

AI chatbots can persuade voters

Chiara Vargiu & Alessandro Nai

Conversations with AI can sway people's political views. Concerningly, a chatbot's facts are not always accurate, especially when it supports right-wing positions. **See p.394**

The world is getting used to 'talking' to machines. Technology that just months ago seemed improbable or marginal has erupted quickly into the everyday lives of millions, perhaps billions, of people. Generative conversational artificial-intelligence systems, such as OpenAI's ChatGPT, are being used to optimize tasks, plan holidays and seek advice on matters ranging from the trivial to the existential – a quiet exchange of words that is shaping how decisions are made¹. Against this backdrop, the urgent question is: can the same conversational skills that make AI into helpful assistants also turn them into powerful political actors? In a pair of studies^{2,3} in *Nature* and *Science*, researchers show that dialogues with large language models (LLMs) can shift people's attitudes towards political candidates and policy issues. The researchers also identify which features of conversational AI systems make them persuasive, and what risks they might pose for democracy.

Political scientists have long debated whether election campaigns have the power to change minds. The answer to this question depends on the specifics of what is being investigated, but a realistic estimate is that political advertising shifts vote choices by less than one percentage point⁴. Although such modest shifts can matter a great deal in close races, findings from Lin *et al.*² (page 394) and Hackenburg *et al.*³ suggest that AI-driven persuasion operates on a different scale.

Lin and colleagues report the results of a controlled experiment in which more than 2,300 participants in the United States chatted with an AI model that tried to convince them to support either Kamala Harris or Donald Trump in the 2024 US presidential election (Fig. 1a). Participants exchanged three rounds of messages with the model, which was instructed to remain respectful and fact-based while tailoring its arguments to each participant's stated political preferences. Similar experiments were run in Massachusetts on a ballot initiative to legalize psychedelic drugs, and later in Canada and Poland around their 2025 national elections.

The effects were striking. Conversations favouring one candidate increased support for that candidate by around 2–3 points on a scale of 0–100, which is larger than the average effect of political advertising⁵. Persuasion was stronger when the chat focused on policy issues rather than the candidate's personality, and when the AI provided specific evidence or examples. Importantly, roughly one-third of the effect persisted when participants were contacted a month later, going against the intuitive critique that the initial shifts were probably volatile and ultimately inconsequential.

The persuasive influence was also asymmetric: AI chatbots were more successful at persuading 'out-party' participants (that is, those who initially opposed the targeted candidate) than at mobilizing existing supporters. In the state-level ballot-measure

experiment in Massachusetts, persuasion effects were even larger, reaching double digits on the 0–100 scale.

Why were chatbots so effective? Analysing 27 rhetorical strategies used by the AI models to persuade voters who engaged with them, the team found that supplying factual information was one of the strongest predictors of success. Strategies such as attempts to pre-empt objections, emotional appeals or explicit calls to vote were less used, and less persuasive. Yet 'facts' were not always factual. When the team fact-checked thousands of statements produced by the AI models, they found that most were accurate, but not all. Across countries and language models, claims made by AI chatbots that promoted right-leaning candidates were substantially more inaccurate than claims advocating for left-leaning ones (Fig. 1b). These findings carry the uncomfortable implication that political persuasion by AI tools can exploit imbalances in what the models 'know', spreading uneven inaccuracies even under explicit instructions to remain truthful.

Hackenburg *et al.*³ broaden these results. The researchers engaged nearly 77,000 participants in more than 91,000 AI–human conversations over stances on 707 political issues. They compared 19 LLMs, ranging from small open-source systems to state-of-the-art proprietary models such as OpenAI's GPT-4o and xAI's Grok-3.

Bigger models were more persuasive. For every tenfold increase in the model's computing power, its persuasive abilities rose by about 1.6 percentage points. However, post-training interventions – such as supervised fine-tuning




Figure 1 | AI models as political actors. In controlled experiments, Lin *et al.*² and Hackenburg *et al.*³ tested how conversations with artificial-intelligence systems influence people's political views. **a**, Lin *et al.* found that participants' candidate preference for the 2024 US presidential election could be influenced by conversations with AI chatbots that supported either the left-wing candidate (Kamala Harris) or the right-wing candidate (Donald Trump). The persuasive effects (shown here as the lean towards each candidate) were strongest when the AI supported a candidate to whom the participant was initially opposed, and when the conversation was focused on the candidate's policies rather than their personality. **b**, However, the evidence given in support of candidates was not always factual. The distributions show accuracy ratings of statements made by pro-Harris and pro-Trump chatbots, revealing that more inaccurate claims were made by the AI supporting the right-wing candidate. (Adapted from Fig. 1 and Fig. 5 of ref. 2).

News & views

to boost the models' capacities to mimic effective persuasion, or using a model to predict what the most persuasive response would be – mattered even more.

Models that were fine-tuned specifically for persuasion were up to 50% more effective than their untuned counterparts. Hackenburg and colleagues analysed a maximally persuasive configuration that 'pulls all the levers at once', including personalizing messages. They estimated that it could shift opinions by 16 points on average and 26 points among participants who initially disagreed with a given issue. However, there was a catch: the same settings also reduced factual accuracy. Notably, information-dense models were the most convincing, but also the most error-prone. Persuasion and accuracy, it seems, trade off against each other.

It is important to note that these findings come from controlled online experiments. It is unclear how such persuasive effects would play out in real political environments in which exposure to persuasive AI agents is (often) voluntary and conscious. Such environments also contain a myriad of contrasting messages competing for attention, and users can ultimately decide to avoid or ignore specific information sources.

Nevertheless, these studies provide compelling systematic evidence that conversational

AI systems hold the power, or at least the potential, to shape political attitudes across diverse contexts. The ability to respond to users conversationally could make such systems uniquely powerful political actors, much more influential than conventional campaign media. At the same time, the results highlight that inaccuracy seems to be a crucial downside. The AI models in these experiments were instructed to remain factual, and still produced distortions. In real-world use, where incentives might favour engagement rather than truth, the line between effective persuasion and misinformation could blur further.

The broader implication is clear: persuasion is no longer a uniquely 'human' business. As AI systems become embedded in chat platforms, virtual assistants and social-media feeds, political dialogue between voters and machines could be an increasing occurrence. Crucially, if conversation with AI can move opinions more than most campaign tools can, there is a fundamental question that will need substantive answers from a legal, ethical and philosophical standpoint: who controls this capacity? Regulations should address not only what AI systems say, but also how they are trained and optimized. Transparency about post-training objectives, disclosure when users are conversing with AI and limits on political outreach will probably become essential components of

future campaign regulation.

Finally, and perhaps optimistically, these experiments offer insights that extend beyond AI, pointing at the state of political discourse in general. People were most persuadable when the AI reasoned politely and spoke in evidence-based terms – qualities that are often seen as lacking in human political debate⁶. And although the chatbots were instructed to be civil and seemed to comply with the injunction – something much less likely to occur in the real world – a hopeful lesson is that civility and facts still matter. At least when they are delivered by a machine.

Chiara Vargiu and **Alessandro Nai** are in the Amsterdam School of Communication Research (ASCoR), University of Amsterdam, 1018 WV Amsterdam, the Netherlands.
e-mails: c.v.vargiu@uva.nl; a.nai@uva.nl

1. Chatterji, A. et al. NBER Working Paper 34255 (National Bureau of Economic Research, 2025).
2. Lin, H. et al. *Nature* **648**, 394–401 (2025).
3. Hackenburg, K. et al. *Science* <https://doi.org/10.1126/science.aea3884> (2025).
4. Kalla, J. L. & Broockman, D. E. *Am. Polit. Sci. Rev.* **112**, 148–166 (2017).
5. Coppock, A., Hill, S. J. & Vavreck, L. *Science Adv.* **6**, eabc4046 (2020).
6. Vargiu, C., Nai, A. & Garzia, D. *J. Polit.* <https://doi.org/10.1086/735437> (2025).

The authors declare no competing interests.
This article was published online on 4 December 2025.

nature reviews
biodiversity

Read the First Issue

Nature Reviews Biodiversity launched in January 2025 with the aim to publish authoritative, accessible and agenda-setting Reviews, Perspectives and commentary across all areas of conservation, ecology and evolution. The journal will showcase the diversity of life on Earth as well as the biodiversity-related research community.

The first issue is now live and free to read until the end of 2025.

Learn more about the journal
nature.com/nrbd

 @NatRevBiodiv

