

Federal Trade Commission

---

# Regarding the Proposed Policy Statement Concerning the Suppression of Accuracy in Artificial Intelligence Systems

Docket ID: FTC-2026-0859

## AUTHORS

**Leah Siskind**

*Director of Impact and AI Research Fellow*

**Nidhi Ummettala**

*Research Intern*

Washington, DC  
July 31, 2026

## Introduction

Artificial intelligence is dramatically shaping Americans' daily lives. Consumers rely on AI systems to answer questions, summarize information, provide financial and legal guidance, assist with education, and support medical decision-making. As these systems increasingly mediate how Americans access information, their integrity becomes a consumer protection issue.

The Federal Trade Commission's (FTC) proposal focuses on circumstances in which AI companies intentionally manipulate model outputs or misrepresent their systems' capabilities. However, it overlooks a significant source of consumer harm: adversarial foreign influence embedded throughout the AI development stack. Foreign governments are actively shaping the data, information sources, and repositories from which AI systems are trained and retrieve information.

The FTC should clarify that material omissions regarding the provenance of information and the sources that fundamentally shape AI-generated outputs, including undisclosed foreign influence, may be relevant under Section 5, a consumer protection law that prohibits deceptive or unfair business practices. In particular, the commission should encourage AI companies to disclose the provenance of information that materially shapes AI-generated outputs and clearly label responses that rely on adversarial state-controlled information sources. Doing so would allow consumers to better evaluate the reliability of those outputs, while strengthening the resilience of America's AI ecosystem against foreign influence operations.

## Consumers Increasingly Rely on AI Systems for Important Decisions

The White House's America's AI Action Plan recognizes AI as foundational infrastructure for the nation's future economic competitiveness and public services, reflecting the technology's growing role across nearly every sector of society.<sup>1</sup>

Recent survey data likewise demonstrates the increasing integration of AI in Americans' daily lives. The Pew Research Center found that about half of U.S. adults now use AI chatbots, an increase from only one-third in 2024. Consumers are using smart devices with AI features, AI-powered chatbots, and other generative AI applications with a growing frequency. They often rely on these systems for information and assistance that previously would have come from traditional search engines, or human sources. Beyond general information searches, Americans increasingly use these systems for matters with significant personal consequences, including election information, news consumption, and health guidance.<sup>2</sup> Even within traditional search engine queries, Pew found that 60 percent of U.S. adults read AI search-engine summaries.<sup>3</sup> Research further demonstrates that users are significantly less likely to click on traditional

---

<sup>1</sup> The White House, "Winning the Race: America's AI Action Plan," July 2025.

(<https://www.whitehouse.gov/wp-content/uploads/2025/07/Americas-AI-Action-Plan.pdf>)

<sup>2</sup> Jennifer Medina, "'Who Should I Vote For?' Voters Turn to A.I. Before Casting Their Ballots," *The New York Times*, July 4, 2026. (<https://www.nytimes.com/2026/07/04/us/politics/voters-ai-chatbots-elections.html>)

<sup>3</sup> Jeffrey Gottfried, William Bishop, Monica Anderson, Michelle Faverio, Eugenie Park, and Colleen McClain, "Americans and AI 2026: Chatbots, Smart Devices and Views on Impact," *Pew Research Center*, June 17, 2026.

(<https://www.pewresearch.org/internet/2026/06/17/americans-and-ai-2026-chatbots-smart-devices-and-views-on-impact/#ai-search-summaries>)

search results when an AI summary appears and only click on the sources cited within the summary 1 percent of the time.<sup>4</sup>

As consumers rely on AI-generated answers rather than evaluating the underlying sources themselves, the reliability of those outputs becomes essential. Research has shown that consumers frequently view AI-generated content as equally credible as content produced by humans. One study found that participants rated human-generated and LLM-generated content with similar levels of competence and trustworthiness, while also perceiving AI responses as clearer and more engaging, making users more likely to rely on those responses when making decisions.<sup>5</sup> The influence that generative AI platforms have is reinforced by their communication styles. Users place significantly greater trust in AI systems that communicate with confidence, even when that confidence is unwarranted. The linguistic fluency and authoritative tone commonly exhibited by LLMs can therefore create an impression of reliability that is misaligned with the system's actual capabilities.<sup>6</sup>

These dynamics create important implications for consumer protection. As AI becomes a primary source of information for users, consumers become more vulnerable to undisclosed influences on model behavior. When those influences remain hidden, consumers receive a misleading impression of the objectivity and reliability of AI-generated information. These risks become especially pronounced where foreign influence embedded throughout the AI stack materially shapes AI-generated information outputs without consumers' knowledge.

### **Hidden Foreign Influence Throughout the AI Stack Creates Risks for Consumers**

Foreign governments and other malicious actors are actively influencing the information ecosystem on which AI systems rely, creating a distinct source of consumer harm that the proposal does not address.<sup>7</sup> Countries like China, Iran, Russia, Turkey, and Qatar leverage a variety of tactics to optimize their state-controlled propaganda to influence and shape outputs in LLMs.<sup>8</sup> Rather than targeting only human audiences, authoritarian regimes are optimizing publicly available information to influence how LLMs are trained and what information they retrieve. This practice, often referred to as model grooming, exploits the dependence of frontier AI systems on large volumes of publicly accessible digital content.<sup>9</sup> Unlike

---

<sup>4</sup> Athena Chapekis and Anna Lieb, "Google users are less likely to click on links when an AI summary appears in the results," *Pew Research Center*, July 22, 2025. (<https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results>)

<sup>5</sup> Martin Huschens, Martin Briesch, Dominik Sobania, and Franz Rothlauf, "Do You Trust ChatGPT? -- Perceived Credibility of Human and AI-Generated Content," *arXiv*, September 5, 2023. (preprint arXiv:2309.02524 <https://doi.org/10.48550/arXiv.2309.02524>)

<sup>6</sup> Luise Metzger, Linda Miller, Martin Baumann, and Johannes Kraus, "Empowering Calibrated (Dis-)Trust in Conversational Agents: A User Study on the Persuasive Power of Limitation Disclaimers vs. Authoritative Style," *Association for Computing Machinery*, May 11, 2024. (<https://dl.acm.org/doi/10.1145/3613904.3642122>)

<sup>7</sup> Leah Siskind, "The quiet way AI normalizes foreign influence," *CyberScoop*, January 15, 2026. (<https://cyberscoop.com/the-quiet-way-ai-normalizes-foreign-influence>)

<sup>8</sup> Leah Siskind, "AI-Amplified Narratives: Measuring Propaganda in LLM Citations," *Foundation for Defense of Democracies*, March 3, 2026. (<https://www.fdd.org/analysis/2026/03/03/ai-amplified-narratives-measuring-propaganda-in-llm-citations>)

<sup>9</sup> Pablo Maristany de las Casas, "Investigation | Talking Points: When Chatbots Surface Russian State Media," *Institute for Strategic Dialogue*, October 27, 2025. (<https://www.isdglobal.org/digital-dispatch/investigation-talking-points-when-chatbots->

traditional influence campaigns, model grooming targets the upstream information environment rather than the end user.

A study from the Foundation for Defense of Democracies (FDD) has shown that state-aligned media is overrepresented in American LLM outputs on questions about international conflicts. The reason for this overrepresentation is straightforward. AI systems prioritize content that is abundant, easily accessible, and frequently cited across the web. State-controlled media outlets often possess these characteristics because they publish at high volume, they do not need paywalls, and they widely disseminate their content through media-sharing deals with outlets in developing countries.<sup>10</sup>

As generative AI tools increasingly rely on publicly available content, foreign actors are no longer focused solely on persuading human audiences. They are increasingly optimizing content to shape AI-generated outputs. Research has also shown that countries with low press freedom that censor dissenting opinions have more favorable coverage in LLM outputs.<sup>11</sup> This framing maps perfectly onto U.S. foreign adversaries: Iran, China, and Russia all have strong control over their domestic media, and independent media is virtually nonexistent. The result is that consumers may unknowingly rely on AI-generated responses that reflect systematically biased information environments. When foreign influence materially shapes AI-generated outputs without consumers' knowledge, users cannot meaningfully assess whether a response reflects reliable information, coordinated state narratives, or manipulated datasets.

### **Model Data Can Be Poisoned, and Consumers Are None the Wiser**

These risks are amplified by a lack of transparency: Users rarely know what data AI models are trained on, or which sources inform their outputs. In one study, researchers poisoned a dataset, resulting in AI systems' generation of false conclusions in approximately half of 450 experimental trials while the AI systems detected the manipulation only 6 percent of the time.<sup>12</sup>

Consumers, however, have no practical ability to evaluate whether the data a model is trained on is unbiased. Due to the lack of reporting requirements in government procurement, federal agencies are also unable to inspect training datasets, retrieval pipelines, or other components of the AI development stack. Nor can they determine whether a particular response substantially relied upon adversarial state-controlled information sources or manipulated datasets.

---

[surface-russian-state-media\)](#)

<sup>10</sup> Samantha Custer, Austin Baehr, Bryan Burgess, Emily Dumont, Divya Mathew, and Amber Hutchinson, "Winning the Narrative: How China and Russia Wield Strategic Communications to Advance Their Goals," *Gates Global Policy Center*, November 2022. (<https://docs.aiddata.org/reports/gf01/gf01-04/Winning-the-Narrative-How-China-and-Russia-Wield-Strategic-Communications-to-Advance-Their-Goals.html>)

<sup>11</sup> Hannah Waight, Eddie Yang, Yin Yuan, Solomon Messing, Margaret E. Roberts, Brandon M. Stewart, and Joshua A. Tucker, "State media control influences large language models," *Nature*, May 13, 2026. (<https://www.nature.com/articles/s41586-026-10506-7>)

<sup>12</sup> Balint Gyevnar, Atoosa Kasirzadeh, and Nihar B. Shah, "Distributed Denial of Science: How Indirect Data Poisoning of AI Systems Can Industrialize Scientific Fraud," *arXiv*, July 12, 2026. (<https://arxiv.org/abs/2607.10712>)

This lack of transparency is significant under Section 5. The FTC has long recognized that the omission of material information may mislead reasonable consumers. Information regarding the provenance and governance of AI-generated content may therefore be material where it would influence a user’s decision to rely upon the outputs of an AI system. Other governments and international bodies have similarly recognized the importance of transparency and traceability requirements for AI systems, increasingly incorporating disclosure and documentation obligations into emerging AI frameworks.<sup>13</sup> For example, the Organization for Economic Cooperation and Development’s AI Principles emphasize traceability and documentation throughout the AI lifecycle, recognizing that trustworthy AI depends upon understanding where information originates and how it shapes model behavior.<sup>14</sup>

AI systems already provide contextual disclosures informing users that responses do not constitute legal, medical, or financial advice. Yet they generally do not disclose when outputs substantially rely on information originating from Chinese state media, Russian government-controlled outlets, Iranian state-affiliated sources, or other foreign adversaries. Such disclosures would not require AI companies to determine the truth or falsity of foreign information, nor would they prohibit developers from citing foreign sources. Instead, they would provide consumers with material context about the provenance of information that materially shapes outputs.

## Recommendations

To improve transparency regarding adversarial propaganda in AI-generated outputs and strengthen the resilience of the American AI ecosystem against foreign influence, the commission should:

- **Clarify that material omissions under Section 5 may include undisclosed adversarial influence that materially affects AI-generated outputs.** Consumers increasingly rely on AI systems because they expect trustworthy and objective information. When hidden foreign influence materially shapes model outputs without disclosure, consumers may be misled regarding their reliability or evidentiary basis.
- **Treat record-keeping about data sources as proof of accuracy claims, not as a new reporting requirement.** AI developers routinely advertise that their tools are accurate and reliable. A developer making those claims should therefore be able to show where its major data sources come from and how it decides what shapes its model’s answers. The commission should clarify — through guidance, not rulemaking — that keeping this kind of record is a reasonable practice that supports the accuracy claims companies already make and demonstrates good-faith effort if questions later arise.
- **Clarify that AI companies are required to disclose to consumers when answers rely on adversarial state sources.** When outputs lean heavily on material traceable to an adversarial government, the tool should say so, letting people judge the reliability for themselves. The FTC already applies this logic in its

---

<sup>13</sup> Executive Office of the President, U.S. Office of Management and Budget, Memo, “Increasing Public Trust in Artificial Intelligence Through Unbiased AI Principles,” December 11, 2025. (<https://www.whitehouse.gov/wp-content/uploads/2025/12/M-26-04-Increasing-Public-Trust-in-Artificial-Intelligence-Through-Unbiased-AI-Principles-1.pdf>)

<sup>14</sup> “Transparency and Explainability (Principle 1.3),” *OECD.AI Policy Observatory*, accessed July 17, 2026. (<https://oecd.ai/en/dashboards/ai-principles/P7>)

Endorsement Guides (16 C.F.R. Part 255), which require disclosing hidden connections behind a recommendation. An undisclosed foreign-state origin is the same problem in a new setting.

## **Conclusion**

The commission's proposed policy statement is an important step toward ensuring AI systems remain subject to Section 5. The final statement should also recognize that hidden foreign influence throughout the AI development stack may constitute a material omission when it affects AI-generated outputs relied upon by consumers. By promoting greater transparency and reasonable governance practices, the commission can strengthen consumer trust, help users make more informed decisions, and enhance the resilience of the American AI ecosystem against foreign information operations while preserving innovation.

Thank you for considering our comments. We look forward to seeing how our input is incorporated into the commission's ongoing policy work.