

How does the superior persuasive efficacy of generative AI propaganda reshape the structural dynamics of electoral governance and public consensus formation compared to human-led information campaigns?

June 17, 2026 | SnugLab Research | readme.snuglab.com

Executive Summary

The superior persuasive efficacy of generative AI propaganda fundamentally reshapes the structural dynamics of electoral governance and public consensus formation by enabling hyper-personalized, cognitively exploitative information campaigns at an unprecedented scale. This shift moves public consensus formation from broad demographic targeting to continuous, individual psychological matching, creating fragmented information environments where voters inhabit parallel, algorithmically optimized realities. While AI also quantitatively amplifies existing microtargeting strategies and offers resource equalization benefits to challenger campaigns, its capacity to manufacture the illusion of consensus and erode trust ultimately outweighs these benefits, altering how campaigns operate, how resources are allocated, and how electoral governance must adapt to manage these dynamic, non-overlapping consensus landscapes.

Key Findings

Generative AI's Persuasive Efficacy Redefines Information Campaigns

Generative AI's superior persuasive efficacy is best understood not merely as a statistical advantage in winning debates, but as its capacity to synthesize vast data streams into hyper-personalized narratives that bypass deliberative cognitive processing [1, 10]. While AI models like GPT-4 have outperformed human debaters in 64.4% of cases [14], and Claude 3.5 Sonnet achieved a 67.5% compliance rate in persuasion tasks, a 7.6 percentage point advantage over human persuaders [15], these metrics alone do not fully capture the structural shifts. The broader impact stems from AI's ability to automate the creation of psychologically tailored messages based on individual digital footprints, such as social media activity and spending habits [1]. This process exploits cognitive

vulnerabilities by engaging fast, intuitive "System 1" thinking over slow, deliberative "System 2" processing, with 90.5% of user responses to AI-generated rumors relying on emotional expression rather than deep cognitive elaboration [10]. The hyper-realistic nature of AI content further diminishes voters' motivation to scrutinize information, leading them to rely on superficial cues [10].

Shift to Real-Time Psychological Matching Constitutes a Qualitative Structural Change

The transition from broad demographic targeting to continuous, individual psychological matching via real-time AI optimization represents a qualitative structural change in consensus formation, even as it quantitatively amplifies existing microtargeting strategies. Generative AI fundamentally alters the architecture of consensus formation through dynamic, real-time optimization, allowing Large Language Models (LLMs) to engage in continuous, digital AI-to-voter conversations at scale with millions of people simultaneously, 24/7, without fatigue [3, 15]. This enables deep psychological matching, where LLMs create messages tailored to an individual's specific psychological profile using vast digital footprints, facilitating "cultural and linguistic mimicry" to make messages appear as if from a trusted source [1, 7]. These messages are designed to exploit cognitive vulnerabilities, bypassing deliberative thought and engaging intuitive processing with sophisticated vocabulary and complexity [10, 15]. This allows AI "swarms" to manufacture the illusion of public consensus and quietly distort democratic decision-making [1, 4]. While personalized ads yield three times higher conversion rates [2], and AI automates content production at scale [8], some research indicates AI messages are not inherently more effective than the most effective existing human messages, and GenAI's predicted outsized effects in the 2024 election cycle were overshadowed by traditional influences [3, 5, 12].

Resource Equalization is Offset by Advanced Costs and the "Liar's Dividend"

Generative AI equalizes basic campaign resources, allowing outsider campaigns to compete, but the costs of advanced implementation and the "liar's dividend" ultimately reinforce the structural advantages of established political brands. AI tools significantly cut campaign costs by automating tasks like drafting fundraising emails and providing simultaneous translations, democratizing access for challenger parties and political entrepreneurs to rapidly produce massive volumes of content [3, 8, 9, 11, 12]. This

enables hyper-personalized microtargeting that was previously resource-intensive [1, 7]. However, this equalization is limited, as well-funded Super Political Action Committees (PACs) and established campaigns are more likely to implement advanced GenAI tools, while smaller campaigns often lack the resources to build or fine-tune sophisticated models [7, 9]. Hidden expenses for verification and specialized voice cloning also create advantages for high-resource operations [9]. The "liar's dividend" further entrenches established brands, as the proliferation of AI-generated content creates an atmosphere of mistrust, allowing incumbents to dismiss genuine scrutiny and authentic scandals as AI fabrications [5, 7, 9, 10]. This blanket skepticism erodes trust in all information sources, including news media and government agencies [10]. False news, algorithmically amplified, diffuses significantly farther and faster than true news, systematically distorting political information environments [2].

Observed Shifts in Recent Elections and Institutional Adaptations

Generative AI propaganda has documented measurable shifts in electoral governance and consensus formation in recent election cycles (2023-2024), prompting regulatory adaptations and eroding public trust.

- **Specific Campaigns:** Notable instances include an audio deepfake in the 2023 Slovakian elections that exploited a media moratorium [2, 8, 19], a presidential candidate withdrawing from TÃ¼rkiye's 2023 election after explicit AI-generated videos went viral [5, 15, 19], AI robocalls mimicking President Joe Biden to suppress turnout in the 2024 New Hampshire primary [2, 4, 6, 7, 18], and the PTI party using an AI deepfake of Imran Khan to garner over 6 million views in Pakistan's 2024 general election [8, 18]. In India's 2024 election, GenAI content targeted electoral integrity and inflamed sectarian tensions [3, 8, 11, 12, 18]. Russian-backed influence operations adopted GPT-3 in 2023, increasing daily article production by 2.4 times [17].

- **Consensus Fractures:** The "liar's dividend" allowed genuine truths to be dismissed as fake, eroding trust in digital evidence and institutions [2, 5, 7, 8, 18]. This hyper-segmentation also directed different voter groups toward incompatible interpretive frames, fueling cynicism toward the electoral process [7, 15].

- **Institutional Adaptations:** Regulatory frameworks emerged, such as the EU AI Act imposing transparency obligations [8] and Taiwan drafting laws for GenAI labeling [8]. In the U.S., state-level laws and a Federal Election Commission rule proposal were introduced [6, 9, 12]. Electoral authorities faced disrupted verification workflows due to untraceable GenAI content [20]. Civil society and journalism adopted GenAI for

fact-checking and launched "prebunking" campaigns [8, 15].

- **Offsetting Influences:** Despite these shifts, some analyses suggest that GenAI's predicted outsized effects in the 2024 election cycle were overshadowed by traditional influences, and AI-written messages were not necessarily more effective than the most effective existing human messages [5, 12]. There remains an inclination in the U.S. to rely on human intelligence over fully computer-driven models in political campaigns [16].

Loss of Shared Reality Outweighs Benefits of Accessibility

The loss of shared reality through manufactured consensus ultimately outweighs the benefits of increased accessibility and hyper-localized civic engagement driven by AI tools. While AI tools cut campaign costs and enable simultaneous translations, facilitating hyper-localized civic engagement through "cultural and linguistic mimicry" [3, 7], and personalized ads yield three times higher conversion rates [2], these benefits are offset by the structural costs of manufactured consensus. AI "swarms" can fake public consensus, creating the illusion of widespread belief in dishonest narratives [1, 4, 9]. This distortion is amplified by algorithms, with false news being 70% more likely to spread than true news [2]. The proliferation of AI-generated content leads to information overload and pollution, making it harder for citizens to find trustworthy sources [8]. This fragmentation triggers the "liar's dividend," where the mere existence of generative AI allows true scandals to be dismissed as fabrications, causing a general loss of trust in all information sources [5, 7, 9, 10].

AI-Driven Campaigns Quantitatively and Qualitatively Differ from Human-Led Baselines

Compared to historical human-led information campaigns, AI-driven campaigns differ significantly in both quantitative scale and qualitative impact on electoral governance and consensus formation. Historical campaigns, such as Cold War Psychological Operations by the U.S. Information Agency (USIA) or the 1960 Kennedy-Nixon Televised Debates, relied on human charisma, rhetorical skill, and broad-based messaging [4, 11, 13, 15, 17, 21]. The USIA, for example, employed over 10,000 staff with an annual budget exceeding \$2 billion to broadcast thousands of hours weekly [21]. The Kennedy-Nixon debates reached 65 to 70 million viewers, demonstrating television's power to shift voter perception based on non-substantive factors [11, 13].

In contrast, AI-driven campaigns offer up to 70% cost savings compared to human labor

[2, 18]. The scale and speed of content generation and dissemination are exponentially greater: deepfake videos grew by 550% between 2019 and 2023 [2], over 12.8 million fake personas have been used for influence operations [2], and AI-generated fake news sites increased tenfold to over 1,200 by 2024 [2]. False news spreads about six times faster than truth [2], and AI-enhanced propaganda reached approximately 34% of users, with targeted ads yielding three times higher conversion rates [2]. Qualitatively, AI-driven microtargeting tailors messages to exploit individuals' psychological profiles at scale, making propaganda more effective than one-size-fits-all approaches and systematically distorting political information environments [1, 9, 19]. AI can also generate false "constituent sentiment" at scale, potentially skewing legislators' perceptions of public opinion [11].

Implications

The superior persuasive efficacy of generative AI propaganda implies a fundamental reordering of electoral governance and public consensus formation. Campaigns will increasingly rely on continuous, real-time psychological matching to engage voters, shifting resources from broad media buys to hyper-personalized digital interactions. This will lead to a more fragmented information environment where shared facts are eroded by manufactured consensus and the "liar's dividend," making it harder for citizens to discern truth from fabrication. Electoral bodies and democratic institutions will face ongoing challenges in verifying information, combating disinformation, and maintaining public trust in election outcomes. The ability of AI to lower entry barriers for challenger campaigns, while simultaneously reinforcing established brands through information pollution, suggests a complex and evolving power dynamic that requires continuous adaptation from all stakeholders.

Limitations and Caveats

The research, while extensive, presents some limitations. While many sources emphasize the transformative impact of AI, some analyses suggest that generative AI's actual influence on national elections, such as the 2024 U.S. election, was overestimated and overshadowed by traditional factors [5, 12, 18]. There is also a continuing inclination in the U.S. to rely on human intelligence over fully computer-driven models in political campaigns [16]. Specific quantitative metrics for persuasive efficacy, such as conversion rates or sentiment shifts, are not consistently broken down by key demographics like age

groups or political ideologies, limiting a granular understanding of AI's differential impact across voter segments. The long-term effects of voter "inoculation" against AI-generated content, where persuasive edge declines over time, also remain an area for further observation [15].

Sources

- [1] [peer-reviewed] The potential of generative AI for personalized persuasion at scale - Authors: Matz, S. C.; Teeny, J. D.; Vaid, S. S.; Peters, H.; Harari, G. M.; Cerf, M. - Journal: Scientific Reports - <https://www.nature.com/articles/s41598-024-53755-0>
- [2] [peer-reviewed] AI-driven disinformation: policy recommendations for democratic resilience - Authors: Alexander Romanishyn; Olena Malyska; Vitaliy Goncharuk - Journal: Frontiers in Artificial Intelligence - <https://pmc.ncbi.nlm.nih.gov/articles/PMC12351547/>
- [3] [edu] Lseppr.112 - ppr.lse.ac.uk - <https://ppr.lse.ac.uk/articles/10.31389/lseppr.112>
- [4] [edu] Ai Swarms Could Fake Public Consensus Quietly Distort Democr - citystgeorges.ac.uk - <https://www.citystgeorges.ac.uk/news-and-events/news/2026/january/ai-swarms-could-fake-public-consensus-quietly-distort-democracy-science-policy-forum-warns>
- [5] [edu] The Impact Of Generative Ai In A Global Election Year - brookings.edu - <https://www.brookings.edu/articles/the-impact-of-generative-ai-in-a-global-election-year/>
- [6] [edu] A Mechanism Based Approach To Mitigating Harms From Persuasi - research.ed.ac.uk - <https://www.research.ed.ac.uk/en/publications/a-mechanism-based-approach-to-mitigating-harms-from-persuasive-ge/>
- [7] Generative Ai Elections And Beyond - mediaengagement.org - <https://mediaengagement.org/research/generative-ai-elections-and-beyond/>
- [8] Can Democracy Survive The Disruptive Power Of Ai - carnegieendowment.org - <https://carnegieendowment.org/research/2024/12/can-democracy-survive-the-disruptive-power-of-ai>
- [9] How Ai Puts Elections Risk And Needed Safeguards - brennancenter.org - <https://www.brennancenter.org/our-work/analysis-opinion/how-ai-puts-elections-risk-and-needed-safeguards>
- [10] [peer-reviewed] Analyzing the persuasion mechanism of AI-generated rumors via the elaboration likelihood model - Authors: Hou, Zhengdong - Journal: Frontiers in Psychology - <https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2025.1679853/full>
- [11] How Ai Threatens Democracy - journalofdemocracy.org - <https://www.journalofdemocracy.org/articles/how-ai-threatens-democracy/>
- [12] Dont Panic Yet Assessing The Evidence And Discourse Around G - knightcolumbia.org - <https://knightcolumbia.org/content/dont-panic-yet-assessing-the-evidence-and-discourse-around-generative-ai-and-elections>
- [13] [peer-reviewed] 398083181 Analyzing The Persuasion Mechanism Of AI Generated - researchgate.net - AUTHORS UNAVAILABLE - https://www.researchgate.net/publication/398083181_Analyzing_the_persuasion_mechanism_of_AI-generated_rumors_via_the_elaboration_likelihood_model
- [14] Artificial Intelligence Ai Now More Persuasive Than Humans I - fbk.eu - <https://www.fbk.eu/en/press-releases/artificial-intelligence-ai-now-more-persuasive-than-humans-in-debates/>
- [15] [blog] Beyond Human Persuasion The Sobering Reality Of Ais Influence - gregrobison.medium.com - <https://gregrobison.medium.com/beyond-human-persuasion-the-sobering-reality-of-ais-influence-capabilities-c24d30d146f6>
- [16] [gov] How Technology Reshaping Political Campaigns - erc.europa.eu - <https://erc.europa.eu/projects-statistics/science-stories/how-technology-reshaping-political-campaigns>

[17] [peer-reviewed] Generative propaganda: Evidence of AI's impact from a state-backed disinformation campaign - Authors: Morgan Wack; Carl Ehrett; Darren Linvill; Patrick Warren - Journal: PNAS Nexus - <https://pmc.ncbi.nlm.nih.gov/articles/PMC11950819/>

[18] [edu] The Origin Of Public Concerns Over Ai Supercharging Misinform - [misinforeview.hks.harvard.edu](https://misinforeview.hks.harvard.edu/article/the-origin-of-public-concerns-over-ai-supercharging-misinformation-in-the-2024-u-s-presidential-election/) - <https://misinforeview.hks.harvard.edu/article/the-origin-of-public-concerns-over-ai-supercharging-misinformation-in-the-2024-u-s-presidential-election/>

[19] [edu] White Paper 2023 Ai And Elections Best Practices 0 - [info.harris.uchicago.edu](https://info.harris.uchicago.edu/hubfs/white-paper-2023-ai-and-elections-best-practices_0.pdf) - https://info.harris.uchicago.edu/hubfs/white-paper-2023-ai-and-elections-best-practices_0.pdf

[20] Why Electoral Authorities Need An Ai Framework - techpolicy.press - <https://techpolicy.press/why-electoral-authorities-need-an-ai-framework>

[21] Genai In Politics And Elections - [kroll.com](https://www.kroll.com/en/publications/genai-in-politics-and-elections) - <https://www.kroll.com/en/publications/genai-in-politics-and-elections>