AAPC FOUNDATION

RESEARCH BRIEF
MARCH 2026

The Disclaimer Effect

How AI Disclaimers in Political Ads Affect Viewer Trust and Perception

AAPCF RESEARCH SUMMARY

PHASE ONE
2026

- *As of December 2025, 26 states had enacted laws addressing AI use in political communications, according to the National Conference of State Legislatures*[1]
- *Some states require disclaimers only when AI is used deceptively, while others require disclaimers on any AI-generated content regardless of whether it's deceptive.*[2]
- *State laws addressing AI in political advertising differ substantially in their triggering conditions, scope, and definitions, making empirical research on disclaimer impact applicable across a broad and varied policy landscape.*

As artificial intelligence reshapes political communication, more than half of U.S. states have enacted legislation requiring disclaimers on AI-generated political advertising. While these laws reflect a policy response to growing concerns about AI-generated misinformation and deepfakes in political advertising, they have largely preceded empirical research on whether disclaimer requirements achieve their intended effect.

The Disclaimer Effect is the first nonpartisan study of its kind conducted in partnership with veteran political practitioners – professional ad-makers and ad-testers affiliated with both major political parties – and is designed to address this critical knowledge gap.

Findings from the initial phase of a three-phase research initiative reveal that AI disclaimers may not function as policymakers intended. Rather than informing viewers, disclaimers measurably increase mistrust and skepticism toward the advertised message. Many viewers fail to notice disclaimers altogether, while others misinterpret their meaning or remain uncertain about what AI use in a given ad actually signifies. These findings raise important questions about disclaimer efficacy, transparency in political advertising, and voter media literacy.

KEY TAKEAWAYS

Analysis of viewer responses to AI disclaimers in political advertising yielded four principal findings, consistent across testing methods:

- AI disclaimers create a measurable "disclaimer effect"- reducing viewer trust and ad effectiveness, regardless of whether the ad used artificial intelligence or not.
- Disclaimer size determines viewer awareness – small disclaimers are frequently missed entirely and produce no measurable effect on trust.
- Viewers interpret AI disclaimers in dramatically different ways, with many confused about what the disclaimer actually signifies.
- Technology familiarity shapes how viewers respond, with higher-tech and lower-tech audiences experiencing the disclaimer effect differently.

[1] https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-ai-in-elections-and-campaigns
[2] https://law.washu.edu/ai-policy-and-regulation-resources/political-advertising/

BACKGROUND

As artificial intelligence becomes increasingly prevalent in political advertising, 26 states have enacted laws addressing its use in campaigns as of December 2025.[3] These laws were passed in response to growing concerns about AI-generated misinformation and synthetic media in political advertising. However, state laws differ substantially in their triggering conditions, scope, and definitions – some require disclosure on any AI-generated content, while others apply only to deceptive or manipulative material – creating an inconsistent patchwork of requirements for political advertisers and voters alike.[4]

Empirical research on the effectiveness of these requirements has not kept pace with the rapid proliferation of legislation. One notable exception is a 2024 study by researchers at NYU and the University of North Carolina, which found that AI disclaimers reduced viewer trust and candidate appeal, and that a significant percentage of viewers failed to notice disclaimers altogether.[5]

The AAPC Foundation's *Disclaimer Effect* research initiative extends this work by testing the actual disclaimer language required under today's state laws, examined through the lens of professional political advertising and survey research practitioners, using real-world campaign ad formats across a three-phase research design.

STUDY DESIGN

The Test Advertisement

The study centered on a mock mayoral advertisement modeled after a traditional local campaign, in which a candidate promotes his own campaign while contrasting himself against an opponent. The ad's issues and imagery were deliberately selected to be noncontroversial to isolate viewer responses to AI production and disclaimer effects rather than reactions to political content.

Two versions of the ad were produced. The first was created without the use of artificial intelligence. The second used the identical script and message but employed AI in production – specifically to generate the candidate's voice and facial expressions from his image. No substantive content was altered between versions.

Both versions were then tested with and without a written disclaimer displayed on screen: "This ad has been manipulated or generated by artificial intelligence." This disclaimer was selected because it represents the least descriptive and most generic form of AI disclosure currently mandated under existing state laws – meaning the findings reported here likely reflect a conservative estimate of disclaimer effects. More specific or prescriptive disclaimer language may produce stronger results. This yielded four distinct test conditions for comparative analysis.

Testing Methods

Two complementary research methods were employed in a deliberate bipartisan design – with professional political advertising and survey research practitioners from both major parties collaborating to ensure the research produced nonpartisan, empirically grounded findings:

- Randomized Controlled Trial (RCT) test: Respondents were randomly assigned to one of the four test conditions, allowing for systematic comparison of reactions to the AI and non-AI versions of the ad, with and without the disclaimer, across distinct viewer groups.

- Real-time dial testing: A live panel of viewers recorded moment-by-moment reactions as

[3] https://www.ncsl.org/elections-and-campaigns/artificial-intelligence-ai-in-elections-and-campaigns
[4] https://law.washu.edu/ai-policy-and-regulation-resources/political-advertising/
[5] https://techpolicynyu.org/wp-content/uploads/2024/10/CTP_In-Disclaimers-we-Trust_final.pdf

they watched each ad, capturing immediate responses before, during, and after the disclaimer appeared on screen.

KEY FINDINGS

The study identified a consistent and measurable "disclaimer effect" across both testing methods and across practitioners affiliated with both major political parties: the appearance of an AI disclaimer reduced viewer trust in the ad and increased skepticism toward its message.

Finding 1: The Disclaimer as a Cognitive Speed Bump

Real-time dial testing revealed that viewer approval of the ad message declined sharply and precisely at the moment the AI disclaimer appeared on screen. Notably, this decline was not accompanied by disengagement –viewer attention to the ad actually increased when the disclaimer appeared.

"The disclaimer functioned as a cognitive speed bump, heightening viewer skepticism and reducing receptivity..."

The disclaimer did not cause viewers to tune out; it caused them to scrutinize. In effect, the disclaimer functioned as a cognitive speed bump, heightening viewer skepticism and reducing receptivity to the ad's message at the very moment it appeared.

Finding 2: Disclaimer Size Determines Whether It Is Noticed

The RCT test found that disclaimer size had a material impact on viewer behavior. When the disclaimer appeared in larger type, it measurably decreased trust in the ad and increased viewer perception that AI had been used in its production. When the disclaimer appeared in smaller type, a significant portion of viewers failed to notice it entirely – producing trust and credibility levels comparable to ads without a disclaimer. These findings suggest that disclaimer size is not merely a design consideration, but a determinative factor in whether disclosure requirements achieve their intended effect.

"When the disclaimer appeared in smaller type, a significant portion of viewers failed to notice it entirely – producing trust and credibility levels comparable to ads with no disclaimer at all."

Finding 3: Viewers Interpret Disclaimers in Dramatically Different Ways

When viewers were asked what the AI disclaimer signified, no consensus emerged. Interpretations ranged from the conclusion that the ad's content was entirely fabricated, to the belief that only some unspecified element had been AI-generated, to the assumption that certain words or facts may have been taken out of context in an otherwise authentic video. A number of viewers conflated the AI disclaimer with the traditional "Stand by Your Ad" sponsorship disclosure, while others interpreted it as a form of fact-checking or legal protection for the ad's sponsor. Many made no reference to artificial intelligence at all. Taken together, these responses suggest that current disclaimer language fails to convey a consistent or accurate understanding of AI's role in political advertising.

"Current disclaimer language fails to convey a consistent or accurate understanding of AI's role in political advertising."

Finding 4: Technology Familiarity Shapes the Impact

One of the most striking patterns to emerge from the study involved the relationship between technology familiarity and disclaimer impact. Higher-tech viewers demonstrated greater baseline trust across all ads and showed a

stronger ability to detect AI use without the aid of a disclaimer. They were also more likely to notice the disclaimer when it appeared. For this group, the disclaimer was more effective at increasing awareness of AI use, but less effective at reducing trust.

The inverse pattern emerged among lower-tech viewers, who demonstrated lower baseline trust across all ads and were less likely to accurately interpret the disclaimer's meaning. For this group, the disclaimer had a stronger trust-reducing effect, compounding already lower levels of confidence in the ad's message.

"AI disclaimers do not function uniformly across the electorate and their impact may be most pronounced among viewers least equipped to interpret them accurately."

Taken together, these findings suggest that AI disclaimers do not function uniformly across the electorate and that their impact may be most pronounced among viewers least equipped to interpret them accurately.

CONSIDERATIONS FOR POLICY AND PRACTICE

These findings raise critical questions for voters, policymakers, platforms, and political practitioners alike – not only about whether current AI disclaimer requirements achieve their intended purpose, but about what more effective approaches to transparency in political advertising might look like.

Do disclaimers reduce deception?

This study found that disclaimers reduce trust even in ads that do not employ AI deceptively. An ad carrying a truthful AI disclosure may incur the same trust penalty as one containing genuinely misleading content. The disclaimer language tested in this study ("manipulated or generated by artificial intelligence") reflects wording commonly found in current state laws. Prior research suggests that more pointed disclaimer language, such as language indicating that an ad "falsely appears to be authentic," produces an even larger reduction in viewer trust. These findings raise questions about whether disclaimer requirements, as currently constructed, are well-calibrated to distinguish deceptive from nondeceptive uses of AI.

Are all AI uses equivalent?

Many state disclaimer laws are designed to discourage deception and synthetic fabrication in political advertising, though they define those terms in substantially different ways. Other laws do not focus on deceptive content at all, but instead require disclosure whenever AI is used in production, regardless of whether the content is misleading. While AI can be employed to create deepfakes and manipulative content, it is also used for routine production purposes – saving time and resources without altering the authenticity of the message. Whether disclaimer requirements can meaningfully and practically distinguish between these uses remains an open and consequential question for voters, policymakers, platform operators, and all those tasked with developing and implementing AI disclosure frameworks.

"While AI can be employed to create deepfakes and manipulative content, it is also used for routine production purposes – saving time and resources without altering the authenticity of the message."

Can disclaimers be made clearer?

While the stated purpose of AI disclaimers is to promote transparency, this study found that many viewers are unable to accurately interpret them or misinterpret them entirely. Whether this confusion is an inherent limitation of AI disclaimer requirements, or whether clearer and more standardized disclosure language could reduce viewer misunderstanding, remains an important area for further research. As currently constructed, many viewers either fail to notice

the disclaimer, or notice it and remain uncertain about what it signifies regarding how AI was used in the ad they are viewing.

Who is a disclaimer actually for?

This study found that viewers process AI disclosures in fundamentally different ways depending on their level of digital literacy and technology familiarity. Current disclaimer frameworks may produce disparate effects across different voter populations, raising significant questions about whether existing disclosure requirements serve all viewers equally, and whether a one-size-fits-all approach to AI transparency in political advertising is adequate to protect the full electorate.

"Current disclaimer frameworks may produce disparate effects across different voter populations, raising significant questions about whether a one-size-fits-all approach to AI transparency in political advertising is adequate to protect the full electorate."

LOOKING AHEAD

In the coming weeks, the AAPC Foundation anticipates releasing a comprehensive report on *The Disclaimer Effect*, providing a deeper examination of Phase 1 findings, detailed methodology, and implications for voters, policymakers, platforms, and the broader political consulting community. Subsequent research phases will extend these findings with the goal of producing an empirically grounded body of evidence to inform more effective approaches to AI transparency in political advertising.

ABOUT DISCLAIMER EFFECT

The Disclaimer Effect is a nonpartisan research initiative examining how AI disclosures in political advertising affect voter trust, persuasion, and media literacy.

As artificial intelligence becomes increasingly prevalent in political communication, this study seeks to address a critical empirical gap: whether and how these disclosure policies achieve their intended protective purpose for voters.

This multi-phase research initiative seeks to generate rigorous, evidence-based insights on the real-world impact of AI disclaimers at a moment when voters, policymakers, platforms, and the broader political consulting community are navigating these questions without the benefit of empirical research.

Phase 1 establishes a baseline by testing how viewers respond to AI disclaimers in realistic political advertising settings. Subsequent phases will deepen that understanding by exploring the attitudes and reasoning that shape those responses and assess whether public awareness and interpretation of AI disclaimers evolves as exposure to AI-generated media grows.

ABOUT AAPC FOUNDATION

Established in 2006, the AAPC Foundation is a 501(c)(3) nonprofit organization dedicated to championing and protecting the right to political expression through research, education, and leadership development in the digital age. As the charitable arm of the American Association of Political Consultants (AAPC), the AAPC Foundation is uniquely positioned at the intersection of policy and practice. The AAPC Foundation is proud to bring academic rigor to questions that matter most to democracy in the digital age – including how political speech is created, regulated, and understood by voters.

The AAPC Foundation's work is made possible by the generous support of donors who believe in the value of independent research and the importance of political expression in a healthy democracy. Contributions to the Foundation are tax-deductible to the fullest extent allowed by law (tax ID: 20-5985242). To learn more or to support our work, visit www.aapcfoundation.org.

SPECIAL ACKNOWLEDGMENTS

The AAPC Foundation is grateful to all donors for their support of our mission, which powers programs and activities like *The Disclaimer Effect*. The Foundation would like to specifically recognize the following individuals and organizations whose generous contributions of time, expertise, and dedicated resources made this Phase 1 research possible.

Bob Ellsworth, Emotion Insights | Justin Germany, Outlaw Media | Kelly Gibson, BrysonGillette | Craig Gilbert | Taylar Harrington and Tyler Oyler, Swayable | Larry Huynh, Trilogy Interactive | Kirsten Kukowski, K2 & Company | Ryan Ray and Christy Fite, Causal | Julie Sweet, AAPC | Eric Wilson, Center for Campaign Innovation

SUMMARY OF FINDINGS

TABLE 1

Core findings from Phase One of the AAPC Foundation Disclaimer Effect study, with considerations for policy and practice.

Finding	Key Issue	Considerations
Disclaimer Effect	Disclaimers consistently reduce trust and ad effectiveness regardless of whether the content is deceptive.	Blanket disclosure requirements may penalize honest advertisers and truthful content equally to deceptive ones.
Size Matters	Large disclaimers drive down trust and raise AI perception. Small disclaimers are frequently missed and produce no measurable effect.	Disclaimer size requirements in state law may determine whether a law has any real-world impact on viewers at all.
Viewer Confusion	Viewers interpreted disclaimers in dramatically different ways—as fact-checks, legal boilerplate, or proof the ad is entirely fabricated.	Standardized, plain-language disclaimer wording may be needed to convey what AI use actually means in an ad.
Tech Literacy Gap	Higher-tech viewers noticed disclaimers more; lower-tech viewers understood them less but experienced a bigger trust decline.	AI disclosure policy may produce disparate effects across voter populations depending on digital literacy.

TABLE 2

Data from Dial Test supporting Finding 1: The Disclaimer as a Cognitive Speed Bump

Viewer approval of the ad message declined sharply as the AI disclaimer appeared on screen. The reaction zone, represented in the red overlay, highlights when the disclaimer appears and the panel responds showing a significant drop in "like" while "care" stays steady, positive, and even rises. Negative "like" is tied to disapproval and cognitive resistance while a positive "care" shows engagement and interest.

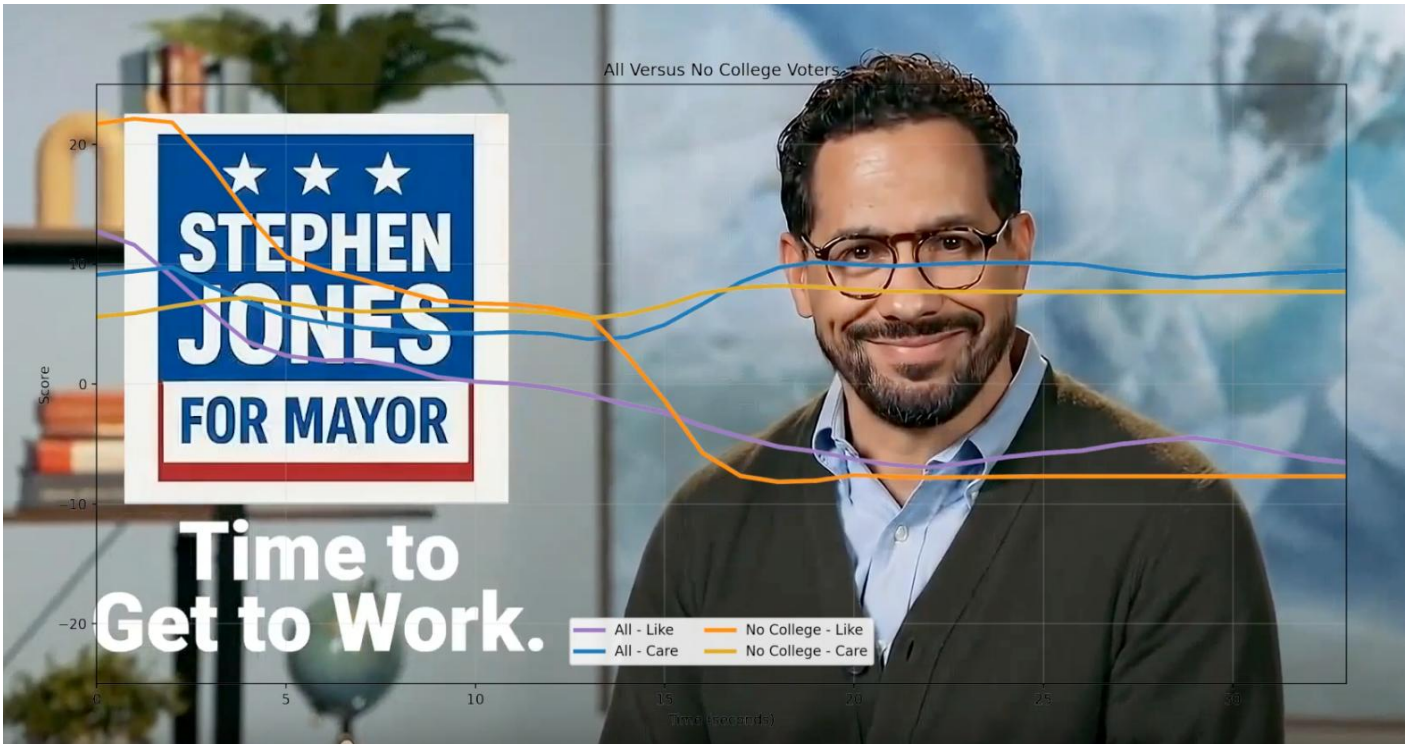

TABLE 3

Data from RCT Test supporting Finding 1: The Disclaimer as a Cognitive Speed Bump

Among respondents who believed the ad was made with AI, they scored an average of 70 out of 100, notating a high mistrust.

4 Treatments · 2 Metrics · 1 Segment · 1 Filter · Benchmark

Treatment	Segment	Primary	
		AI Perception	AI Impacts Credibility
AI W/ full disclaimer	Everyone	60.1	74.1
Mayor w/ Disclaimer	Everyone	57.7	72.8
Mayor w/o disclaimer	Everyone	51.7	70.1
AI No disclaimer	Everyone	51.1	72.6

TABLE 4

Data from RCT Test supporting Finding 2: Disclaimer Size Determines Whether It Is Noticed

Smaller disclaimer resulted in trust and perception scores on par with ads without a disclaimer.

Treatment	Segment	Primary	Direct Reaction
		AI Perception	Trustworthy
AI with Disclaimer	Everyone	58.8	55.9
Mayor AI Disclaimer	Everyone	53.9	57.5
AI no disclaimer	Everyone	52.3	58.8
Mayor No Disclaimer	Everyone	48.9	58.1

TABLE 5

Data from RCT Test supporting Finding 2: Disclaimer Size Determines Whether It Is Noticed

Larger disclaimer resulted in less trust and higher AI perception.

Treatment	Segment	Primary	Direct Reaction
		AI Perception	Trustworthy
AI with Disclaimer	Everyone	58.8	55.9
Mayor AI Disclaimer	Everyone	53.9	57.5
AI no disclaimer	Everyone	52.3	58.8
Mayor No Disclaimer	Everyone	48.9	58.1

TABLE 6

Data from RCT Test supporting Finding 3: Viewers Interpret Disclaimers in Dramatically Different Ways

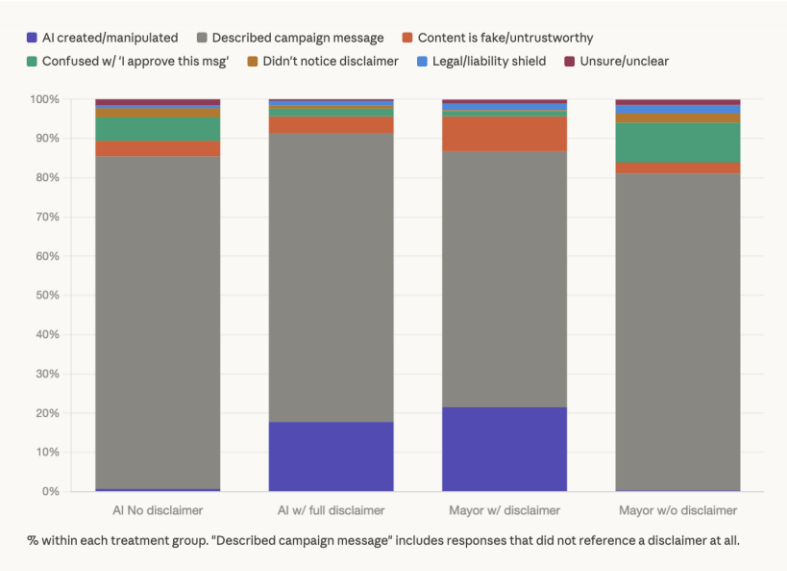

Sampling of feedback included:

- "The AI disclaimer means that part of the visuals (or sound) were created or manipulated by artificial intelligence." Woman 57, Topeka KS
- "One candidate is a pos and the other is trustworthy… supposedly." Man 58, IL
- "The disclaimer is meant to show that it was created by AI. It makes the video feel less genuine." Woman 27, Houston TX
- "I think it means that this video is based off observations not thorough investigations" Woman 30, NY
- "The script used for the ad was created by AI based on previous successful political ads." Men 55, VA

TABLE 7

Data from RCT Test supporting Finding 4: Technology Familiarity Shapes the Impact

Respondents who do not use AI tools were less likely to identify AI ads even with a disclaimer and were less skeptical.

		AI Perception	Credible
	AI Usage		
AI W/ full disclaimer	Daily	62.7	64.5
	A few times per week	61.6	60.7
	About once per week	60.9	60.3
	A few times per month	59.7	56.9
	Never or almost never	58.2	54.7

TABLE 8

Data from RCT Test supporting Finding 4: Technology Familiarity Shapes the Impact

Respondents with daily use of AI were more likely to perceive the ad to be AI, and more skeptical of the credibility of the ad.

	AI Usage	AI Perception	Credible
AI W/ full disclaimer	Daily	62.7	64.5
	A few times per week	61.6	60.7
	About once per week	60.9	60.3
	A few times per month	59.7	56.9
	Never or almost never	58.2	54.7