
Gold Label: fake

PANDORA:
Participant ID: 1
Participant Demographic Group: d1
Claim: News: Updated CDC Covid Numbers Show Only Six Percent of Total US Deaths Actually Dies Solely From Virus.
Participant Choice: true
Gold Label: fake

Figure 8: Dataset Examples - MIST and PANDORA

liefs for each participant (keeping it separate from the evaluation data). This leads to 27.6 claim judgments as observed beliefs.

(2) Imputed Data - inferred from the World Values Survey Wave 7[5] distributions. Imputed data represent demographic belief priors (group–level), inferred from WVS, conditioned solely on demographic attributes. We map these imputed belief items to our belief taxonomy using exploratory factor analysis. This yields 126 imputed belief questions and corresponding demographic distributions. Table 1 shows examples, and Appendix A.2 contains mapping details and all questions.

Fig 11 shows an example from each MIST and PANDORA datasets. Both datasets are similar, except for the size. Furthermore, MIST does not contain the living area information of the participants, therefore, we do not use it for fine-tuning purposes.

C Prompt-Based Conditioning

C.1 Prompt Details

Fig 10 shows the susceptibility accuracies across the belief and demographic configurations for both

[5] https://www.worldvaluessurvey.org/WVSDocumentationWV7.jsp

Prompt

You are a persona grounded by attributes: <demographic group>.

Past beliefs and priors for this persona (for context, do not re-evaluate them): <beliefs>.

When judging a claim, stay consistent with this persona's prior beliefs where reasonable.

Figure 9: Prompt Conditioning Experiments

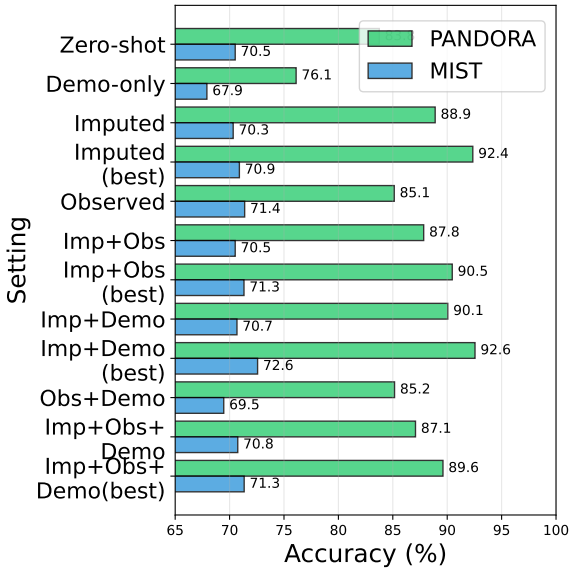

Figure 10: Susceptibility Accuracies across settings and datasets

datasets. For `Imputed` cases, we average across all belief dimensions, and `Imputed (best)` cases, we select the belief dimension that achieves the highest susceptibility accuracy. Findings show that incorporating belief information consistently improves performance over the zero-shot and demographic-only baselines. We next analyze the impact of specific belief types, demographic ablations, datasets, and model performances.

C.2 Belief-Dimension Analysis

Table 4 shows the belief-dimension-wise scores for each dataset and demographic dimension.

Across both datasets, single dimensions consistently outperform the combined prior, and this is especially stark on PANDORA, where all-dimensions collapses to much lower accuracy while the best

single dimensions remain high.

On Prolific, the top-performing dimensions are fairly stable: emotion-related dominates Gender, while moral-and-value-beliefs is strongest and most frequent for Age and Rural/Urban, and heuristic/cognitive beliefs for Education. On MIST, gains from beliefs are smaller but more consistent: winners are spread across - top dimensions like moral-and-value, emotion-related, heuristics, etc., still dominate. In addition, worldview-and-identity also does better in MIST.

C.3 Thematic Analysis

We perform topic clustering on the news claims to identify latent topical grouping and investigate whether LLM susceptibility accuracies vary across demographics within each topic. We perform this on Dataset 2, as it is much larger and diverse than Dataset 1. Using non-negative matrix factorization on claim texts, we identify five dominant topics: (1) global power, population, and future threats; (2) political leaders and historical rankings; (3) science, health, and technology claims; (4) politics, government, and information-control narratives; and (5) public opinion, moral values, and social belief statements. Topic-level analysis shows that susceptibility accuracy is not uniform across demographics. Topics related to political rankings and public opinion (Topics 2 and 5) achieve uniformly high accuracy across all demographic groups, resulting in negligible age, gender, or education gaps. Contrastingly, science/health claims and government-related narratives (Topics 3 and 4) exhibit the largest demographic variation, with older and higher-education groups generally achieving higher accuracy. Gender-based differences are minimal across all topics. Table 5 shows results across some topics and models showing differences across demographics.

These findings suggest that demographic susceptibility in LLMs is not just a broad, identity-driven phenomenon but also contextual and topic-sensitive. This also aligns with interdisciplinary work in psychology and communication suggesting that demographic susceptibility to misinformation is highly context-dependent. Prior studies show that gender alone is a weak predictor of misinformation belief, with analytic thinking and reasoning style playing a much larger role (Marinescu et al., 2018). In contrast, age and education demographics differ more for complex or ambiguous claims, particularly in science and policy domains, consistent with work on cognitive processing, and motivated reasoning (Kahan, 2017). From a cognitive perspective, this pattern supports dual-process theories of reasoning, whereby strong lexical reduce individual differences, while inference-heavy claims amplify them (Evans and Stanovich, 2013).

D Fine-tuning Analysis

D.1 Phase 1 results

For Phase 1, Table 6 shows that Qwen has the lowest KL divergence, followed by Mistral and Llama for both settings - head tuning and full-fine-tuning. With LoRA-ft, the trends remain similar but slightly lower than head-ft. This indicates that Qwen more faithfully captures population-level belief distributions, a critical property for reliable downstream simulation of demographic misinformation susceptibility. Other models also perform competitively, suggesting that most models LLMs can capture belief distributions to a reasonable extent.

D.2 Phase 2 results

Table 7 shows susceptibility accuracies and F1 scores across different setups and models. Across methods, the key finding is that cross-study (MIST-2) generalization improves drastically once we move from 1-phase to 2-phase pipelines. The 1-phase baseline model perform decent, however, the performance drops sharply on MIST-2. For MIST-2, our BAFT approach performs the best, showing that belief adapter + lightweight susceptibility head captures more transferable signals aligned with out-of-domain data.

Comparing the 2-phase variants, we find that BAFT remains the strongest on MIST-2, however, LoRA-ft slightly improves on PANDORA+MIST-1 but gives weaker transfer. This may be due to mild overfitting when more parameters are updated. Across models, Qwen consistently performs the best for 2-phase settings. Overall, the 2-phase training approach is helpful improving robustness and also adapting to out-of-domain data.

D.3 Shortcut Reliance Experiments

Table 8 shows that shortcut reliance is low, especially under head tuning. We measure flip-rate: the fraction of evaluation examples for which the model's predicted label changes when a shortcut feature is removed or perturbed, and Probability Delta: the average absolute change in the model's

Dataset	Dimension	Gender	Age	Education	Rural/Urban
Prolific-dataset	worldview_and_identity	89.20625	90.19875	69.54372184	90.71375
	epistmetic_and_trust_beliefs	85.515	89.2575	84.96031741	89.0475
	cognitive_style_beliefs	89.0875	88.69	84.76190471	87.6575
	conspiracy_mindset	83.25375	89.0475	85	86.985
	moral_and_value_beliefs	90.15875	90.19875	84.36507924	90.72375
	emotion_related	88.3811111	89.545	85.31746026	89.08625
	heurisitics	86.78625	90.00125	86.38888875	90.555
	alldimensions	73.74625	88.96875	85.95238095	89.26
MIST-dataset	worldview_and_identity	71.66125	72.49375	71.65924606	
	epistmetic_and_trust_beliefs	71.68	71.5875	71.47649656	
	cognitive_style_beliefs	71.27458333	71.91125	71.35596075	
	conspiracy_mindset	71.1875	71.64	71.26027066	
	moral_and_value_beliefs	71.02625	71.99	70.52219001	
	emotion_related	71.2175	71.43	70.75732404	
	heurisitics	71.85875	72.09125	71.80576158	
	alldimensions	71.50958333	71.74125	71.18524694	

Table 4: Belief Dimension Analysis

predicted probability for the original class after the shortcut is removed. On PROLIFIC+MIST-1, head-tuned models (llama, qwen and mistral) show near-zero flip rates, indicating that removing the shortcut signal has little effect on predicted labels. Full fine-tuning also tends to have small flip rates and probability deltas in-domain, but they are more frequently non-zero (e.g., llama/mistral flip more than head tuning), consistent with the idea that updating more parameters can slightly increase sensitivity to spurious cues.

On MIST-2, the most notable shortcut sensitivity appears for qwen under full fine-tuning (flip rate 0.1098, larger prob delta), whereas the corresponding head-tuned Qwen model shows 0 flip rate. This pattern supports the our previous findings that two-phase-head-only training is more robust: it preserves performance while reducing reliance on shortcut features that do not transfer.

D.4 Training Details

D.4.1 2-phase Head-FT details

Phase 1 trains a belief adapter to predict demographic-conditioned WVS response distributions. The model is a frozen encoder with a trainable linear head mapping the last-token hidden state to 10 logits, followed by softmax. Training uses AdamW (lr 5×10^{-4}, batch size 16, 2 epochs) and optimizes a scale-aware KL divergence loss, computed only over the valid bins 1..K for each example. For evaluation, we test generalization to unseen belief items, reporting KL distributional fit and majority-category accuracy.

Belief Modeling Prompt

You are modeling belief distributions for a single demographic group.

Demographic : d1

Question: q1

There are k possible response options, numbered 1 through k. Predict the probability for each option.

Figure 11: Prompt for Phase-1 Belief Modeling

Phase 2 trains a susceptibility prediction head on top of a frozen belief adapter learned in Phase 1. The model combines a frozen base encoder and frozen belief head to produce both a semantic representation of the input and a demographic-conditioned belief probability vector. These two are concatenated and fed into a lightweight, trainable classification head. We format the inputs as instruction-following prompts that include a demographic persona and up to two belief examples. During training, only the susceptibility head parameters are updated using cross-entropy loss, while the base model and belief adapter remain frozen. The model is optimized with AdamW (learning rate (5×10^{-4}), batch size 8) for 2 epochs, with

Demographic	Model	Topic/Theme	Effect and Notes
Age	Llama-3-8B	New study & ideology (T1)	+2.8 ppts (60+ <30). Older: 0.984 vs younger: 0.956.
	Llama-3-8B	Global threats & population (T2)	+2.0 ppts. Older: 0.961 vs younger: 0.941.
	Llama-3-8B	Marijuana/new-study (T4)	+1.7 ppts. 60+ group: 0.947 vs <30: 0.930.
	Qwen-2.5-14B	Marijuana/new-study (T4)	+2.5 ppts. 60+ group: 0.917 vs <30: 0.892.
	Mistral-7B	All topics	<0.5 ppts. Age differences negligible.
Education	Llama-3-8B	Sleep/blue-light health (T2)	+8.4 ppts. Completed: 0.918 vs not: 0.833 (six cases).
	Qwen-2.5-14B	Politics/elections (T3)	+16.7 ppts. Completed perfect; not-completed: 0.833 (six cases).
	Llama-3-8B	Marijuana/new-study (T4)	+0.9 ppts. Difference: 0.962 vs 0.953.
	Mistral-7B	Marijuana/new-study (T4)	+0.2 ppts. Effect negligible.
	All models	Other topics	0.0–0.2 ppts. No meaningful education effect.
Gender	Llama-3-8B	Global influence/future threats (T0)	≤ 2 ppts ($\lvert F - M \rvert$). Nearly identical accuracies.
	Llama-3-8B	Science & health claims (T2)	≤ 2 ppts. No consistent female–male separation.
	Llama-3-8B	Public opinion/social beliefs (T4)	≈ 0 ppts. Near-ceiling for both genders.
	Mistral-7B	Global influence/future threats (T0)	≈ 0 ppts. Accuracies overlap almost exactly.
	Mistral-7B	Marijuana/public opinion (T4)	≤ 1 ppt. No directional trend.
	Qwen-2.5-14B	All topics	0 ppts. Identical accuracies across all topics.

Table 5: Thematic Analysis across Demographic

Model	Head Tuning	Full Fine-tuning
Qwen	0.0511	0.0732
Mistral	0.087	0.107
Llama	0.287	0.487

Table 6: KL divergence comparison across different models and tuning methods.

Model	PANDORA+MIST-1		MIST-2	
	Acc	F1	Acc	F1
Baseline (1-phase)				
llama	0.750	0.751	0.680	0.673
qwen	0.760	0.761	0.720	0.712
mistral	0.740	0.761	0.650	0.653
LoRA-FT (2-phase)				
llama	0.809	0.809	0.876	0.861
qwen	0.800	0.800	0.893	0.892
mistral	0.800	0.790	0.887	0.867
BAFT				
llama	0.793	0.793	0.884	0.896
qwen	0.792	0.791	0.924	0.906
mistral	0.792	0.792	0.884	0.923

Table 7: Fine-tuning performances

reproducible initialization via fixed random seeds. We finally evaluate at each epoch using accuracy and macro-F1.

D.4.2 2-phase LoRA-FT details

Here, we use Low-Rank Adaptation (LoRA) for parameter-efficient fine-tuning in both phases. In Phase 1, LoRA enables the model to adapt its representations to capture demographic belief distributions. In Phase 2, LoRA is again applied to the base model during susceptibility training, allowing the model to adjust task-relevant representations while keeping the number of trainable parameters small.

E Significance Tests Across Experiments

Prompt-based Conditioning (Susceptibility Accuracy): Conditioning LLMs with belief information improves demographic susceptibility simulation compared to zero-shot and demo-only prompts. We averaged our findings across 3 individual runs, and improvements of belief settings are statistically significant (paired t-test, $p < 0.05$ for both PANDORA and MIST-1).

Belief Dimension Analysis. Certain belief dimensions are more predictive of susceptibility than others. We use repeated-measures ANOVA and performance variation vary significantly by belief dimensions, more for PANDORA, with a few non-significant variation for MIST (congitive, and alldimensions).

Post-Training Adaptation In PANDORA+MIST-1, accuracy differences between baseline and the 2-phase methods are modest (e.g., 0.75 vs. 0.793 for llama), and z-tests show that both BAFT and LoRA fine-tuning significantly outperforms the baseline (p=0.02–0.03). On MIST-2 the gains are

Method	Model	PANDORA+MIST-1			MIST-2		
		Flip Rate	Prob Delta	Acc Drop	Flip Rate	Prob Delta	Acc Drop
Full-FT	Llama	0.0239	0.0148	0.0049	0.0000	0.0000	0.0000
	Qwen	0.0173	0.0476	0.0012	0.1098	0.0700	0.1027
	Mistral	0.0323	0.0024	0.0000	0.0000	0.0129	0.0000
Head Tuning	Llama	0.0000	0.0125	0.0000	0.0000	0.0129	0.0000
	Qwen	0.0080	0.0261	0.0012	0.0000	0.0427	0.0000
	Mistral	0.0000	0.0071	0.0000	0.0000	0.0069	0.0000

Table 8: Shortcut reliance comparison between Full Fine-tuning and Head Tuning across different models on PANDORA+MIST-1 and MIST-2 datasets.

pronounced: for llama the baseline accuracy is 0.680 while BAFT reaches 0.884 and LoRA-FT 0.876. Two-proportion z-tests show that both BAFT and LoRA-FT significantly outperform the baseline ($p<0.001$).

F Model Choices, Implementation Details and Computational Resources

We conduct experiments using instruction-tuned LLMs including LLaMA, Qwen, and Mistral, implemented with the Hugging Face Transformers and PEFT libraries. Our framework adopts a modular design with explicit belief heads and susceptibility heads, and employs LoRA-based fine-tuning where base-model adaptation is required, while relying on lightweight head training in BAFT to decouple belief modeling from susceptibility prediction. All experiments are run on NVIDIA A40 GPUs with a maximum sequence length of 1024 tokens and fixed random seeds to ensure reproducibility and computational efficiency.

G Reproducibility

We open-source our codes and data, which are uploaded to the submission system. This would help future work to reproduce our results and explore BeliefSim: demographic-aware misinformation susceptibility simulation in LLMs using belief-priors, and BAFT.

Table 9: World Values Survey Questions mapped to Misinformation Taxonomy Dimensions

Dimension	Q ID	Question
WorldView and Identity Beliefs		
	Q6	Importance of religion (1=Very important to 4=Not at all important)
	Q4	Importance of politics (1=Very important to 4=Not at all important)
	Q171	Frequency of religious service attendance (1=More than once/week to 7=Never)
	Q173	Religious self-identification (1=Religious person, 2=Not religious, 3=Atheist)
	Q240	Political left-right scale (1-10)
	Q254	National pride (1=Very proud to 4=Not at all proud, 5=Not from US)
	Q255-Q259	Closeness to village/region/country/continent/world (1=Very close to 4=Not close at all)
	Q19, Q23	Unwanted neighbors: different race, different religion (1=Don't want, 2=Want)
	Q235-Q239	Views on government styles: strong leader, technocracy, army rule, democracy, religious law (1=Very good to 4=Very bad)
	Q241-Q249	Essential characteristics of democracy: tax/subsidize, religious law interpretation, free elections, unemployment aid, army takeover, civil rights, income equality, obedience, women's rights (1-10 scale)
Epistemic and Trust Beliefs		
	Q57	General trust in people (1=Most can be trusted, 2=Need to be careful)
	Q58-Q63	Trust in family, neighborhood, people you know, first-time meetings, other religions, other nationalities (1=Trust completely to 4=Don't trust at all)
	Q69, Q71, Q75	Confidence in police, government, universities (1=A great deal to 4=None at all)
	Q158-Q163	Views on science and technology: health/comfort, opportunities, faith vs science, moral breakdown, daily relevance, world better/worse (1-10 scale)
Cognitive Style Beliefs		
	Q152-Q153	Priority ranking: economic growth, defense forces, participation in decisions, beautification (select most and next most important)
	Q154-Q155	Priority ranking: order, participation, fighting prices, free speech (select most and next most important)
	Q156-Q157	Priority ranking: stable economy, humane society, ideas over money, fight crime (select most and next most important)
	Q176	Moral uncertainty: trouble deciding which moral rules are right (1-10 scale)
Conspiracy Mindset		
	Q112	Perceived corruption level in country (1=No corruption to 10=Abundant corruption)
	Q113-Q117	Corruption beliefs about state authorities, business executives, local authorities, civil service providers, journalists/media (1=None to 4=All of them)
	Q118	Frequency of bribery needed for services (1=Never to 4=Always)
Moral and Value Beliefs		
	Q176	Moral uncertainty (1-10 scale)
	Q177-Q195	Justifiability of: claiming unentitled benefits, fare evasion, stealing, tax cheating, bribery, homosexuality, prostitution, abortion, divorce, premarital sex, suicide, euthanasia, wife beating, child beating, violence, terrorism, casual sex, political violence, death penalty (1=Never justifiable to 10=Always justifiable)
Emotion Related		
	Q44-Q45	Views on technology development and respect for authority (1=Good, 2=Don't mind, 3=Bad)
	Q46	Happiness level (1=Very happy to 4=Not at all happy)
	Q47	Health status (1=Very good to 5=Very poor)
	Q48	Freedom of choice and control over life (1=No choice to 10=Great deal of choice)
	Q49-Q50	Life satisfaction and financial satisfaction (1=Completely dissatisfied to 10=Completely satisfied)
	Q52	Felt unsafe from crime in last 12 months (1=Often to 4=Never)
	Q131	General security feeling (1=Very secure to 4=Not at all secure)
	Q146-Q148	Worry about war, terrorist attack, civil war (1=Very much to 4=Not at all)
Heuristic		
	Q94-Q104	Active membership in: church/religious, sports/recreation, arts/music/education, labor union, political party, environmental, professional, humanitarian/charity, consumer, self-help/mutual aid, women's groups (1=Inactive, 2=Active, 3=Don't belong)
	Q201-Q208	Media/communication usage: newspaper, TV news, radio news, mobile phone, email, internet, social media, talking with friends (1=Daily to 5=Never)
	Q18-Q26	Unwanted neighbors: drug addicts, different race, AIDS, immigrants, homosexuals, different religion, heavy drinkers, unmarried couples, different language (1=Don't want, 2=Want)
	Q29-Q31	Gender attitudes: men as better political leaders, university more important for boys, men as better business executives (1=Strongly agree to 4=Strongly disagree)
	Q33-Q35	Job/gender attitudes: men's right to jobs when scarce, natives over immigrants, woman earning more causes problems (1=Strongly agree to 5=Strongly disagree)