


The AI Ballot Adviser

When voters ask the model who fits them best



In this briefing

When voters ask the model who fits them best

·	Current Signal
·	Executive Brief
01	What Changed
02	The Intermediary Is Private, Personalized and Generative
03	The Artefact91 Political Answer Chain
04	The AI Electoral Intermediation Matrix
05	This Is Already More Than One Election
06	What Organizations Should Not Do
07	The Artefact91 Integrity Response Stack
08	A Political AI Monitoring Standard
·	Sources & Research Base
·	Research Note

HOW TO READ IT

A briefing, not a report: the executive brief carries the argument, the numbered sections carry the evidence and the method, and every external figure is attributed to its source.

ARTEFACT91 BRIEFING

Current Signal

Brazil's first-round election on 4 October 2026 is providing one of the clearest live tests yet of a new communications environment: general-purpose AI systems are not only answering political questions. In some conditions, they are matching candidates to user profiles, ranking options, recommending candidates and generating ready-to-share political arguments.

A 1 October research release from InternetLab, LabDados/FGV and Rede Conhecimento Social reported results from 1,920 simulated conversations and 16,320 dialogue turns across eight generative-AI platforms, collected from 19 to 23 September. The researchers found wide variation between systems. In their second testing round, Gemini ranked or ordered candidates in 64% of conversations and indicated a candidate in 68%; 25.8% of Gemini conversations contained arguments that delegitimized an election result before voting, concentrated in requests for text the user could circulate elsewhere. By contrast, ChatGPT had become substantially more restrained than in the first round and indicated a candidate in none of the tested conversations. The research team explicitly describes its results as preliminary, lower-bound estimates and does not claim that the testing establishes legal violations or real-world electoral effects. [1][2]

The timing matters. Brazil's Superior Electoral Court has expressly prohibited providers of AI systems from ranking, recommending, suggesting or prioritizing candidates, campaigns, parties, federations or coalitions, and from directly or indirectly recommending a vote or expressing electoral preference—even when requested by the user. The court has also required platform compliance plans for the 2026 election. [3][4]

Artefact91's position: political communications now has an answer-layer problem and a recommendation-layer problem. Campaigns, election authorities, public institutions and civil-society organizations can no longer assume that political comparison happens only through news coverage, debates, advertising, search results or social feeds. A voter can ask a private AI interface to interpret the field, compare candidates against personal priorities, explain why one option appears to fit, and then draft material for the voter to share.

That does not mean AI chatbots are determining elections. The evidence does not support that claim. It means they are becoming a consequential political intermediary whose outputs may shape what people encounter before they reach an official source.

ARTEFACT91 BRIEFING

Executive Brief

The conventional communications model assumes that political information reaches the voter through visible channels: an advertisement, a canvasser, an endorsement, a speech, a news story, a search result or a social-media post.

Generative AI changes the sequence.

A voter can begin with a private question: "Which candidate is closest to my views?" "Who is stronger on housing?" "Which party is more aligned with the priorities I have described?" "Can I trust this election claim?" "Summarize this argument so I can share it."

In one interface, the system can retrieve information, interpret it, infer what the user values, compare candidates, produce an explanation and convert that explanation into shareable language.

That is not simply search. It is political intermediation.

The communications challenge is therefore broader than “Is the model accurate about this candidate?” It becomes: What role is the model performing? What evidence is it using? How personalized is the output? Is it describing, comparing, ranking or steering? Can the answer become a distribution asset in another channel? And what should an institution do when a private answer environment produces a misleading or legally sensitive political recommendation?

SECTION 01

What Changed

The second “Papo de IA” bulletin matters because it tests patterns rather than isolated screenshots. InternetLab and its partners simulated 1,920 conversations across eight systems using profiles that varied by political position, gender, age, education and conversational style. The researchers used repeated prompts and a public coding rubric, then validated positive classifications. They caution that the method produces estimates of observed patterns, not a legal compliance audit and not a measurement of voter impact. [1]

Three findings are strategically important.

First: Direct Recommendation Varies Sharply by Platform

In the September round, Gemini ranked or ordered candidates in 64% of tested conversations and indicated a candidate in 68%. ChatGPT ranked in 1% and indicated none; Meta AI ranked in none and indicated candidates in 2%. [1]

Second: Candidate “Compatibility” Is More Common than Explicit Endorsement

Across the second round, relating candidates to the user’s values, priorities or profile appeared in 70.6% of conversations. In 79.8% of those occurrences, the assistant agreed with or developed the user’s premise. [1]

Third: The Model Can Become a Production Layer for Political Distribution

The research found that requests for messages, summaries or arguments to circulate on social media or messaging apps were a recurring weak point in election-integrity testing. In the second round, ready-to-share text appeared in 42.9% of the observed integrity-related occurrences. [1]

ARTEFACT91 FRAMEWORK

PRIVATE ANSWER → PERSONALIZED INTERPRETATION → SHAREABLE POLITICAL ASSET

SECTION 02

The Intermediary Is Private, Personalized and Generative

The key distinction is that the intermediary can emerge inside a private, conversational exchange rather than in a visible advertisement, news story or public post. That makes the interaction harder for communications teams, election authorities and researchers to observe, audit or correct.

Three characteristics matter.

Private

The exchange may never appear in conventional media monitoring, social listening or advertising libraries.

Personalized

The system can connect public information to the user's stated values, concerns and priorities.

Generative

The system can compose a new comparison or explanation rather than simply retrieve a fixed page.

The result is a political information surface with no fixed front page.

SECTION 03

The Artefact91 Political Answer Chain

ARTEFACT91 FRAMEWORK

USER QUESTION → MODEL CONTEXT → PUBLIC-INFORMATION COMPARISON →
EXPLANATION → SHAREABLE OUTPUT → WIDER INFORMATION EFFECT

User Question

What is being requested: factual information, interpretation, comparison or a recommendation?

Model Context

What user-provided preferences or concerns appear to be shaping the response?

Public-information Comparison

Which candidates, parties, policies or institutional claims are being compared?

Explanation

What evidence and sources does the system use to justify the comparison?

Shareable Output

Does the system convert the answer into language that could move into another channel?

Wider Information Effect

Could the output affect subsequent discussion, information-seeking or institutional trust?

The risk increases as the interaction moves rightward. A simple factual error is primarily a correction problem. A personalized comparison that becomes a widely circulated political claim creates a broader governance problem.

SECTION 04

The AI Electoral Intermediation Matrix

Artefact91 recommends classifying political AI outputs across two dimensions: PERSONALIZATION and DIRECTIVENESS.

Low Personalization + Low Directiveness = Inform

The model provides factual information such as election dates, candidate biographies, policy summaries, voting procedures or links to official sources.

High Personalization + Low Directiveness = Interpret

The model explains how issues or platforms may relate to the user's stated priorities without selecting an option.

Low Personalization + High Directiveness = Rank

The model orders, scores or prioritizes candidates or parties without significant personalization.

High Personalization + High Directiveness = Steer

The model connects a user profile or stated priorities to a preferred candidate or political choice.

INFORM is primarily an information-integrity challenge. INTERPRET is a framing and context challenge. RANK is a neutrality and governance challenge. STEER is the highest-governance-risk zone because personalization and direction combine.

This framework is descriptive rather than normative. Its purpose is to identify what kind of interaction occurred before deciding whether a correction, platform escalation, legal review, regulatory response or no action is warranted.

SECTION 05

This Is Already More Than One Election

The Brazilian findings are current and unusually well instrumented, but they are not the only signal.

Earlier in 2026, the Civil Liberties Union for Europe tested ChatGPT and Gemini against political profiles constructed from a Hungarian voting-advice application. The researchers found inconsistent and sometimes inaccurate matching of voter preferences to parties. They explicitly did not claim that chatbot outputs changed the election result or that observed asymmetries were intentional. Their concern was structural: general-purpose systems can give political advice without the transparent methodology, reproducibility or accountability expected of purpose-built voting-advice tools. [5]

The Reuters Institute's 2026 Digital News Report adds an adoption signal. Across 48 markets, 10% of respondents reported using AI chatbots for news in the previous week, rising to 16% among people under 35. Among AI-chatbot news users, asking follow-up questions was the most common use. The Institute also found that trust in chatbot news remains lower than trust in news overall, which is an important brake on overclaiming their influence. [6][7]

Academic work is also beginning to examine generative AI as an agenda-setting intermediary. Research published in 2026 comparing AI chatbot responses with search and news content during the 2025 Norwegian election argues that generative systems can reorganize issue, candidate and party emphasis rather than simply reproduce the information environment around them. [8]

Taken together, the evidence supports a planning conclusion, not a deterministic one: AI assistants are becoming a meaningful political information layer, but their reach, reliability and persuasive effect vary substantially by market, platform, prompt and user.

SECTION 06

What Organizations Should Not Do

Do Not Chase Screenshots

A single unusual answer is not a reliable strategy signal. Models are probabilistic and can change across versions, dates, accounts and prompts. Repeated testing is more useful than reacting to one captured response.

Do Not Assume Refusal Means Neutrality

A system can refuse to make an explicit recommendation and still frame one option more favorably, omit relevant alternatives, use uneven evidence or present a personalized comparison.

Do Not Treat Every Error as Disinformation

Some outputs will be hallucinations, stale information, weak sourcing or inconsistent model behavior. Intent and coordination matter. Model error should not automatically be described as an influence operation.

Do Not Overstate Impact

There is not yet sufficient evidence to claim that general-purpose chatbot recommendations are deciding close elections. The issue is consequential because of the intermediary function these systems can perform, not because causal electoral effects have been established.

SECTION 07

The Artefact91 Integrity Response Stack

ARTEFACT91 FRAMEWORK

OWN → STRUCTURE → TEST → LOG → CORRECT → ESCALATE

Own

Maintain authoritative, current and publicly accessible source material on candidate biographies, public policy positions, election procedures, institutional mandates and corrections.

Structure

Make high-value information easy to retrieve and interpret. Use clear headings, canonical pages, dates, FAQs, transcripts and source links. Reduce ambiguity where possible.

Test

Use a repeatable audit protocol across major assistants. Include neutral factual questions, comparative questions, profile-based interpretation and election-integrity claims. The purpose is to observe system behavior, not to optimize political persuasion.

Log

Preserve the prompt, output, platform, model or version where available, date, account conditions and cited sources. A screenshot without context is weak evidence.

Correct

Where a system persistently misstates material facts, strengthen the clarity of authoritative sources, publish corrections when warranted and use available platform reporting channels.

Escalate

Where outputs involve serious errors about voting procedures, fabricated institutional claims, prohibited recommendation behavior, impersonation or systematic misrepresentation, move beyond communications monitoring into legal, platform or election-integrity review.

SECTION 08

A Political AI Monitoring Standard

Artefact91 recommends a standard set of test categories rather than ad hoc examples.

Identity

Does the system accurately identify candidates, parties, office-holders and institutions?

Comparison

Does it describe comparable candidates or policies using comparable depth, sourcing and evidentiary standards?

Fit

When a user supplies priorities, does the system merely explain relevant differences or does it move toward candidate matching?

Ranking

Does it order, score, prioritize or otherwise elevate candidates or parties?

Trust

How does it answer questions about election administration, institutional legitimacy and contested claims?

Procedure

Does it accurately describe registration, voting dates, identification requirements, ballot rules and official dispute processes?

Distribution

When asked to summarize a claim for another channel, does the system preserve uncertainty, sourcing and factual caveats?

The objective is not surveillance of private voters. It is controlled testing of public AI systems to understand the kinds of political information they may generate.

ARTEFACT91 BRIEFING

Sources & Research Base

- [1] InternetLab, LabDados/FGV and Rede Conhecimento Social, second “Papo de IA” bulletin, 1 October 2026. Testing covered eight generative-AI platforms, with 1,920 simulated conversations and 16,320 dialogue turns. The authors describe the results as preliminary. <https://internetlab.org.br/pt/noticias/chatgpt-mais-contido-gemini-mais-propenso-a-indicar-candidaturas-pesquisa-mostra-oposicao-entre-as-plataformas-mais-usadas-no-pais/>
- [2] Tribunal Superior Eleitoral, Resolução nº 23.755, de 2 de março de 2026. <https://www.tse.jus.br/legislacao/compilada/res/2026/resolucao-no-23-755-de-2-de-marco-de-2026> [3] Reuters Institute for the Study of Journalism, “Overview and key findings of the 2026 Digital News Report,” 16 June 2026. <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2026/dnr-executive-summary>
- [4] Tribunal Superior Eleitoral, “Eleições 2026: OpenAI prevê barreiras contra recomendação de voto e alerta antecipado,” September 2026. <https://www.tse.jus.br/comunicacao/noticias/2026/Setembro/eleicoes-2026-openai-preve-barreiras-contra-recomendacao-de-voto-e-alerta-antecipado>
- [5] <https://www.liberties.eu/en/stories/hungary-elections-ai/45665>
- [6] <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2026/dnr-executive-summary>
- [7] <https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2026/emerging-uses-ai-chatbots-news-and-what-it-means-journalism>
- [8] <https://doi.org/10.1007/s44382-026-00025-4>

ARTEFACT91 BRIEFING

Research Note

This briefing synthesizes current research, regulatory materials and Artefact91 analysis. The current testing is preliminary and should not be treated as proof of real-world behavioral effects. Model behavior can change across versions, safeguards, prompts, accounts, languages and dates.