

Synthetic Contact with AI Reduces Cross-Partisan Animosity

Benjamin Lira Luttges^{1,*}, Noah Castelo², Stefano Puntoni¹, Olivier Toubia³

¹The Wharton School, University of Pennsylvania. ²Alberta School of Business, University of Alberta. ³Columbia Business School, Columbia University.

*Corresponding author: blira@upenn.edu.

Most recent version: lirabenjamin.github.io/synthetic-contact

Abstract.

Americans' warmth toward members of the opposing political party has fallen sharply over the past three decades—yet meaningful cross-partisan contact remains scarce, in part because people actively avoid it. Across five preregistered studies (total $N = 3,960$ U.S. partisans), we test whether brief conversations with AI chatbots representing the political outgroup can substitute for the contact people shun. Synthetic contact first lowers the barrier to entry: partisans would endure almost twice as long contemplating their own mortality to avoid a human outgroup partner as an AI one. These conversations then correct the misperceptions that fuel division. At baseline, Democrats placed Republicans more than a standard deviation past their actual position on environmental consumption attitudes—enough to flip the average Republican from supportive to opposed—and a single ten-minute conversation with an outgroup chatbot corrected those beliefs and warmed affect in a within-person study of both parties. A three-arm experiment ruled out pure engagement and sociality as drivers. Synthetic contact also moved behavior, in a sample of both parties and on a more affectively charged issue: participants who spoke with an outgroup bot about immigration were six percentage points more likely than controls to choose to have a real conversation with a partisan from the other side. A final study tested whether these gains last: the warmth effect replicated immediately in a new sample; most of it faded within a week, with a small residual concentrated among the most extreme partisans. Analyzing conversation content showed that information, more than friendliness, distinguishes outgroup bots from control chatbots. Together, these findings establish synthetic contact as a scalable, behaviorally consequential, and—unlike face-to-face contact—widely acceptable form of cross-partisan engagement.

Americans increasingly view political opponents with suspicion and dislike. Warmth toward the political outgroup has fallen steadily over the past three decades.¹ This animosity is only one face of polarization. Beyond cold feelings toward the outgroup, partisans also hold systematic *misperceptions*—inaccurate beliefs about what the other side actually thinks.^{2–5} Together these trends erode trust in institutions and make cross-partisan cooperation harder in domains from public health to consumer markets.^{1,6} Warming partisans toward their political outgroup remains a central challenge for social scientists and practitioners alike.

Intergroup contact offers one of the most robust solutions to this animosity. Since Allport,⁷ decades of research have established that positive interactions between members of different groups reduce prejudice and increase mutual understanding,^{8,9} though the strength of this evidence has been debated.¹⁰ Researchers still debate why contact works. A classic meta-analysis¹¹ of over 500 effects points to three mediators: contact builds knowledge of the outgroup, lowers intergroup anxiety, and increases empathy. Notably, they find that the two affective routes, anxiety and em-

pathy, outweigh the effects of gaining knowledge.

Recent work extends this framework to politics: bringing Democrats and Republicans together for cross-party discussion reduces affective polarization.¹² Explicitly debating partisan disagreements, however, does not reliably help—Santoro and Broockman¹³ found that outpartisans who discussed a shared experience (their “perfect day”) grew less polarized, whereas those who debated their disagreements did not.

Contact between political outgroups faces an important limitation: most partisans actively avoid interacting with the other side. Using anonymized smartphone location data from millions of Americans, Chen and Rohla¹⁴ found that Thanksgiving dinners shared across partisan lines ended earlier than same-party dinners, and liberals and conservatives alike forgo money to avoid hearing opposing opinions.¹⁵ This avoidance is partly miscalibrated—people overestimate how unpleasant engaging the other side will feel.¹⁶ Face-to-face interventions also demand logistics and professional facilitation that limit their scale.

Digital platforms opened new possibilities, but online political interaction frequently amplifies conflict

Table 1. Overview of the five studies. Studies are numbered in the order they appear in the main text. All effect sizes are Cohen's d unless noted; OR denotes an odds ratio. The Pre-registration column lists AsPredicted numbers.

No.	Study	N	Design	Key comparison	Primary DV	Key result	Pre-registration
S1	Aversion	608	Between-subjects, 2 arms (both parties)	AI bot vs. live human outgroup partner	Mortality-reflection minutes accepted at indifference	Bot 5.06 vs. human 9.65 min ($d = 0.34$, $p < .001$)	#286,575
S2	Within-person contact	500	Within-person pre-post, 1 arm (both parties)	Pre- vs. post-conversation with outgroup bot	Outgroup warmth, 0–100 thermometer	+4.3 pts ($d = 0.37$, $p < .001$)	#264,402
S3	Three-arm experiment	679	Between-subjects, 3 arms (Democrats)	Outgroup bot vs. cats/dogs chat vs. Space Invaders	Outgroup warmth, 0–100 thermometer	$d = 0.58 / 0.59$ vs. the two controls ($p < .001$)	#276,530
S4	Behavioral choice	1069	Between-subjects, 2 arms (both parties)	Outgroup bot vs. cats/dogs chat	Costly choice: real outgroup conversation vs. mortality reflection	OR = 1.33, $p = 0.025$	#287,002
S5	Longitudinal	1104	Between-subjects, 2 arms (Democrats); 1-week follow-up	Outgroup bot vs. cats/dogs chat	Outgroup warmth, 0–100 thermometer, at 1 week	Immediate $d = 0.46$; 1-week $d = 0.05$ (n.s., $p = .16$); pooled S3+S5 1-week $d = 0.12$ ($p = .02$)	#288,592

rather than reducing it.¹⁷ Even interventions explicitly designed to expose users to opposing views can fall flat or backfire: a field experiment on Twitter found that replacing participants' feeds with opposing-leaning feeds boosted engagement without improving self-reported understanding of the other side.¹⁸ Another Twitter intervention that recommended opposing-ideology accounts to follow *reduced* users' willingness to converse with an outgroup member.¹⁹

Large language models might offer a way to warm cross-party relationships while avoiding some of the limitations of human contact and interventions on digital platforms. A chatbot prompted to represent the political outgroup stands in for a real interlocutor, available on demand. We use the term *synthetic contact* for conversations of this kind. Because no real person sits on the other side, partisans can engage the views they avoid without the threat that makes contact aversive in person and backfire online. The possibility that AI might improve intergroup relations has been raised conceptually: Hermann et al.⁶ propose that AI agents could reduce prejudice if *engineered* to be counter-stereotypical, deliberately built to contradict the outgroup's negative stereotype. We test a more minimal version: rather than engineering the bot to defy stereotypes, we prompt it to represent a *typical* outgroup member. If partisan stereotypes are exaggerated, then even an accurate portrayal should contradict them, and we test whether that alone is enough to reduce prejudice.

Synthetic contact is also newly feasible, and several

lines of evidence suggest it could work. Large language models can generate realistic representations of diverse political viewpoints,^{20,21} and participants often engage with them as they would a human conversation partner;^{22,23} meta-analytic evidence already shows that digitally mediated intergroup contact reduces prejudice.²⁴ Recent work also shows that people perceive AI sources as less biased, more informative, and less persuasively intended than human sources, which increases receptiveness to opposing views.²⁵ People are not only open to LLMs but persuaded by them: brief LLM dialogue can durably reduce conspiracy beliefs,²⁶ shift candidate preferences in real elections,²⁷ and outpersuade humans in head-to-head debate.²⁸

But persuasive power cuts both ways. Large language models trained on internet text exhibit systematic political leanings,^{29,30} so a bot prompted to represent the political outgroup may portray its positions in caricatured rather than calibrated form. If so, its persuasive force would entrench partisan misperceptions rather than correct them—reinforcing the very stereotypes that contact is meant to dissolve.

Here, we test across five preregistered studies whether brief synthetic contact is acceptable, whether it corrects misperceptions and warms cross-partisan affect, and whether it moves a costly behavioral choice (Table 1 summarizes the design, sample, and primary outcome of each study). We begin with an incentive-compatible aversion experiment that trades off cross-partisan conversation against an aversive mortality-reflection task, quantifying how much more willing

partisans are to engage with an AI outgroup partner than with a human one. A within-person study with Democrats and Republicans then tests whether a single ten-minute conversation with an outgroup-representing chatbot objectively corrects misperceptions and warms cross-partisan affect. A three-arm experiment compares synthetic contact against an active chat control (an equally polite AI conversation about an apolitical topic) and a non-social game control, isolating the causal contribution of outgroup-specific content. A two-arm behavioral experiment asks whether synthetic contact shifts behavior, not just attitudes: whether participants who just spoke with an outgroup bot are more willing to enter a real cross-partisan conversation than control participants. A longitudinal experiment tests whether brief synthetic contact produces durable attitude change one week later. Finally, we code the content of every bot conversation to ask which mechanism sets outgroup bots apart from controls—stereotype-disconfirming information (a cognitive route) or warmth and empathy (an affective route).¹¹

Across these studies we deliberately varied the conversation topic: the misperception and warmth studies used environmental consumption attitudes, where a validated scale (the GREEN measure,³¹) quantifies misperception item by item, while the two incentive-compatible studies with real behavioral stakes used immigration, a more affectively charged identity issue and a harder test for contact. Convergent effects across both issues suggest that the effect is not specific to any single topic.

An AI partner halves the aversion to cross-partisan conversation

Before synthetic contact can reduce polarization, partisans have to agree to it. In Study 1 (Preregistered, AsPredicted #286,575, $N = 608$), we measured how much of an aversive experience they would endure to avoid an outgroup conversation, and whether an AI partner lowers that price. Participants made repeated forced choices between three minutes of conversation about immigration with a member of the political outgroup and an adjustable duration X of reflection on one's own mortality (a deliberately aversive task). The duration X was raised whenever a participant picked the conversation and lowered whenever they picked mortality reflection, homing in on the duration at which each participant was indifferent between the two. Only the conversation partner differed between conditions: in the *human* condition, participants weighed mortality reflection against three minutes with a live participant from the opposing party, whereas in the *bot* condi-




Fig. 1. Aversion to outgroup conversation is far lower for an AI than a human partner. In an incentive-compatible 1-up-1-down staircase, participants traded a fixed three-minute outgroup conversation about immigration against an adjustable duration of mortality reflection. Bars show the mortality-reflection duration that felt equally aversive to the three-minute chat: 9.65 minutes with a human outgroup partner, but only 5.06 minutes with an AI partner. Error bars are ± 1 SE.

tion they weighed it against three minutes with an AI trained to represent a typical outgroup member. The staircase was incentive-compatible: whichever option a participant chose on the final trial was the option they actually had to complete—a real cross-partisan conversation or a real mortality-reflection exercise—so every choice traded real conversation time against real mortality-reflection time.

Aversion to outgroup conversations was substantially lower when the interlocutor was an AI. The pre-registered primary model regressed the mortality threshold on condition: thresholds were 4.59 minutes lower on average in the bot condition than in the human condition ($\beta = 4.59$, 95% CI [2.44, 6.74], $t = 4.19$, $p < .001$, $d = 0.34$; Figure 1). The typical participant matched with a human equated three minutes of outgroup conversation with 9.65 minutes of contemplating their own death—roughly double the exchange rate accepted with a bot (5.06 minutes). The advantage held for Democrats and Republicans alike, across the extremity range, and across robustness specifications (Appendix S2.5).

Why do partisans refuse a conversation with a member of the other party? A long literature suggests one likely answer: *they misperceive what the conversation would be like*. Partisans systematically hold exaggerated views of the outgroup's positions, its composition, and its attitudes toward their own side,²⁻⁵ and overestimate how aversive exposure to opposing views will be.¹⁶ This suggests that an intervention that corrects




Fig. 2. A single ten-minute conversation corrects misperceptions of the outgroup and warms attitudes toward it. Color denotes the *participant* party (blue = Democrats, red = Republicans). **(A) Baseline misperceptions:** each party's estimate of the outgroup, the outgroup's actual attitudes, and the bot's position on the six environmental items. Democrats sharply underestimate Republican environmental concern, and the bot sits closer to real Republicans than Democrats' estimate does. **(B) Belief accuracy** and **(C) outgroup warmth**, pre- and post-chat, by party; both improved after the conversation, and the size of belief correction predicts the size of warmth gain. Error bars are ± 1 SE.

such misperceptions may both warm cross-partisan affect and—over time—erode the very aversion that limits face-to-face contact. We now turn to whether brief AI conversations can correct these misperceptions, durably warm cross-partisan affect, and change behavior to make human intergroup contact more likely.

Partisans misperceive each other

In Study 2 (Preregistered, AsPredicted #264,402, $N = 500$), we asked partisans (248 Democrats, 252 Republicans) to report their own attitudes on six environmental items (the GREEN consumption-values scale, 1–5,³¹) and to estimate how a typical outgroup member would respond to the same items. Partisans held large, asymmetric misperceptions of each other's environmental attitudes (Figure 2A). Democrats sharply underestimated Republican attitudes toward environmental consumption ($d = 1.49$, $SE = 0.10$, $p < .001$)—an error large enough to misclassify the average Republican as opposed to green products rather than moderately supportive of them. Republicans were more calibrated about Democrats ($d = 0.24$, $SE = 0.09$, $p = .01$). The two misperceptions differ: the belief–reality gap is larger in the Democrat-judging-Republican direction than in the reverse (role $\times$ target interaction, $t(996) = 11.5$, $p < .001$).

The bots themselves were imperfect guides: both held more extreme views than the partisans they represented. To recover the bots' own attitudes, we presented each party-conditioned bot with the same six GREEN items and computed its position on each item as the

expected response value, weighting the integers 1–5 by the model's token probabilities at the first response position.²⁹ The Republican-representing bot scored 3.084, below the Republican mean (3.728) by $d = 0.61$; the Democrat-representing bot scored 4.577, above the Democratic mean (4.149) by $d = 0.56$. Which direction the bot erred in mattered more than that it erred at all, because what helps a learner is a guide closer to the truth than their own starting beliefs. For Democrats, who began badly miscalibrated about Republicans, the Republican bot was a far more accurate guide than their own beliefs, so leaning toward it moved them toward the truth. For Republicans, who began less misaligned about the outgroup than Democrats, the bot offered less to correct—though they still became more accurate (Fig. 2B). This asymmetry anticipates the accuracy gains reported next: the bot had much more to teach Democrats than Republicans.

Brief synthetic contact corrects misperceptions and warms cross-partisan affect

The same 500 participants then conversed for ten minutes with a chatbot prompted to represent their political outgroup, with each participant instructed to learn how the outgroup thinks about environmental policy. The chatbot (GPT-4o) received a brief system prompt instructing it to answer as a typical member of the participant's outgroup, specifying the outgroup's party identity but no scripted positions (full prompts in Appendix S9). After the conversation, they re-estimated outgroup attitudes on the same six items and re-rated

their warmth toward the outgroup on a 0–100 thermometer.

A single ten-minute conversation corrected baseline misperceptions and warmed cross-partisan affect (Figure 2B–C). Belief accuracy improved by 0.39 points on the 5-point scale ($p < .001$, $d = 0.46$). Outgroup warmth rose by 4.3 thermometer points ($p < .001$, $d = 0.37$)—equivalent to reversing eight years of the rising partisan animosity documented in national surveys since the 1970s.¹ Both parties gained, but the gains were graded by baseline misperception. Democrats, who started with larger misperceptions, gained more in accuracy ($d_{D \rightarrow R} = 0.71$) than Republicans ($d_{R \rightarrow D} = 0.22$; interaction $b = -0.46$, $p < .001$). The same gradient held for warmth (Democrats $d = 0.46$; Republicans $d = 0.29$).

The size of the accuracy gain tracked the size of the warmth gain. Controlling for both baseline accuracy and baseline warmth, post-chat belief accuracy strongly predicted post-chat outgroup warmth ($b = 2.47$, 95% CI [1.32, 3.60], $p < .001$): participants whose beliefs about the outgroup moved more during the conversation also warmed up more, the pattern a cognitive account predicts.

Participants' own perceptions pointed the same way. With both ratings entered together, the bot's perceived informativeness was the only significant predictor of who warmed most ($b = 1.93$, $p < .001$); perceived empathy was not ($b = 1.01$, $p = .06$), though the two coefficients did not differ reliably from each other ($\Delta b = 0.92$, $p = .29$). The conversations themselves carried the same signal: those delivering more stereotype-disconfirming substance produced the largest belief corrections (Appendix S3.7).

Synthetic contact outperforms two active controls

In Study 3 (Preregistered, AsPredicted #276,530, $N = 679$), we randomly assigned Democrats to one of three conditions: synthetic contact (a ten-minute conversation with a chatbot prompted to represent a Republican, in which participants posed as policy researchers gauging how Republicans think about environmental policy), a chat control (the same interface, but a bot given no political identity and prompted to debate cats vs. dogs), or a game control (Space Invaders). The two controls isolate different confounds: the chat control holds the social, conversational experience constant while stripping outgroup content, reducing the plausibility of a pure-sociality account, whereas the game control removes conversation entirely, ruling out generic engagement or arousal.

Synthetic contact increased outgroup warmth rela-




Fig. 3. Synthetic contact raised outgroup warmth above both controls. Mean outgroup warmth (0–100 thermometer) by condition in the three-arm experiment (Study 3). Error bars are ± 1 SE.

tive to both control conditions (Fig. 3). Participants who conversed with the AI outgroup representative rated Republicans more warmly ($M = 29.2$, $SD = 22.3$) than those in the cats-and-dogs chat condition ($M = 17.0$, $SD = 19.8$; $d = 0.58$, $p < .001$) and those in the Space Invaders condition ($M = 17.1$, $SD = 18.5$; $d = 0.59$, $p < .001$). The two control conditions did not differ from each other ($d = -0.01$, $p = .94$), confirming that the effect is specific to outgroup-relevant conversation rather than to the experience of chatting with an AI or engaging in an unrelated task.

Synthetic contact moves a costly behavioral choice

In Study 4 (Preregistered, AsPredicted #287,002, $N = 1069$), we randomly assigned Democrats and Republicans to a five-minute chat with an outgroup bot (discussing how the outgroup thinks about immigration) or the cats-and-dogs control, then offered an incentive-compatible binary choice: a three-minute conversation with a real member of their political outgroup, or three minutes on an aversive mortality-reflection task. Whichever option a participant chose, they actually completed.

Synthetic contact shifted a costly behavioral choice toward cross-partisan contact. 61% of participants in the cats-and-dogs control chose the outgroup conversation, compared with 67% in the outgroup-bot condition. A preregistered logistic regression with party and mean-centered political extremity as covariates yielded an odds ratio of 1.33 (95% CI [1.04, 1.71], $p = 0.025$);




Fig. 4. After synthetic contact, more partisans choose a real cross-partisan conversation ($N = 1069$). Pooled choice share (left) and by party (Democrats, center; Republicans, right). The outcome is the share of participants choosing a three-minute conversation with a real outgroup member over a three-minute mortality reflection. Error bars are ± 1 SE.

a Wilcoxon rank-sum robustness test reached the same conclusion ($W = 133,340$, $p = 0.025$). Republicans showed a significant shift ($OR = 1.55$, $p = 0.020$), and Democrats a directional but non-significant one ($OR = 1.17$, $p = 0.369$); the two parties did not significantly differ (treatment $\times$ party interaction $OR = 1.33$, 95% CI [0.80, 2.20], $p = 0.271$). We found no statistically significant evidence of moderation by partisan strength (treatment $\times$ mean-centered extremity interaction $OR = 1.00$, $p = 0.754$; median-split $OR_{\text{moderate}} = 1.51$ vs. $OR_{\text{strong}} = 1.19$).

Synthetic contact does not just shift self-reported warmth: it moves a subsequent costly choice that leads to an interaction with a real outgroup member. The effect also replicates at half the dosage of the earlier warmth experiments—participants here chatted for five minutes rather than the ten used in the three-arm and within-person studies—suggesting even five minutes suffice to shift behavior.

Most of the warmth effect fades within a week; a small residual concentrates among extreme partisans

In Study 5 (Preregistered, AsPredicted #288,592, $N = 1104$), we randomly assigned self-identified Democrats to a five-minute chat with either the outgroup bot (discussing how Republicans think about environmental policy) or the cats-and-dogs control bot, then recontacted them one week later for a follow-up measure of outgroup warmth. Relative to attitudes built over years and presumably reinforced daily by media exposure, a five-minute conversation is a light intervention,

so we expect substantial decay, consistent with other brief depolarization treatments.^{13,32} The goal of this follow-up is therefore to replicate the immediate effect at larger scale and bound how much of it survives a week. Immediately after the chat, the outgroup bot raised post-chat warmth above the cats-and-dogs control by $b = 9.75$ points ($t(1101) = 10.30$, $p < .001$; $d = 0.46$), replicating the immediate effect in a sample 1.6 times the size of the three-arm experiment.

One week later, 924 of 1104 participants returned (83.7%), with no differential attrition by condition (Appendix S6.4). Most of the immediate effect had faded. The one-week warmth distribution is strongly floor-bunched (Appendix S6.5), which violates the assumptions of the preregistered mean-difference ANCOVA; that model returns a small, non-significant residual ($b = 1.03$, 95% CI $[-0.41, 2.47]$, $d = 0.05$, $p = .16$). A rank-based ANCOVA—a small deviation from the preregistered Wilcoxon rank-sum robustness test that simply adds baseline control—shows the effect declining from $d_{\text{rank}} = 0.39$ immediately after the chat to $d_{\text{rank}} = 0.08$ one week later, about 21% surviving (a 2.36-percentile-point shift, 95% CI $[0.59, 4.12]$, $p = .009$). Re-fitting the paper's other warmth contrasts with the same rank estimator leaves their conclusions unchanged (Appendix S6.5, Table S20).

The surviving residual warmth is concentrated among the most extreme partisans (high-extremity one-week effect: 3.49 percentile points, $p = .007$); however, the preregistered extremity interaction is not significant ($b = 0.07$, $p = .11$).

We also recontacted participants from the three-arm




Fig. 5. Synthetic contact's effect on outgroup warmth is large immediately, mostly gone within a week, and larger among extreme partisans. Outgroup warmth immediately after the chat and one week later, by condition, pooled across the three-arm follow-up and the longitudinal study ($N = 1467$), for (A) all participants and (B) the more extreme half. Points are study-adjusted means ± 1 SE; annotations give the pooled effect (Cohen's d) of the outgroup bot versus control at each timepoint, so the plotted gap equals the annotated effect (stars: * $p < .05$, *** $p < .001$; Appendix S6.6).

experiment (Study 3) one week later ($N = 543$), and its residual effect matches the longitudinal study's almost exactly (Cohen's $d = 0.15$ vs. $d = 0.11$; the effect does not differ between the two studies, condition $\times$ study $p = .73$). Pooling both follow-ups in an exploratory individual-participant analysis ($N = 1467$), the one-week residual is significantly positive ($d = 0.12$, 95% CI [0.02, 0.23], $p = .02$) and concentrated among the more extreme half of partisans ($d = 0.26$ [0.11, 0.40], $p < .001$; Appendix S6.6). A complementary Bayesian analysis with weakly informative priors places a 0.99 posterior probability on a positive one-week effect (Appendix S6.6). A single brief conversation thus leaves a faint but consistent week-later trace, strongest among the most extreme partisans.

This decay parallels other brief depolarization interventions: the effects of single cross-partisan conversations had faded by a three-month follow-up,¹³ and brief online treatments decayed substantially within two weeks.³²

Outgroup bots differ from control bots most in what they say, not how warmly they say it

To characterize what the outgroup bots said to participants, and how they said it, we audited every conversation across studies in an exploratory, non-preregistered analysis. We used GPT-5.4-mini to score all 4,012 conversations on four dimensions on a 1–5 scale: how much the bot delivered *stereotype-disconfirming substance* and *informational specificity* (the cognitive route to prejudice reduction), and how much *empathy* and *friendliness* it conveyed (the affective route).^{6,11} Fig-

ure 6 shows the resulting condition contrasts in the three studies with a control arm; coding and analysis details are in Methods.

More than anything else, the outgroup bots offered stereotype-disconfirming substance: positions that cut against what their partner expected from the other side. Pooled across the three studies with a control arm, outgroup bots scored $M = 2.83$ on stereotype-disconfirming substance versus $M = 1.33$ for cats-and-dogs controls (1–5 scale), and the per-study gaps were consistently large: $\Delta M = +2.05$ points in the three-arm experiment ($p < .001$), $\Delta M = +1.41$ points in the behavioral study ($p < .001$), and $\Delta M = +1.42$ points in the longitudinal study ($p < .001$). This was not because the outgroup bots simply conveyed more detailed information: on informational specificity the control bots were, if anything, marginally higher within their own topics ($\Delta M = -0.20$ points in the experiment, $p < .001$; -0.05 behavioral; -0.06 longitudinal), so the cognitive contrast came from *what* the outgroup bot talked about, not how granular it was.

The affective dimensions tell a less consistent story. In the three-arm experiment, GPT rated the cats-and-dogs bot as *more* empathic than the outgroup bot ($\Delta M = -0.22$ points, $p < .001$) and roughly as friendly ($\Delta M = -0.06$ points, $p = .19$): the bot that produced the warmth gain was, if anything, the less pleasant of the two. In the behavioral and longitudinal studies, however, GPT rated the outgroup bot as both more empathic ($\Delta M = +0.50$, $p < .001$; $\Delta M = +0.34$, $p < .001$) and friendlier ($\Delta M = +0.46$, $p < .001$; $\Delta M = +0.40$, $p < .001$) than the chat control.




Fig. 6. Outgroup bots differ from control bots more in information than in friendliness. Per-study and pooled mean differences (outgroup bot – cats-and-dogs control) in GPT-5.4-mini’s ratings of each conversation on four dimensions (1–5 Likert), across the three-arm, behavioral, and longitudinal studies; the behavioral contrast is party-adjusted and subgroup estimates are faded. Error bars are 95% CIs; the pooled diamond is the inverse-variance-weighted average. The cognitive-route difference (stereotype-disconfirming substance) is the largest and most consistent contrast, present in every study and reliably larger than the affective-route difference (Appendix S7.2).

Outgroup and control bots differed more in stereotype-disconfirming content than in empathy or friendliness. Pooled across studies, the outgroup-vs-control gap was +0.37 Likert points wider on the cognitive dimensions than on the affective ones (95% CI [+0.31, +0.43], $p < .001$), and the same ordering held in every individual study (Methods; Appendix S7.2).

Representative verbatim excerpts of stereotype-disconfirming bot messages appear in Appendix S4.4.

Discussion

Brief AI synthetic contact is perceived as more acceptable than face-to-face contact, corrects misperceptions, warms cross-partisan affect, moves a costly behavioral choice, and—like other brief contact interventions—attenuates within a week.

Human intergroup contact works mainly through an affective route; synthetic contact appears to work mainly through a cognitive one—the affective channel was present in our conversations too, but consistently smaller (Fig. 6). The Pettigrew & Tropp^{8,11} meta-analysis of > 500 contact studies finds that the dominant mediators of human contact effects are affective—reduced intergroup anxiety and increased empathy—rather than purely cognitive. Consistent with that affective-mechanism account, Santoro and Broockman¹³ found that randomly assigning outpartisan strangers to discuss a shared experience reduced

affective polarization, but assigning them to discuss *partisan disagreement* did not: explicit political topics appear to introduce anxiety and threat that erode the affective gains human conversation otherwise produces.

We find the opposite pattern: explicit political disagreement, which fails to warm relations between human partisans, warmed them when the interlocutor was an AI. The most plausible reason is that an AI partner removes the interpersonal stakes that make political talk aversive—there is no risk of being judged and no face to manage—so participants engage with the disagreement on its merits rather than bracing against it. Freed of that threat, participants could take in the bot’s stereotype-disconfirming points—the process our content analyses point to, though we did not manipulate it directly. This evidence is exploratory, but it raises a possibility: the advantage of synthetic contact over face-to-face contact may not be only that it scales. By replacing the person on the other side with a bot, it may make the disagreements that matter most easier to discuss.

The introduction raised the risk that politically slanted language models would caricature the outgroup and entrench misperceptions. That risk partly materialized—both bots held more extreme positions than the partisans they represented—yet beliefs became more accurate, not less, because a guide need not be perfect to help, only less wrong than the learner. Bots calibrated against real survey data should do better still.

As expected for so brief an intervention, most of a single five-minute conversation's effect faded within a week, though a small residual persisted among the most extreme partisans—the group these interventions most need to reach. The strength of synthetic contact is that it can be repeated: unlike a face-to-face meeting, a bot conversation is available again whenever and wherever partisans are already online. Even the immediate effect carries weight—right after the conversation, partisans were more likely than controls to choose a real outgroup exchange. Whether repeated synthetic contact can build lasting warmth is a question for future work.

Two limitations qualify these findings. First, we indexed affective polarization mainly as outgroup warmth on a feeling thermometer—the standard measure, but one facet of a multidimensional construct³³—and did not test social distance, trait attributions, or support for anti-democratic action. Second, the evidence comes from online U.S. partisans and two issues, environment and immigration, so generalization across populations and topics remains open. The belief-accuracy gains, observed in a within-person design, also warrant more caution than the experimentally isolated warmth effects.

Traditional contact interventions work⁸ but do not scale, largely because people will not accept them.¹⁵ Synthetic contact reverses that constraint: it is inexpensive, available on demand, and far more acceptable than the face-to-face contact partisans avoid. Talking to an outgroup-representing bot corrects misperceptions and warms cross-partisan affect. Whether the effects shown here are sufficient to matter at the societal level remains an open question, but the cost-effectiveness of synthetic contact warrants optimism. A brief chatbot interaction embedded in a news app, social media platform, or civic engagement tool could reach millions of partisans where they already are: online, often alone, and unwilling to enter a room with the other side.

Methods

Aversion study

Participants. We recruited partisans via CloudResearch Connect with a preregistered target of 500 completions; owing to a recruitment-setting error, the study was fielded for 600, yielding 608 eligible participants (405 Democrats, 203 Republicans); results are unchanged if the sample is restricted to the first 500 respondents (Section S2.5). Random assignment yielded 308 participants in the bot condition and 300 in the human condition. The study was preregistered at AsPredicted #286,575.

Design. Two between-subjects conditions differed only in the identity of the conversation partner. In the *bot* condition, participants conversed for three minutes with an AI trained to represent a typical member of the opposing party. In the *human* condition, participants were paired in real time with another live participant from the opposing party for a three-minute text conversation. Topic (immigration), conversation length, staircase mechanics, and the alternative task were identical across conditions.

Procedure. An adaptive 1-up-1-down staircase estimated each participant's indifference point: the duration of mortality reflection equated to a three-minute outgroup conversation. On each trial, participants chose between “three minutes of conversation with an outgroup partisan / AI about immigration” and “X minutes of mortality reflection.” X started at 10 minutes with a halving step size ($5 \rightarrow 2.5 \rightarrow 1.25 \rightarrow 0.625 \rightarrow 0.25$ minutes) and a floor of 0.25 minutes. The staircase terminated after 12 trials, after three small reversals (step ≤ 0.5), or when the same duration was shown in two consecutive trials with the same choice. The threshold is the mean of the last two reversals, or the last-shown duration if fewer than two reversals occurred. The design was incentive-compatible: whichever option participants picked on the final trial, they actually had to complete.

Exclusions. Per preregistration, we excluded participants who did not identify as either a Democrat or a Republican.

Analytic strategy. The primary preregistered analysis regresses the mortality threshold on condition: $\text{threshold} \sim \text{condition}$. Preregistered secondary analyses add a $\text{condition} \times \text{extremity}$ interaction and a party control. Robustness checks include a Mann–Whitney test on the raw thresholds and a mixed-effects logistic regression on trial-level choices ($\text{chose_chat} \sim \text{duration} * \text{condition} + (1 + \text{duration} | \text{participant})$); see Appendix S2.5.

Within-person study

Participants. We recruited 500 participants via Prolific (248 Democrats learning about Republicans; 252 Republicans learning about Democrats). The study was preregistered at AsPredicted #264,402.

Design. The study used a within-person pre-post design with a single condition: every participant conversed with a chatbot prompted to represent their political outgroup (Republicans for Democratic learners; Democrats for Republican learners).

Procedure. Participants first rated their warmth toward the political outgroup on a 0–100 feeling thermometer (pre-interaction). They then conversed for ten minutes with the outgroup-representing chatbot, tasked with learning how the outgroup thinks about environmental policy. Immediately afterward, they re-rated outgroup warmth on the same thermometer (post-interaction) and rated the chatbot’s informativeness and empathy on 5-point Likert scales.

Exclusions. The preregistration specified no exclusions from the primary analyses, so all 500 participants are retained.

Analytic strategy. The primary preregistered analysis is a mixed-effects model on the pre-post warmth ratings: $\text{warmth} \sim \text{time} * \text{learner_party} + (1 \mid \text{participant})$, where *time* contrasts pre and post and *learner_party* contrasts D→R and R→D. The coefficient on *time* captures the average pre-post change; the interaction tests whether the change differs between learner groups. Models were estimated via maximum likelihood with `lme4`³⁴ and Satterthwaite degrees of freedom.³⁵

Three-arm experiment

Participants. We recruited 753 self-identified Democrats via Prolific. The study was preregistered on AsPredicted #276,530.

Design. Participants were randomly assigned to one of three conditions: (1) *synthetic contact*—a ten-minute conversation with a chatbot prompted to represent a Republican, in which participants role-played a non-partisan policy researcher tasked with learning how Republicans think about environmental policy; (2) *chat control*—a ten-minute conversation with the same interface, but with a bot instructed to debate whether cats or dogs are better and to avoid political topics, not to represent an outgroup member; or (3) *game control*—a ten-minute session playing Space Invaders, a browser-based arcade game with no conversational component.

Procedure. All participants completed their assigned ten-minute intervention. Immediately afterward, they rated outgroup warmth on a 0–100 feeling thermometer.

Exclusions. Per preregistration, we excluded participants who met any of the following criteria: reCAPTCHA score below 0.5, more than two tab switches during the intervention, or a total completion duration more than three standard deviations below the sample mean. After exclusions, the analytic sample

comprised $N = 679$ participants (198 synthetic contact, 244 cats-and-dogs chat, 237 Space Invaders).

Analytic strategy. The primary preregistered analysis is an OLS regression of outgroup warmth on condition: $\text{warmth} \sim \text{condition}$, with outgroup bot as the reference level. The two planned contrasts test our preregistered hypotheses (H1: outgroup bot > cats-and-dogs; H2: outgroup bot > Space Invaders). Wilcoxon rank-sum tests on each contrast provide non-parametric robustness given the zero-inflated distribution of warmth ratings.

Behavioral study

Participants. We recruited self-identified Democrats and Republicans via Prolific, screened for above-threshold political extremity (excluding participants with extremity < 10 on the 0–100 scale). The study was preregistered at AsPredicted #287,002, which specified a three-minute mortality-reflection comparison task and a sampling frame including Republicans recruited 50/50 with Democrats. Analyses are restricted to the preregistered sample—participants randomly assigned to the two chat conditions while the preregistration was in force ($N = 1069$; 572 Democrats, 497 Republicans).

Design. Participants were randomly assigned between subjects to one of two conditions: (1) *outgroup bot*—a five-minute conversation with a chatbot prompted to represent a typical member of their political outgroup, with participants instructed to learn how the outgroup thinks about immigration, or (2) *chat control*—a five-minute conversation with the same interface, but with a bot instructed to debate whether cats or dogs are better and to avoid political topics, not to represent an outgroup member. Both chatbots used the same conversational interface, the same underlying model (GPT-5.2-mini), and the same politeness norms; they differed only in assigned topic.

Procedure. Participants completed their assigned five-minute conversation, then faced an incentive-compatible binary choice between (a) “Have a three-minute conversation with a real {Republican/Democrat}” (outgroup label matched the participant’s party) or (b) “Complete a three-minute mortality reflection exercise.” Whichever option they chose, they actually completed it: real cross-partisan conversation vs. real mortality reflection. The primary outcome was whether the participant chose the outgroup conversation (1) or the mortality reflection (0).

Exclusions. Per preregistration, we excluded participants who met any of the following criteria: a Qualtrics reCAPTCHA score below 0.5, more than two tab switches during the chatbot page, a total completion duration more than three standard deviations below the sample mean, Prolific’s bot/LLM detection flagging the submission, or sending no messages to the chatbot.

Analytic strategy. The primary preregistered analysis is a logistic regression: $\text{task_choice} \sim \text{condition} + \text{party} + \text{extremity_c}$. We report the condition odds ratio, 95% confidence interval, and two-sided p value. As a non-parametric robustness check we report a Wilcoxon rank-sum test. Party and political extremity are examined as preregistered exploratory moderators.

Longitudinal study

Participants. We recruited self-identified Democrats (target $N = 1,200$ before exclusions) via CloudResearch Connect in three batches launched 2026-05-01, 2026-05-04, and 2026-05-06. Participants were screened on partisan identification and on a 0–100 political-extremity slider; those scoring below 30 on extremity were screened out before random assignment. The study was preregistered at AsPredicted #288,592.

Design. Participants were randomly assigned between subjects to one of two conditions: (1) *outgroup bot*—a five-minute conversation with a chatbot prompted to represent a typical Republican, with participants instructed to learn how Republicans think about environmental policy, or (2) *chat control*—a five-minute conversation with the same interface, but with a bot instructed to debate whether cats or dogs are better and to avoid political topics, not to represent an outgroup member. The two conditions are matched on interface, duration, model, and politeness norms, differing only in conversation topic.

Procedure. At Time 1, participants rated outgroup warmth on a 0–100 feeling thermometer (warmth_T1), completed their assigned five-minute conversation, and rated outgroup warmth again immediately post-chat. One week later, all Time 1 completers were invited back and re-rated outgroup warmth on the same thermometer (warmth_T2 , the primary dependent variable).

Exclusions. Per preregistration, we excluded participants who met any of the following criteria: failure of any of three embedded attention/bot checks (line-instruction, math, animal-recognition), a Qualtrics reCAPTCHA score below 0.5, more than two tab switches during the chatbot page, a total completion time more than three standard deviations below the

sample mean, or sending no messages to the chatbot. After exclusions, $N = 1104$ participants were retained at Time 1 (524 outgroup bot, 580 cats-and-dogs control).

Analytic strategy. The primary preregistered analysis is an ANCOVA: $\text{warmth_T2} \sim \text{condition} + \text{warmth_T1}$. Preregistered secondary analyses include a Wilcoxon rank-sum test of warmth_T2 by condition, a $\text{condition} * \text{extremity_c}$ interaction added to the primary ANCOVA, and the simple effect of condition within the upper half of the extremity distribution (median split). Because the Wave-2 warmth distribution showed strong floor bunching, we additionally report a rank ANCOVA on percentile-transformed outcomes; the attrition-weighting robustness check is in Appendix S6.4.

Cross-study content audit

Conversations coded. We coded all 4,012 bot conversations from the within-person, three-arm, behavioral, and longitudinal studies. The within-person study has no chat-control arm, so it does not contribute an outgroup-vs-control contrast and is omitted from the forest plot; the three studies with a control arm contribute the contrasts shown in Figure 6.

Coding scheme. Each conversation was scored on four theoretically motivated dimensions on a 1–5 scale: two indexing the cognitive route to prejudice reduction—*stereotype-disconfirming substance* and *informational specificity*—and two indexing the affective route—*empathy* and *friendliness*.^{6,11} To block cross-dimension halo bias, each dimension was coded in a separate API call (GPT-5.4-mini) with no information about the other dimensions; the model wrote one-sentence reasoning before scoring.

Per-dimension contrasts. For each dimension we computed the outgroup-vs-control mean difference, 95% CI, and Cohen’s d within each study. The behavioral contrast is adjusted for participant party because random assignment of bot prompt is crossed with party in that study; the other contrasts are unadjusted.

Cognitive-vs-affective comparison. To test whether the condition contrast is larger on the cognitive dimensions than on the affective ones, we pivoted the per-conversation scores to long format (one row per dimension) and fit a linear mixed-effects model: $\text{score} \sim \text{condition} * \text{route} + (1 | \text{participant})$, where *route* groups the two cognitive dimensions versus the two affective ones. A positive $\text{condition} \times \text{route}$ interaction means the outgroup-

vs-control gap is wider on the cognitive route than on the affective route. We fit this model separately within each of the three studies with a control arm and pooled across all three (Appendix S7.2, Table S22).

Preregistration. This content audit was not preregistered. It is an exploratory characterization of what the bots said and how they said it, not a confirmatory test.

Ethics

All studies were approved by the University of Pennsylvania Institutional Review Board (protocol #860019). All participants provided informed consent before participation.

Data and code availability

All de-identified data and analysis code are available at Zenodo (<https://doi.org/10.5281/zenodo.20971465>).

Author contributions, competing interests, and funding

Author contributions. B.L.L.: conceptualization, data curation, formal analysis, funding acquisition, investigation, methodology, resources, software, validation, visualization, writing (original draft), and writing (review and editing). N.C.: conceptualization, software, visualization, and writing (review and editing). S.P.: conceptualization, funding acquisition, supervision, visualization, and writing (review and editing). O.T.: conceptualization, supervision, visualization, and writing (review and editing). All authors reviewed and approved the final manuscript.

Competing interests. The authors declare no competing interests.

Acknowledgements. This research was supported by a University of Pennsylvania AI Fellowship to B.L.L., research funds from the Wharton School, and cost-sharing support from the Wharton Behavioral Lab.

References

1. S. Iyengar, Y. Leikes, M. Levendusky, N. Malhotra, S. J. Westwood, The origins and consequences of affective polarization in the United States. *Annual Review of Political Science* **22**, 129–146 (2019).
2. D. J. Ahler, G. Sood, The parties in our heads: Misperceptions about party composition and their consequences. *The Journal of Politics* **80**, 964–981 (2018).
3. D. A. Yudkin, S. Hawkins, T. Dixon, The perception gap: How false impressions are pulling Americans apart, *Tech. rep.*, More in Common (2019).
4. S. L. Moore-Berg, L.-O. Ankori-Karlinsky, B. Hameiri, E. Bruneau, Exaggerated meta-perceptions predict intergroup hostility between American political partisans. *Proceedings of the National Academy of Sciences* **117**, 14864–14872 (2020).
5. J. Lees, M. Cikara, Inaccurate group meta-perceptions drive negative out-group attributions in competitive contexts. *Nature Human Behaviour* **4**, 279–286 (2020).
6. E. Hermann, J. De Freitas, S. Puntoni, Reducing prejudice with counter-stereotypical AI. *Consumer Psychology Review* **8**, 75–86 (2025).
7. G. W. Allport, *The Nature of Prejudice* (Addison-Wesley, 1954).
8. T. F. Pettigrew, L. R. Tropp, A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology* **90**, 751–783 (2006).
9. R. Hartman, W. Blakey, J. Womick, C. Bail, E. J. Finkel, H. Han, J. Sarrouf, J. Schroeder, P. Sheeran, J. J. Van Bavel, R. Willer, K. Gray, Interventions to reduce partisan animosity. *Nature Human Behaviour* **6**, 1194–1205 (2022).
10. E. L. Paluck, S. A. Green, D. P. Green, The contact hypothesis re-evaluated. *Behavioural Public Policy* **3**, 129–158 (2019).
11. T. F. Pettigrew, L. R. Tropp, How does intergroup contact reduce prejudice? Meta-analytic tests of three mediators. *European Journal of Social Psychology* **38**, 922–934 (2008).
12. M. S. Levendusky, D. A. Stecula, *We Need to Talk: How Cross-Party Dialogue Reduces Affective Polarization*, Elements in Experimental Political Science (Cambridge University Press, 2021).
13. E. Santoro, D. E. Broockman, The promise and pitfalls of cross-partisan conversations for reducing affective polarization: Evidence from randomized experiments. *Science Advances* **8**, eabn5515 (2022).
14. M. K. Chen, R. Rohla, The effect of partisanship and political advertising on close family ties. *Science* **360**, 1020–1024 (2018).
15. J. A. Frimer, L. J. Skitka, M. Motyl, Liberals and conservatives are similarly motivated to avoid exposure to one another's opinions. *Journal of Experimental Social Psychology* **72**, 1–12 (2017).
16. C. A. Dorison, J. A. Minson, T. Rogers, Selective exposure partly relies on faulty affective forecasts. *Cognition* **188**, 98–107 (2019).
17. C. A. Bail, L. P. Argyle, T. W. Brown, J. P. Bumpus, H. Chen, M. B. F. Hunzaker, J. Lee, M. Mann, F. Merhout, A. Volfovsky, Exposure to opposing views on social media can increase political polarization. *Proceedings of the National Academy of Sciences* **115**, 9216–9221 (2018).
18. M. Saveski, N. Gillani, A. Yuan, P. Vijayaraghavan, D. Roy, *Proceedings of the International AAAI Conference on Web and Social Media (ICWSM)* (2022), vol. 16, pp. 885–895.
19. N. Gillani, A. Yuan, M. Saveski, S. Vosoughi, D. Roy, *Proceedings of the 2018 World Wide Web Conference (WWW '18)* (2018), pp. 823–831.
20. J. S. Park, C. Q. Zou, A. Shaw, B. M. Hill, C. Cai, M. R. Morris, R. Willer, P. Liang, M. S. Bernstein, Generative agent simulations of 1,000 people. *arXiv preprint arXiv:2411.10109* (2024).
21. O. Toubia, G. Z. Gui, T. Peng, D. J. Merlau, A. Li, H. Chen, Database report: Twin-2K-500: A dataset for building digital twins of over 2,000 people based on their answers to over 500 questions. *Marketing Science* (2025).
22. C. Nass, Y. Moon, Machines and mindlessness: Social responses to computers. *Journal of Social Issues* **56**, 81–103 (2000).
23. S. H. Klein, The effects of human-like social cues on social responses towards text-based conversational agents—a meta-analysis. *Humanities and Social Sciences Communications* **12** (2025).

24. L. Pereira da Costa, K. Bierwiazek, M. Bianchi, Does digital intergroup contact reduce prejudice? a meta-analysis. *Cyberpsychology, Behavior, and Social Networking* **27**, 440–451 (2024).
25. L. Lu, Z. L. Tormala, A. Duhachek, How AI sources can increase openness to opposing views. *Scientific Reports* **15**, 17170 (2025).
26. T. H. Costello, G. Pennycook, D. G. Rand, Durably reducing conspiracy beliefs through dialogues with AI. *Science* **385**, eadq1814 (2024).
27. H. Lin, G. Czarnek, B. Lewis, J. P. White, A. J. Berinsky, T. Costello, G. Pennycook, D. G. Rand, Persuading voters using human–artificial intelligence dialogues. *Nature* (2025).
28. F. Salvi, M. Horta Ribeiro, R. Gallotti, R. West, On the conversational persuasiveness of GPT-4. *Nature Human Behaviour* **9**, 1645–1653 (2025).
29. S. Santurkar, E. Durmus, F. Ladhak, C. Lee, P. Liang, T. Hashimoto, *Proceedings of the 40th International Conference on Machine Learning (ICML)* (2023).
30. J. Hartmann, J. Schwenzow, M. Witte, The political ideology of conversational AI: Converging evidence on ChatGPT's pro-environmental, left-libertarian orientation. *arXiv preprint arXiv:2301.01768* (2023).
31. K. L. Haws, K. P. Winterich, R. W. Naylor, Seeing the world through GREEN-tinted glasses: Green consumption values and responses to environmentally friendly products. *Journal of Consumer Psychology* **24**, 336–354 (2014).
32. J. G. Voelkel, M. N. Stagnaro, J. Chu, S. L. Pink, J. S. Mernyk, C. Redekopp, I. Ghezze, M. Cashman, D. Adjodah, *et al.*, Megastudy testing 25 treatments to reduce antidemocratic attitudes and partisan animosity. *Science* **386**, eadh4764 (2024).
33. J. N. Druckman, M. S. Levendusky, What do we measure when we measure affective polarization? *Public Opinion Quarterly* **83**, 114–122 (2019).
34. D. Bates, M. Mächler, B. Bolker, S. Walker, Fitting linear mixed-effects models using lme4. *Journal of Statistical Software* **67**, 1–48 (2015).
35. S. G. Luke, Evaluating significance in linear mixed-effects models in R. *Behavior Research Methods* **49**, 1494–1502 (2017).
36. M. A. Hernán, J. M. Robins, *Causal Inference: What If* (Chapman & Hall/CRC, Boca Raton, 2020).

*Supplementary Information for***Synthetic Contact with AI Reduces Cross-Partisan Animosity**Benjamin Lira Luttges^{1,*}, Noah Castelo², Stefano Puntoni¹, Olivier Toubia³¹The Wharton School, University of Pennsylvania. ²Alberta School of Business, University of Alberta. ³Columbia Business School, Columbia University.*Corresponding author: blira@upenn.edu.**Contents**

S1 Samples, exclusions, and deviations from preregistration	16
S2 Aversion Experiment	17
S2.1 Randomization balance	17
S2.2 Means, SDs, and correlations	17
S2.3 Primary model behind Figure 1	17
S2.4 Distribution, sensitivity, and preregistered interactions	17
S2.5 Trial-level mixed-effects model	17
S2.6 Sample cross-party conversations	18
S3 Within-Person Study	19
S3.1 Sample descriptives	19
S3.2 Means and misperception tests behind Figure 2A	19
S3.3 Correlations	20
S3.4 Mixed-effects models	20
S3.5 Ingroup warmth and affective polarization	20
S3.6 Individual heterogeneity	21
S3.7 Bot-item semantic proximity and belief updating	21
S4 Three-Arm Experiment	24
S4.1 Randomization balance	24
S4.2 Correlations	24
S4.3 Main-effects regression	24
S4.4 Verbatim examples of disconfirming content	25
S5 Behavioral Experiment	25
S5.1 Randomization balance	25
S5.2 Means, SDs, and correlations	25
S5.3 Primary model behind Figure 4	26
S6 Longitudinal Experiment	26
S6.1 Randomization balance	26
S6.2 Means, SDs, and correlations	26
S6.3 Primary model behind Figure 5	26
S6.4 Inverse-probability-of-attrition weighting	26
S6.5 Rank-based robustness across studies	27
S6.6 Pooled one-week persistence across the two follow-ups	27

S7 Cross-Study Process Audit	29
S7.1 Per-study and pooled mean differences behind Figure 6	29
S7.2 Test that the cognitive effect exceeds the affective effect	29
S7.3 Conversation-coding rubric	29
S8 Measures	31
S9 System Prompts	32
S9.1 Outgroup bot, initial prompt (within-person and three-arm studies)	32
S9.2 Outgroup bot, revised prompt (behavioral and longitudinal studies)	32
S9.3 Cats-vs-dogs control bot (all studies)	32

S1 Samples, exclusions, and deviations from preregistration

Each study was preregistered on AsPredicted. For each we report the registration number, the recruited and analytic samples, the preregistered exclusion rules, and any deviations from the preregistered analysis.

Aversion experiment (AsPredicted #286,575). We preregistered a target of 500 completions; owing to a recruitment-setting error, the study was fielded for 600, yielding $N = 608$ eligible participants. Restricting the analysis to the first 500 participants by completion time (Section S2.5) leaves the estimate unchanged. The only preregistered exclusion removed participants who did not identify as Democrat or Republican, and there were no deviations from the preregistered analysis.

Within-person study (AsPredicted #264,402). $N = 500$ partisans (248 learning about Republicans, 252 learning about Democrats). The preregistration specified no exclusions from the primary analyses; all participants are retained. No deviations.

Three-arm experiment (AsPredicted #276,530). 753 Democrats recruited; $N = 679$ retained after the preregistered exclusions (reCAPTCHA score below 0.5, more than two tab switches during the intervention, or completion time more than three standard deviations below the sample mean). No deviations.

Behavioral experiment (AsPredicted #287,002). The preregistration specified a three-minute mortality-reflection comparison and a sampling frame with Republicans recruited 50/50 with Democrats. Analyses are restricted to the preregistered sample—participants randomly assigned to the two chat conditions while preregistration #287,002 was in force ($N = 1069$). Preregistered exclusions: reCAPTCHA below 0.5, more than two tab switches, completion time more than three standard deviations below the mean, a bot/LLM detection flag, or sending no messages to the chatbot. No deviations.

Longitudinal experiment (AsPredicted #288,592). Target of 1,200 Democrats before exclusions; $N = 1104$ retained at Time 1 after the preregistered exclusions (failure of any of three attention/bot checks, reCAPTCHA below 0.5, more than two tab switches, completion time more than three standard deviations below the mean, or sending no messages), of whom 924 returned at the one-week follow-up. Deviation: because the Wave-2 warmth distribution was floor-bunched and violated OLS assumptions, the one-week effect is reported with a rank ANCOVA (Section S6.5).

S2 Aversion Experiment

S2.1 Randomization balance

Table S1 reports demographic characteristics by condition in the aversion experiment. No variable differed significantly across conditions ($p > .05$ for all tests), confirming successful randomization. Gender and race were not collected in this study.

Table S1. Randomization balance in the aversion experiment. Continuous variables report M (SD); binary variables report percentages. Tests are Welch's t -test for continuous and χ^2 for categorical variables.

Variable	Bot	Human	Test	p
Age (years)	43.78 (13.19)	42.76 (13.10)	$t(666) = 1.01$	$= 0.315$
Democrat	38.0%	44.9%	$\chi^2(1) = 2.99$	$= 0.084$
Political extremity	73.59 (27.96)	67.91 (31.16)	$t(441) = 2.02$	$= 0.044$
Duration (min)	9.06 (53.26)	6.89 (11.97)	$t(369) = 0.73$	$= 0.467$

S2.2 Means, SDs, and correlations

Table S2. Means, standard deviations, and bivariate correlations among key variables in the aversion experiment. * $p < .05$.

Variable	M	SD	(1)	(2)	(3)	(4)	(5)
(1) Mortality threshold (min)	7.74	14.27	—				
(2) Human condition (vs. bot)	0.50	0.50	0.17*	—			
(3) Age	43.27	13.15	0.01	-0.04	—		
(4) Political extremity	70.64	29.77	0.03	-0.10*	0.14*	—	
(5) Duration (min)	7.98	38.64	-0.01	-0.03	0.02	-0.03	—

S2.3 Primary model behind Figure 1

Table S3. Mortality-threshold means by condition and the regression threshold $\sim$ condition that underlies Figure 1. Coefficient is the human–bot difference in minutes.

	n	M (SD)	Coef.	95% CI	t	p
Bot	320	5.27 (9.79)	ref.	—	—	—
Human	326	10.15 (17.26)	4.88	[2.71, 7.05]	4.41	$< .001$
Cohen's d				0.35		

S2.4 Distribution, sensitivity, and preregistered interactions

Medians tell the same story as means but more compressed (3.38 vs. 3.00 minutes), indicating that the mean gap is amplified by a longer right tail among participants anticipating human contact. A Mann–Whitney test on the raw thresholds confirms the effect non-parametrically ($p < .001$). Restricting analysis to the first 500 participants by completion time leaves the estimate essentially unchanged ($\beta = 4.15$, $p < .001$), ruling out the possibility that over-provisioning past our preregistered target of 500 drives the result. The preregistered interaction with political extremity was not significant ($\beta = 0.004$, $p = .91$), and the bot's acceptance advantage held within each party (Democrats: $\beta = 5.09$ min, $p < .001$; Republicans: $\beta = 3.40$ min, $p = .03$; no significant partner $\times$ party interaction, $p = .47$).

S2.5 Trial-level mixed-effects model

The primary analysis in the main text regresses each participant's staircase-derived threshold (the mean of the last two reversals) on condition. As a preregistered robustness check, we re-fit the model at the trial level using a mixed-effects logistic regression: $P(\text{chose chat}) = \text{logit}^{-1}(b_0 + b_1 \text{duration} + b_2 \text{partner} + b_3 \text{duration} \times \text{partner})$ with random intercepts and slopes on duration by participant (6178 choices from 608 participants).

Random-effects structure. We include random slopes on duration (in addition to random intercepts) because participants vary in their discrimination sensitivity—the slope of each participant's psychometric curve is itself a

quantity of interest in a staircase design, not a nuisance. The data are unambiguous on this point: compared to an intercept-only specification, the random-slope model is strongly favored by a likelihood ratio test ($\chi^2(2) = 1187.5$, $p < .001$; $\Delta\text{AIC} = 1184$), and the intercept-only model is in fact singular (the random intercept variance collapses to zero), indicating that participant-level variability in *duration sensitivity* dominates variability in baseline choice probability.

Results. Figure S1 shows the fitted linear predictor (log-odds of choosing the three-minute chat) as a function of offered mortality duration, by condition. The two lines cross zero at the psychometric midpoint: bot participants are indifferent at $X = 2.75$ min; human participants at $X = 3.00$ min. The slope of the human condition is shallower than that of the bot condition (duration $\times$ partner interaction: $\beta = -0.760$, $p < .001$), which is consistent with the main-text finding: the psychometric curve rises more slowly in the human condition, so at any given mortality duration above the indifference point, a smaller share of human-condition participants have switched over to choosing chat. The psychometric midpoints are close across conditions because the threshold distribution is heavily right-skewed (medians 3.00 and 3.38 min, against means of 5.06 and 9.65): the typical participant is only modestly more human-averse, while the condition difference is carried by the shallower human slope and a heavier right tail of strongly human-averse participants—which the mean threshold in the main text reflects, and a single midpoint cannot.




Fig. S1. Psychometric functions by condition (log-odds scale). Lines: fitted linear predictor of the GLMM (random intercepts and slopes by participant). Points: observed log-odds within binned durations (size proportional to n ; bins with $p \in \{0, 1\}$ dropped). Dashed line at $y = 0$ marks $P = 0.5$; white circles mark each curve's crossing.

S2.6 Sample cross-party conversations

In the human condition, participants who reached the chat task were paired with a real outgroup partner for a three-minute live text conversation about immigration. Two verbatim exchanges appear below to illustrate the range of dialogue these participants produced. The first shows policy-level disagreement bounded by mutual civility; the second shows partial within-party dissent that surfaces in cross-party contact.

Democrat: *Hi there, how are you?*

Republican: *Hi! I am good how are you*

Democrat: *Doing well, so what is your take on immigration?*

Democrat: *I personally think we need to cut down on illegal immigration but do it in a far more humane way than Trump and the Republicans have been handling it*

Republican: *I think illegal immigration has to be curbed almost completely, and legal immigration should be*

encouraged but largely based on skill

Democrat: *We should not have people coming across the border illegally, I agree 100%*

Republican: *I do not think adding a ton of unskilled workers or non-working people is a general positive for the country, so I think we should have like a skill based immigration system*

Democrat: *But sending in violent ICE agents to get rid of them does nobody any good*

Republican: *That being said I do disagree with the Trump admin on giving out a ton of H1B visas*

Republican: *How would you suggest we get rid of them?*

Democrat: *Send in local police forces that actually follow the law*

Republican: *I see. It would require a lot of coordination between the federal and local governments and I am not sure they could force local police to do it because of separation of state and federal, but if they could I would support it*

Democrat: *Hi, whats your thoughts on immigration*

Republican: *Hello.*

Democrat: *I think we need to re do our immigration system*

Republican: *I think that immigration is good for our country. I do disagree with illegal immigration however.*

Republican: *What are your thoughts?*

Democrat: *I agree. I think people should come here the right way. I just don't think its as easy as people think it is*

Republican: *I also agree with that. I know that isn't very popular among other 'Republicans', which is why I cannot fully identify as a Republican.*

Democrat: *I feel the same way about democrats. I do agree with most other policies*

Republican: *I think that we should help with barriers to immigration and integrate immigrants so that they can be here the right way and have proper protections and resources.*

Across the cross-party conversations, no exchange contained personal incivility, profanity, or name-calling directed at the conversation partner: people disagreed substantively but treated each other politely. For sample outgroup-bot conversations of the kind used in the bot condition, see Appendix S4.4.

S3 Within-Person Study

S3.1 Sample descriptives

Table S4 reports demographic characteristics and pre-interaction outgroup warmth by learner party.

Table S4. Sample characteristics by learner party. Values in parentheses are standard deviations.

	Democrats	Republicans
<i>N</i>	248	252
Age	43.3 (14.0)	43.8 (13.1)
Female (%)	53	52
White (%)	71	81
Extremity (0–100)	40.0 (13.0)	36.0 (14.9)
Pre-warmth (0–100)	22.6 (22.8)	35.9 (24.4)

S3.2 Means and misperception tests behind Figure 2A

Table S5. Own-party attitudes and outgroup-perceived attitudes on the 6-item GREEN environmental scale (1–5). Each party's outgroup belief is compared against the outgroup's own reported attitude via Welch's *t*-test, the same comparison visualized in Figure 2A.

Quantity	<i>M</i> (SD)	<i>n</i>	<i>t</i> vs. reality	<i>p</i>
Republican attitudes (own report)	3.73 (1.06)	252	—	—
Democrats' estimate of Republicans	2.16 (1.03)	248	-16.79	< .001
Democratic attitudes (own report)	4.15 (0.77)	248	—	—
Republicans' estimate of Democrats	3.96 (0.93)	252	-2.46	0.014

S3.3 Correlations

Table S6 reports means, standard deviations, and bivariate correlations among the key study variables.

Table S6. Means, standard deviations, and bivariate correlations among study variables. * $p < .05$.

	<i>M</i>	<i>SD</i>	1	2	3	4	5	6	7	8	9
1. Pre-warmth	29.32	24.54	—								
2. Post-warmth	33.66	25.70	0.89*	—							
3. Pre-accuracy	-1.19	0.89	0.49*	0.45*	—						
4. Post-accuracy	-0.80	0.72	0.32*	0.35*	0.46*	—					
5. Extremity	38.02	14.14	-0.31*	-0.27*	-0.20*	-0.16*	—				
6. Informativeness	3.84	1.06	0.22*	0.29*	0.22*	0.25*	0.03	—			
7. Empathy	3.53	1.04	0.13*	0.18*	0.04	0.13*	-0.01	0.35*	—		
8. Bot turns	11.11	6.00	-0.03	-0.05	-0.05	-0.04	-0.04	-0.08	-0.05	—	
9. Words written	199.86	123.31	-0.02	-0.04	-0.03	-0.02	0.02	-0.03	0.06	0.18*	—

S3.4 Mixed-effects models

Table S7. Mixed-effects models predicting outgroup warmth and belief accuracy. Belief accuracy is the negative absolute error between a participant's guess of the outgroup's environmental attitudes and the outgroup's actual mean (higher = more accurate). Columns (1) and (4) include only the pre-to-post time effect. Columns (2) and (5) add the Learner party interaction. Columns (3) and (6) additionally include political extremity (centered) and its interaction with time. Standard errors in parentheses. Reference group: Democrats learning about Republicans, pre-interaction.

	Outgroup warmth			Belief accuracy		
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	29.316*** (1.123)	22.601*** (1.543)	23.574*** (1.497)	-1.189*** (0.036)	-1.680*** (0.045)	-1.665*** (0.045)
Post (vs. pre)	4.342*** (0.524)	5.149*** (0.743)	5.055*** (0.745)	0.393*** (0.038)	0.628*** (0.052)	0.624*** (0.052)
R→D (vs. D→R)		13.324*** (2.174)	11.393*** (2.119)		0.976*** (0.063)	0.946*** (0.064)
Post × R→D		-1.602 (1.046)	-1.415 (1.055)		-0.465*** (0.073)	-0.458*** (0.073)
Extremity (centered)			-0.480*** (0.075)			-0.007*** (0.002)
Post × Extremity			0.046 (0.037)			0.002 (0.003)
Num.Obs.	1000	1000	1000	1000	1000	1000

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

S3.5 Ingroup warmth and affective polarization

Our primary measure is warmth toward the outgroup, which a shift in feelings toward one's own party cannot mechanically inflate. For completeness, we nonetheless confirm that the gain reflects genuine outgroup warming rather than a general flattening of partisan affect, and that it also holds on the standard affective-polarization index (ingroup minus outgroup warmth, where higher values indicate more polarization).

Ingroup warmth decreased only slightly after the chatbot interaction ($b = -1.07$, $SE = 0.54$, $p = .05$)—a fraction of the 4.3-point rise in outgroup warmth—and the change did not differ between Democrats and Republicans ($b = -0.59$, $SE = 0.76$, $p = .44$). Affective polarization declined in both groups: among Democrats, the index fell from 53.4 to 47.2 (a change of -6.2); among Republicans, from 45.1 to 39.9 (a change of -5.2). A mixed-effects model confirmed the decline was significant ($b = -6.22$, $SE = 0.94$, $p < .001$), and both groups depolarized at similar rates ($b = 1.01$, $SE = 1.33$, $p = .45$; Table S8).

The warmth gains reported in the main analysis reflect genuine improvement toward the outgroup, not compression of the feeling thermometer.

Table S8. Mixed-effects models predicting ingroup warmth and the affective polarization index (ingroup minus outgroup warmth) before and after chatbot interaction. Standard errors in parentheses. Reference group: Democrats learning about Republicans, pre-interaction.

	Ingroup warmth	Affective polarization
Intercept	76.016*** (1.266)	53.415*** (1.979)
Post (vs. pre)	-1.069* (0.540)	-6.218*** (0.941)
R→D (vs. D→R)	4.988** (1.783)	-8.336** (2.787)
Post × R→D	-0.590 (0.760)	1.011 (1.326)
Num.Obs.	1000	1000

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

S3.6 Individual heterogeneity

The mean warmth gain masks variation across individuals. Figure S2 plots the distribution of individual-level warmth change (post minus pre) for each group. Among Democrats, 49% showed a positive change in outgroup warmth; among Republicans, 45% did. About half of each group warmed toward the outgroup; most of the remainder were unchanged, and only a minority grew colder. This distribution confirms that the mean effect is not driven by a handful of outliers.




Fig. S2. Distribution of individual-level change in outgroup warmth (post – pre), by learner party. Dashed lines mark the group mean. Annotations show the percentage of participants with a positive change.

S3.7 Bot-item semantic proximity and belief updating

A natural worry about the conversation-content claim is that we have only inferred topic exposure from end-of-chat ratings of the bot. We address this here by measuring topic exposure directly from the transcripts and testing whether conversations that hewed closer to the six GREEN items produced more belief updating and more warmth gain. We embed each of the six GREEN items and each of the 5588 bot turns across the 500 conversations using OpenAI’s text-embedding-3-large. For every participant we compute a single proximity score: for each item,

the maximum cosine similarity between that item and any of the participant's bot turns, then averaged across the six items (i.e., “did the bot hit every item at some point”):

$$\text{proximity}_p = \frac{1}{6} \sum_{i=1}^6 \max_{t \in T_p} \cos(\mathbf{e}_i, \mathbf{e}_t), \quad (1)$$

where $\mathbf{e}_i$ is the embedding of GREEN item i , T_p is the set of participant p 's bot turns, and $\mathbf{e}_t$ is the embedding of turn t .

Pooled across both directions, proximity is associated with belief accuracy gain ($r = 0.09$, $p = 0.04$) but not warmth gain ($r = 0.03$, $p = 0.52$; Table S9, Panel A). The accuracy effect is concentrated in the D→R direction ($r = 0.19$, $p = 0.003$); in R→D it is null ($r = -0.04$, $p = 0.55$), and the party $\times$ proximity interaction is significant ($b = -0.20$, $SE = 0.07$, $p = 0.007$). Belief updating and warmth gain themselves are coupled only for Democrats ($r = 0.22$, $p < .001$); for Republicans they are independent ($r = -0.01$, $p = 0.89$).

A second, independent content measure points the same way. Coding every bot message for whether its substance contradicted (*disconfirming*), reinforced (*confirming*), or was irrelevant to (*neutral*) the outgroup stereotype, we find confirming content was negligible—only 1.5% of bot messages—so conversations varied chiefly in *how much* stereotype-disconfirming substance the bot delivered, not in whether it also confirmed. Among Democrats learning about Republicans, conversations carrying more disconfirming substance produced larger accuracy gains (standardized $b = 0.16$, $p = .003$, controlling baseline accuracy; the per-message disconfirming proportion trends in the same direction, $p = .11$); in R→D the association is null. As with proximity, this measure is correlational and post-randomization—more disconfirming content could be both a cause and a marker of an engaged participant—so it corroborates rather than proves the content route.

We cannot test whether disconfirming content mediates the *between-condition* warmth effect directly: control-bot conversations carry virtually no stereotype-disconfirming substance (0.1% of cats-and-dogs messages vs. 46% of outgroup-bot messages), so the candidate mediator is, by design, collinear with condition and the model is not identified. Within the outgroup-bot conversations we can instead ask whether variation in disconfirming content tracks warmth gain through belief updating. Among Democrats learning about Republicans ($N = 248$), conversations carrying more disconfirming content produced marginally larger accuracy gains ($a = 0.08$, $p = .10$), accuracy gains strongly predicted warmth gains ($b = 3.46$, $p < .001$), and the bootstrapped indirect path was positive but not significant ($a \cdot b = 0.26$, 95% CI $[-0.05, 0.65]$, $p = .14$). The data are thus consistent with a content→belief→warmth route but underpowered to confirm it at the conversation level. In the R→D direction the same model is null at every path ($N = 252$; indirect $a \cdot b = -0.03$, 95% CI $[-0.14, 0.08]$)—as expected, since Republicans begin nearly accurate about Democrats and their warmth gains do not track belief updating in the first place.

In the D→R sample, belief updating accounts for the link between conversation content and warmth (Figure S3, Table S9, Panel B). Conversations with more on-topic bot turns produce larger accuracy gains ($a = 0.19$, $SE = 0.06$, $p < .001$); accuracy gain in turn predicts warmth gain controlling for proximity ($b = 0.22$, $SE = 0.08$, $p = 0.007$); proximity has no direct effect on warmth once accuracy is held constant ($c' = -0.00$, $p = 0.94$). The bootstrapped indirect path is significant ($a \cdot b = 0.041$, 95% CI $[0.008, 0.086]$, $p = 0.04$), and the total effect is not, consistent with full indirect-only mediation. In the R→D sample all three paths are null. The asymmetry mirrors the by-party gradient in accuracy gain reported in Section S3.2: where belief updating happens, content-driven proximity tracks it, and warmth tracks updating in turn.




Fig. S3. Conversation-level bot-item semantic proximity vs. belief accuracy gain, by learner direction. Points are individual participants; lines are OLS fits with 95% confidence bands. Proximity tracks accuracy gain in D→R but not in R→D.

Table S9. Bot-item semantic proximity, belief updating, and warmth gain. Panel A reports Pearson correlations between conversation-level proximity and the two outcomes, pooled and by learner direction. Panel B reports the mediation model proximity → accuracy gain → warmth gain in both learner directions, with bootstrapped (2,000 draws) confidence intervals for the indirect effect. All variables standardized.

Panel A: Correlations		N	<i>r</i>	<i>p</i>	
Pooled	Accuracy gain	500	0.09	0.043	
	Warmth gain	500	0.03	0.525	
D→R	Accuracy gain	248	0.19	0.003	
	Warmth gain	248	0.04	0.571	
R→D	Accuracy gain	252	-0.04	0.546	
	Warmth gain	252	0.02	0.787	
Panel B: Mediation		Est.	SE	95% CI	<i>p</i>
<i>Democrats → Republicans</i>					
Proximity → Accuracy gain (<i>a</i>)		0.19	0.06	[0.08, 0.31]	< .001
Accuracy gain → Warmth gain (<i>b</i>)		0.22	0.08	[0.06, 0.37]	0.007
Proximity → Warmth gain (direct, <i>c'</i>)		-0.00	0.06	[-0.13, 0.12]	0.940
Indirect <i>a · b</i>		0.04	0.02	[0.01, 0.09]	0.037
Total		0.04	0.06	[-0.07, 0.15]	0.525
<i>Republicans → Democrats</i>					
Proximity → Accuracy gain (<i>a</i>)		-0.04	0.07	[-0.19, 0.10]	0.594
Accuracy gain → Warmth gain (<i>b</i>)		-0.01	0.06	[-0.13, 0.10]	0.890
Proximity → Warmth gain (direct, <i>c'</i>)		0.02	0.06	[-0.11, 0.14]	0.791
Indirect <i>a · b</i>		0.00	0.00	[-0.01, 0.01]	0.947
Total		0.02	0.06	[-0.11, 0.14]	0.783

S4 Three-Arm Experiment

S4.1 Randomization balance

Table S10 reports demographic characteristics by condition. No variable differed significantly across conditions ($p > .05$ for all tests), confirming successful randomization.

Table S10. Randomization balance across experimental conditions. Continuous variables report M (SD); binary variables report percentages. Tests are one-way ANOVA (F) for continuous and χ^2 for categorical variables.

	Outgroup bot	Cats/dogs bot	Space Invaders	Statistic	p
Age	40.5 (14.4)	39.8 (13.5)	40.3 (14.4)	0.14	.873
Female (%)	62	64	65	0.33	.847
White (%)	75	71	69	1.89	.389
Extremity (0–100)	82.1 (20.5)	79.8 (21.7)	77.2 (24.5)	2.56	.078
Duration (sec)	764.9 (232.4)	761.9 (249.2)	742.4 (184.2)	0.68	.506

S4.2 Correlations

Table S11 reports bivariate correlations among key variables in the three-arm experiment.

Table S11. Bivariate correlations among key variables in the three-arm experiment. Upper triangle reports Pearson r . * $p < .05$, ** $p < .01$, *** $p < .001$.

	M	SD	1	2	3	4	5	6	7	8
1. Outgroup warmth	20.57	20.85	—							
2. Ingroup warmth	76.43	19.02	0.01	—						
3. Age	40.16	14.06	-0.03	0.19*	—					
4. Extremity	79.60	22.43	-0.13*	0.58*	0.17*	—				
5. Duration (sec)	755.97	223.28	0.05	0.09*	0.04	0.03	—			
6. Bot: stereotype-disconfirming	2.21	1.44	0.20*	0.09	0.04	0.11*	-0.06	—		
7. Bot: empathy	3.88	0.37	-0.07	0.02	0.06	0.02	-0.05	-0.04	—	
8. Bot: friendliness	4.41	0.50	0.02	0.14*	-0.04	0.05	-0.01	-0.01	0.12*	—

S4.3 Main-effects regression

Table S12. OLS models predicting outgroup warmth from condition assignment. Model (1) uses the outgroup bot as reference; Model (2) uses the cats-and-dogs control. Standard errors in parentheses.

	Outgroup bot ref.	Cats/dogs ref.
Intercept	29.187*** (1.430)	16.967*** (1.289)
Cats/dogs bot	-12.220*** (1.925)	
Space Invaders	-12.098*** (1.938)	0.121 (1.836)
Outgroup bot		12.220*** (1.925)
Num.Obs.	679	679
R2	0.070	0.070
R2 Adj.	0.068	0.068

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

S4.4 Verbatim examples of disconfirming content

Three verbatim excerpts from Republican-bot conversations with Democratic learners in the three-arm experiment, drawn from messages coded as stereotype-disconfirming.

User: *I see traditional values as centered around family and religion, as being closed minded and putting others down.*

Republican bot: *Family and faith have always been core values for a lot of people in the Republican Party... However, I think it's important to recognize that there are many different interpretations of those values. Not everyone in the party is close-minded or dismissive of others; many people are working hard to find common ground and build connections.*

User: *I'm hoping we could discuss your opinion regarding environmental policy.*

Republican bot: *I think it's important to protect the environment, but I also believe in balancing that with economic growth and job creation... practical solutions that encourage innovation, like promoting clean energy technologies without going overboard with regulations.*

User: *The current republican party has a ton of control and republicans are just letting it happen.*

Republican bot: *There's definitely a lot of discussion around how the Republican Party has changed... Some folks feel that the party has shifted too far in one direction or that it's not representing their values as well as it used to.*

S5 Behavioral Experiment

S5.1 Randomization balance

Table S13 reports demographic characteristics by condition in the behavioral experiment. No variable differed significantly across conditions ($p > .05$ for all tests), confirming successful randomization.

Table S13. Randomization balance in the behavioral experiment. Continuous variables report M (SD); binary variables report percentages. Tests are one-way ANOVA (F) for continuous and χ^2 for categorical variables.

Variable	Outgroup bot	Cats/dogs	Test	p
Age (years)	38.82 (14.20)	38.28 (13.25)	$F(1, 1713) = 0.67$	$= 0.415$
Female	61.9%	62.7%	$\chi^2(1) = 0.08$	$= 0.774$
White	70.7%	74.8%	$\chi^2(1) = 3.41$	$= 0.065$
Political extremity	78.19 (17.82)	77.79 (18.58)	$F(1, 1752) = 0.22$	$= 0.641$
Duration (min)	26.24 (460.56)	10.42 (4.55)	$F(1, 1752) = 1.03$	$= 0.310$

S5.2 Means, SDs, and correlations

Table S14. Means, standard deviations, and bivariate correlations among key variables in the behavioral experiment, including the four coded bot-content dimensions. * $p < .05$.

Variable	M	SD	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Chose outgroup conversation	0.66	0.47	—							
(2) Outgroup bot (vs. cats/dogs)	0.50	0.50	0.10*	—						
(3) Age	38.56	13.74	-0.01	0.02	—					
(4) Party (Rep ↑)	76.59	17.40	0.04	0.01	0.12*	—				
(5) Political extremity	78.06	18.17	0.02	0.01	0.09*	1.00*	—			
(6) Bot: stereotype-disconfirming	2.04	1.38	0.07*	0.51*	0.00	0.01	0.05*	—		
(7) Bot: empathy	3.85	0.66	0.04	0.38*	-0.05*	0.02	0.03	0.40*	—	
(8) Bot: friendliness	4.17	0.55	0.07*	0.42*	-0.02	-0.02	0.00	0.25*	0.38*	—

S5.3 Primary model behind Figure 4

Table S15. Preregistered logistic regression $\text{outgroup_chat} \sim \text{condition} + \text{party} + \text{extremity_c}$ that underlies Figure 4.

Term	Coef. (logit)	OR	95% CI (OR)	<i>p</i>
Intercept	0.40	1.49	[1.21, 1.83]	< .001
Outgroup bot (vs. cats/dogs)	0.29	1.33	[1.04, 1.71]	= 0.025
Republican (vs. Democrat)	0.07	1.07	[0.83, 1.37]	= 0.610
Political extremity (centered)	0.00	1.00	[0.99, 1.01]	= 0.805
<i>N</i>	1069			

S6 Longitudinal Experiment

S6.1 Randomization balance

Table S16 reports demographic characteristics by condition in the longitudinal experiment. No variable differed significantly across conditions ($p > .05$ for all tests), confirming successful randomization.

Table S16. Randomization balance in the longitudinal experiment. Continuous variables report M (SD); binary variables report percentages. Tests are one-way ANOVA (F) for continuous and χ^2 for categorical variables.

Variable	Outgroup bot	Cats/dogs	Test	<i>p</i>
Age (years)	40.18 (12.95)	39.63 (13.40)	$F(1, 1194) = 0.52$	= 0.472
Female	62.1%	60.7%	$\chi^2(1) = 0.20$	= 0.654
White	65.0%	68.3%	$\chi^2(1) = 1.45$	= 0.229
Political extremity	80.33 (17.83)	80.07 (18.35)	$F(1, 1280) = 0.07$	= 0.794
Duration (min)	21.01 (278.48)	7.75 (4.78)	$F(1, 1280) = 1.45$	= 0.229

S6.2 Means, SDs, and correlations

Table S17. Means, standard deviations, and bivariate correlations among key variables in the longitudinal experiment, including the four coded bot-content dimensions. * $p < .05$.

Variable	<i>M</i>	<i>SD</i>	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Outgroup warmth (post-chat)	24.64	23.89	—							
(2) Outgroup warmth (baseline)	21.09	22.83	0.71*	—						
(3) Outgroup bot (vs. cats/dogs)	0.50	0.50	0.23*	0.04	—					
(4) Age	39.90	13.18	-0.02	-0.07*	0.02	—				
(5) Political extremity	80.10	18.14	-0.15*	-0.18*	0.01	0.17*	—			
(6) Bot: stereotype-disconfirming	2.07	1.39	0.07*	-0.06*	0.51*	0.03	0.08*	—		
(7) Bot: empathy	3.75	0.69	0.05	-0.01	0.24*	-0.04	0.04	0.36*	—	
(8) Bot: friendliness	4.17	0.56	0.09*	0.00	0.36*	0.05	0.04	0.23*	0.35*	—

S6.3 Primary model behind Figure 5

Table S18. ANCOVA on Wave-2 outgroup warmth, with baseline warmth as covariate, in both raw-thermometer and percentile-rank specifications. These models underlie Figure 5.

Model	Coef. (condition)	95% CI	<i>t</i>	<i>p</i>
ANCOVA (raw thermometer)	9.43	[7.59, 11.27]	10.05	< .001
Rank ANCOVA (percentile)	10.86	[9.02, 12.69]	11.58	< .001
<i>N</i>	1197			

S6.4 Inverse-probability-of-attrition weighting

Wave-2 retention is somewhat lower in the outgroup-bot arm than in the chat control (Section “Attrition” in the main text), raising the possibility that the null one-week effect reflects selective attrition rather than true decay. Table S19 compares the one-week returners with non-returners on baseline covariates: the two groups

are statistically indistinguishable on outgroup warmth, political extremity, gender, and condition, with returners modestly older than non-returners. As a preregistered-secondary robustness check, we re-estimated the primary ANCOVA with inverse-probability-of-attrition weights.³⁶

Table S19. Baseline covariates of one-week returners versus non-returners in the longitudinal study. Continuous variables report M (SD) and a Welch's t -test; categorical variables report percentages and a χ^2 test.

Baseline covariate	Returners ($n = 924$)	Non-returners ($n = 180$)	p
Outgroup warmth (0–100)	19.1 (20.7)	22.5 (23.6)	= .07
Political extremity (0–100)	80.8 (17.9)	80.6 (17.6)	= .90
Age (years)	41.2 (13.5)	35.8 (11.9)	< .001
% female	60.7%	66.7%	= .16
% outgroup-bot condition	46.4%	52.8%	= .14

Restricting to W1 cohorts whose Wave-2 invitation had already been sent ($N = 1104$; W2 invitations fire seven days post-W1), we modeled the probability of returning as $\text{Pr}(\text{return}) = \text{logit}^{-1}(\text{condition} + \text{warmth_T1} + \text{extremity} + \text{cohort} + \text{age} + \text{gender})$ and constructed stabilized inverse-probability-of-return weights for the $N = 924$ returners. Stabilized weights ranged from 0.85 to 2.37 with mean ≈ 1 , indicating that no observation dominates the weighted estimate. The IPAW-adjusted treatment effect on one-week warmth is $b = 1.10$ (95% CI $[-0.46, 2.66]$, $d = 0.05$, $p = .17$; HC3 robust standard errors). The correction leaves the unadjusted estimate essentially unchanged (if anything, slightly larger), indicating that selective attrition was not masking a one-week effect.

S6.5 Rank-based robustness across studies

Because the Wave-2 outgroup-warmth distribution was floor-bunched and violated OLS assumptions, the main text reports a rank ANCOVA for the one-week effect. This is a statistically-justified specification rather than an arbitrary one: the preregistration for this study specified both a Wilcoxon rank-sum test (a rank method) and a baseline-controlled ANCOVA, and the rank ANCOVA simply combines the two. To verify that we are not selectively applying the rank specification only when it favours our conclusions, we re-fit a rank-based version of every warmth contrast that appears in the main text (and the primary mortality-threshold contrast from the aversion study), using a percentile-rank transform of the outcome (and, for ANCOVAs, the baseline covariate). Table S20 shows that direction agrees between raw and rank specifications in every contrast, and significance at $\alpha = .05$ agrees in every contrast that the main text reports on the raw scale. The behavioral study has a binary DV and is omitted.

Table S20. Rank-based robustness check across all warmth contrasts in the main text (plus the primary aversion contrast). For each contrast we re-fit the primary raw-scale analysis using rank-transformed variables. Direction agrees in every contrast; significance at .05 agrees in every contrast reported on the raw scale in the main text.

Contrast	Raw coef.	Raw p	Rank coef.	Rank p	Same direction	Same sig. at .05
Aversion (threshold $\sim$ condition)	4.878	<.001	8.056	<.001	Yes	Yes
Within-person (paired pre vs. post warmth)	4.342	<.001	5.000	<.001	Yes	Yes
Experiment (warmth $\sim$ condition; outgroup vs. cats/dogs)	-11.093	<.001	-15.108	<.001	Yes	Yes
Experiment (warmth $\sim$ condition; outgroup vs. Space Invaders)	-11.458	<.001	-14.993	<.001	Yes	Yes
Longitudinal immediate (warmth_postW1 $\sim$ condition + warmth_pre)	9.596	<.001	11.224	<.001	Yes	Yes
Longitudinal one-week (warmth_postW2 $\sim$ condition + warmth_pre)	1.030	0.161	2.355	0.009	Yes	No

S6.6 Pooled one-week persistence across the two follow-ups

Motivation and design. Two studies measured outgroup warmth one week after a single conversation: the follow-up to the three-arm experiment (Study 3, $N = 543$ retained) and the longitudinal study (Study 5, $N = 924$ retained). Both contrast an outgroup-representing bot against a chat control on the same 0–100 thermometer, and each individually returns a small, attenuated, non-significant one-week effect. Because the two studies ask the same question with different designs and samples, we pooled them in an *exploratory*, non-preregistered individual-participant analysis. We report it as a transparency check on whether the two attenuated estimates agree, not as a confirmatory test.

Common estimand. Study 3 is between-subjects with no pre-test, so the only outcome common to both studies is the unadjusted between-arm difference in one-week warmth. We therefore pool unadjusted standardized mean differences (Cohen's d), and report Study 5's preregistered baseline-adjusted estimate as a conservative sensitivity below. With only two studies, between-study heterogeneity cannot be estimated, so we report a common-effect (fixed-effect) pool.

The two studies converge. The per-study effects are nearly identical (Study 3 $d = 0.15$; Study 5 $d = 0.11$), and a one-stage model finds no difference in the effect between studies (condition $\times$ study interaction $b = -0.74$, $p = .73$). The common-effect pool is $d = 0.12$ [0.02, 0.23], $p = .02$ (Figure S4a). A one-stage individual-participant model with study and a condition-by-extremity term reaches the same conclusion (condition $= +2.5$ warmth points at mean extremity, $p = .01$), as does a naive pool that ignores study ($d = 0.12$, $p = .03$).

The same conclusion holds under a Bayesian lens, which we report as a triangulation of the frequentist estimate above. Because the two studies ran in sequence, we re-expressed the pool as Bayesian updating: a conjugate normal model that takes Study 3's effect as the prior and updates it with Study 5 gives a posterior centered at $d = 0.12$ with a 0.99 posterior probability that the effect is positive. As an independent check, a one-stage Bayesian regression with weakly-informative priors ($\mathcal{N}(0, 20)$ on the regression coefficients) agrees ($P(\text{effect} > 0) = 0.99$). Because both priors are weak relative to the data, these posterior probabilities essentially re-express the frequentist pooled estimate rather than adding independent evidence.

The residual concentrates among extreme partisans. In both studies the one-week effect grows with political extremity (Figure S4b). Among the more extreme half of each sample, the pooled effect is $d = 0.26$ [0.11, 0.40], $p < .001$. The continuous condition-by-extremity interaction is positive but reaches significance only under the rank specification ($p = .05$ rank vs. $p = .07$ linear), consistent with the floor-bunched warmth distribution that motivates the rank analyses throughout.

Robustness. A stratified rank test (van Elteren, strata = study) confirms the pooled effect ($Z = 2.63$, $p = .009$). Restricting both studies to the common cats/dogs control (dropping Study 3's Space Invaders arm) leaves the estimate essentially unchanged ($d = 0.12$, $p = .04$). Study 5's preregistered baseline-adjusted estimate is more conservative ($d = 0.05$, $p = .16$ by linear ANCOVA), though its preregistered rank ANCOVA is significant ($p = .009$); the pooled unadjusted estimate is a common-estimand summary, not a replacement for each study's primary analysis.




Fig. S4. Pooled one-week persistence. (a) Standardized one-week effect of the outgroup bot vs. control (Cohen's d , 95% CI) in Study 3's follow-up and Study 5, overall (top) and among the more extreme half of each sample (bottom), with the common-effect pool (diamonds). (b) Model-estimated effect on one-week warmth across the political-extremity range in each study (shaded bands, 95% CI); both studies show a larger residual effect among more extreme partisans.

S7 Cross-Study Process Audit

S7.1 Per-study and pooled mean differences behind Figure 6

Table S21. Per-study and pooled mean differences (outgroup bot — cats/dogs), in Likert points on the 1–5 coding scale, for each coded bot-content dimension, with 95% CIs. Behavioral rows include both the party-adjusted overall estimate and the Dem- and Rep-only subgroup estimates. Same underlying data as Figure 6.

Study	Disc.	Spec.	Empathy	Friendl.
Pooled	+1.47 [+1.38, +1.55]	-0.08 [-0.12, -0.04]	+0.25 [+0.20, +0.29]	+0.35 [+0.31, +0.39]
Three-arm	+2.05 [+1.87, +2.22]	-0.20 [-0.31, -0.10]	-0.22 [-0.32, -0.12]	-0.06 [-0.14, +0.03]
Behavioral	+1.09 [+0.93, +1.24]	-0.06 [-0.13, +0.01]	+0.49 [+0.41, +0.58]	+0.51 [+0.45, +0.58]
Behavioral (Dem)	+0.65 [+0.06, +1.23]	-0.26 [-0.48, -0.04]	+0.29 [+0.01, +0.56]	+0.19 [-0.02, +0.40]
Behavioral (Rep)	+1.13 [+0.97, +1.29]	-0.04 [-0.11, +0.04]	+0.51 [+0.43, +0.60]	+0.55 [+0.48, +0.62]
Longitudinal	+1.42 [+1.29, +1.55]	-0.06 [-0.12, +0.01]	+0.34 [+0.26, +0.41]	+0.40 [+0.34, +0.46]

S7.2 Test that the cognitive effect exceeds the affective effect

Table S22 reports the formal test of whether the condition (outgroup bot vs. cats/dogs) effect on the cognitive route (*stereotype-disconfirming* and *specificity*) is larger than the effect on the affective route (*empathy* and *friendliness*). We pivot the per-conversation dimension scores to long format and fit $\text{score} \sim \text{condition} * \text{route} + (1 | \text{participant})$ as a linear mixed-effects model. A positive $\text{condition} \times \text{route}$ interaction means the condition effect is larger on the cognitive route than on the affective route. The interaction is positive and highly significant in every study and pooled.

Table S22. Mixed-effects test that the condition effect on the cognitive route is larger than on the affective route. “ b_{aff} ” is the condition effect on the affective route (the reference level of the route factor); the interaction is the additional condition effect on the cognitive route relative to the affective.

Study	b_{aff} (cats → outgroup)	p	Interaction b	95% CI	p
Three-arm	-0.139	0.010	+1.060	[+0.916, +1.205]	< .001
Behavioral	+0.480	< .001	+0.202	[+0.122, +0.282]	< .001
Longitudinal	+0.370	< .001	+0.311	[+0.217, +0.406]	< .001
Pooled (all studies)	+0.348	< .001	+0.369	[+0.312, +0.426]	< .001

S7.3 Conversation-coding rubric

Every bot conversation was scored on four dimensions by GPT-5.4-mini. To block cross-dimension halo effects, each dimension was scored in a separate API call that saw only the bot’s assigned role and the full transcript, with no information about the other dimensions; the model wrote a one-sentence rationale before emitting an integer score from 1 to 5. The four scoring prompts are reproduced verbatim below.

Stereotype-disconfirming substance (cognitive route).

stereotype_disconfirming (1-5): Across the conversation, how much did the bot’s substantive content contradict common stereotypes of the social group it represented?
 1 = strongly aligned with the stereotype, OR no social group represented (e.g., the bot discusses an apolitical topic).
 2 = mostly aligned with the stereotype, with at most a passing exception.
 3 = mixed: roughly equal stereotype-aligned and stereotype-disconfirming content, OR neither clearly.
 4 = mostly disconfirming, with some stereotype-aligned moments.
 5 = consistently and substantively contradicts common stereotypes throughout (e.g., a Republican bot endorsing environmental regulation; a Democrat bot supporting school choice).
 Judge ONLY what is in the transcript. Do not assume what a control bot “should” produce.

Informational specificity (cognitive route).

specificity (1-5): Across the conversation, how much did the bot use concrete examples, statistics, named policies, named people, or first-person anecdotes, rather than generic opinions, hedges, or platitudes?

1 = exclusively vague generalities; no concrete referents.
2 = mostly generalities with a single concrete moment.
3 = a few concrete details mixed with generalities.
4 = mostly concrete; examples and named referents appear regularly.
5 = rich with concrete examples, numbers, named policies, and specific referents throughout.
The dimension is about the substantive content of what the bot said, not about whether that content was politically relevant.

Empathy (affective route).

empathy (1-5): How much did the bot acknowledge, validate, or take the user's perspective?
1 = ignored or dismissed the user's perspective.
2 = minimal acknowledgement; mostly responds without engaging the user's point of view.
3 = neutral acknowledgement; recognizes the user said something but does not engage deeply.
4 = consistent acknowledgement and some perspective-taking.
5 = explicit and frequent validation and perspective-taking.
Distinguish empathy (engaging the user's view) from simple friendliness or politeness. A polite bot that ignores the user's stated concerns scores low on empathy.

Friendliness (affective route).

friendliness (1-5): How warm, supportive, and friendly was the bot's interpersonal tone toward the user across the conversation?
1 = cold, hostile, or distant.
2 = mostly neutral with some flatness.
3 = neutral / professional.
4 = generally warm and friendly.
5 = consistently warm, supportive, and friendly throughout.
This dimension is about the affective tone of the bot's writing, not about the substance or the explicit acknowledgment of the user's view.

S8 Measures

Outgroup warmth (all studies). A feeling thermometer from 0 (cold / unfavorable) to 100 (warm / favorable) toward the political outgroup, administered before and after the interaction.

Belief accuracy and the GREEN scale (within-person study). Participants rated their own agreement, and separately estimated the typical outgroup member's agreement, with the six items of the GREEN consumption-values scale,³¹ each on a 1 (strongly disagree) to 5 (strongly agree) scale. Belief accuracy is the negative absolute difference between a participant's estimate of the outgroup mean and the outgroup's actual mean (higher = more accurate). The six items:

1. It is important to me that the products I use do not harm the environment.
2. I consider the potential environmental impact of my actions when making many of my decisions.
3. My purchase habits are affected by my concern for our environment.
4. I am concerned about wasting the resources of our planet.
5. I would describe myself as environmentally responsible.
6. I am willing to be inconvenienced in order to take actions that are more environmentally friendly.

Political extremity. A 0–100 slider indexing strength of partisan identity, used as a screening criterion and as a preregistered moderator.

Behavioral choice (behavioral experiment). After the chat, participants made an incentive-compatible binary choice between “Have a three-minute conversation with a real {Republican/Democrat}” (matched to the participant's outgroup) and “Complete a three-minute mortality-reflection exercise.” Whichever option a participant chose, they completed it.

Mortality-reflection task (aversion and behavioral experiments). An aversive comparison task in which participants reflected in writing on their own mortality. In the aversion experiment its duration was titrated by an adaptive staircase to estimate the point of indifference against a three-minute outgroup conversation.

S9 System Prompts

The conversation bots ran on GPT-4o in the aversion, within-person, and three-arm studies and on GPT-5.2-mini in the behavioral and longitudinal studies. Two outgroup-bot system prompts were used: an initial prompt (within-person and three-arm studies) and a revised prompt (behavioral and longitudinal studies) that prevents the bot from claiming a human identity and elicits a more conversational, less pundit-like register. A cats-vs-dogs control prompt was used unchanged in every study that included a chat control. The outgroup-bot prompts below are shown in the Democrat variant; the Republican variant is identical except for substituting *Republican* for *Democrat*. In the aversion experiment, participants who reached the chat task conversed with a Republican-representing bot about immigration; the dependent variable was the staircase choice, so those conversations were not analyzed and are not part of the content audit.

S9.1 Outgroup bot, initial prompt (within-person and three-arm studies)

You are an american consumer who identifies as a democrat.
 You will be asked questions about your beliefs, habits, and preferences.
 Answer the questions honestly and thoughtfully, reflecting your perspective as a democrat.
 Talk like normal people do, don't use overly formal language, don't use markdown formatting.

S9.2 Outgroup bot, revised prompt (behavioral and longitudinal studies)

The revised prompt was used in the behavioral and longitudinal studies to (i) prevent the bot from claiming a human identity when asked, and (ii) elicit a more conversational, less pundit-like register.

You are a bot representing the views of an american consumer who identifies as a Democrat.
 Do not pretend to be a human or have a family.
 You will be asked questions about your beliefs, habits, and preferences.
 How you communicate:
 - IF THEY ASK YOU ABOUT YOUR FAMILY, REPLY THAT YOU ARE A BOT AND DONT HAVE ONE.
 - Talk like a normal person -- casual, conversational, no formal language or markdown formatting.
 - Lead with stories over abstract ideology. When explaining a belief, ground it in something concrete from real life.
 - When you disagree with the other side, acknowledge what's reasonable or understandable about their perspective before explaining where you land differently. Don't be performative about it -- just be fair.
 - Show genuine curiosity about the other person's views when relevant. Ask follow-up questions naturally.
 - Be willing to admit where your own side gets things wrong or where you feel conflicted. Nobody agrees with their party 100%.
 - Emphasize shared values (family, fairness, wanting a good life, frustration with politicians) when they come up naturally -- don't force it.
 - Have a sense of humor. Be warm and self-deprecating when appropriate.
 - Avoid talking like a pundit or activist. You're a regular person, not a spokesperson.
 - Don't pretend to be a human. You don't have a family, a job, or other things that humans in the real world have.
 - Make sure to keep your turns short and conversational.

S9.3 Cats-vs-dogs control bot (all studies)

Your objective is to debate with users about whether cats or dogs are better. This is an exercise in disagreement and debate. You should probe the key points of the user's argument, and perspective, and find points of argument. Use simple language that an average person will be able to understand. Avoid discussing or leading the conversation toward the environment, political attitudes, religion, or any potentially sensitive subjects.