



The Matter of Disclosure and Transparency of Artificial Intelligence-Generated Content in Political Advertisements

→ Neil Chilson, Head of AI Policy

MB Docket No. 24-211

SUBMITTED | September 19, 2024

The Abundance Institute is a mission-driven non-profit dedicated to creating the policy and cultural environment where emerging technologies can develop and thrive in order to perpetually expand widespread human prosperity. This comment is designed to assist the agency as it explores these issues. The views expressed in this comment are those of the author(s) and do not necessarily reflect the views of the Abundance Institute.

About

Thank you for the opportunity to comment on the Federal Communications Commission’s Notice of Proposed Rulemaking on “Disclosure and Transparency of Artificial Intelligence-Generated Content in Political Advertisements.”¹

The Abundance Institute is a new, mission-driven nonprofit dedicated to creating a policy and cultural environment where emerging technologies can germinate, develop, and thrive in order to perpetually expand widespread human prosperity. Our work on AI includes research, advocacy, testimony before federal and state legislatures, and expert convenings and events.

We also operate the AI Election Observatory (aielectionobservatory.com), where we aggregate and analyze media coverage of uses of AI in the upcoming U.S. election.

I. Introduction

ChatGPT’s release in November 2022 rocketed artificial intelligence into public discussion and spurred a vibrant new sector of easy-to-use consumer-ready, content generation tools. Hundreds of millions of people have used these tools to generate text, images, audio, and video. Right behind these eager users were the regulators, worrying about how these new tools might be misused. Politicians around the country, worried that they might end up the subject of a deepfake video, have proposed a wide range of

¹ Federal Communications Commission, Notice of Proposed Rulemaking, *Disclosure and Transparency of Artificial Intelligence-Generated Content in Political Advertisements*, 89 F.R. 63381 (Aug. 5, 2024), <https://www.federalregister.gov/d/2024-16977> (“NPRM”).

legislation restricting the use of such content in election-related communications.² Some such bills have passed.

The FCC has jumped into this fray with its NPRM. The Commission's ostensible goal is to "provide greater transparency regarding the use of artificial intelligence-generated content in political advertising."³ Unfortunately, the proposed rule is unnecessary, overly broad, and poses significant legal and practical costs that outweigh its intended benefits. The FCC should table the rule; no revisions would suffice to ensure it is balanced, practical, and constitutionally sound.

II. The Unnecessary and Potentially Harmful Singling Out of AI-Generated Content

There is no reason for the FCC to single out "AI generated content" for increased regulation. AI is not new. Since the launch of ChatGPT, AI has risen in prominence in the public mind. But the study and application of AI are almost as old as computing itself, and AI has been involved in content generation for decades. New "generative" AI tools have made it easier for the average person to create certain kinds of sophisticated text, image, audio, and video content. However, there is nothing inherently deceptive about the AI content generation process that somehow tilts the broadcast media environment sufficiently to require regulatory intervention.

² Public Citizen, *Tracker: State Legislation on Deepfakes in Elections*, <https://www.citizen.org/article/tracker-legislation-on-deepfakes-in-elections/>.

³ NPRM ¶ 1.

A. There Is No Evidence of Increased Deception in Broadcast Ads Due to AI

The NPRM offers no concrete evidence or specific data demonstrating that deceptive political advertising is increasing due to AI. It focuses on AI-powered "deepfakes" and other potentially misleading content.⁴

The document expresses concerns about the potential for AI to be used in creating misleading content but does not cite studies or statistics showing an actual rise in deceptive ads. It mentions that the use of AI in political advertising is expected to grow in future election cycles but does not explain how this projection provides evidence of current widespread deception.⁵ It does reference some examples and concerns raised by various sources about the potential misuse of AI in political advertising, but these are largely speculative or isolated incidents rather than evidence of a broader trend.⁶

For instance, the NPRM describes a January 2024 incident involving an AI-generated robocall in New Hampshire that impersonated President Biden.⁷ However, this single incident is poor evidence of a widespread regulatory gap in broadcast political advertising. Indeed, that incident involved telephone calls, not broadcast media, and violated existing New Hampshire laws.⁸

4 NPRM ¶ 10.

5 NPRM ¶ 9.

6 NPRM ¶10.

7 *Id.*, n.46.

8 National Public Radio, *FCC Investigates AI-Generated Deepfake Robocall Targeting Biden in New Hampshire* (May 23, 2024), <https://www.npr.org/2024/05/23/nx-s1-4977582/fcc-ai-deepfake-robocall-biden-new-hampshire-political-operative>.

Our own research suggests that while media coverage of AI in elections has surged, there have been very few actual incidents of expressly generative AI content used in political advertising. Our database of 35,972 media articles contains only four instances of use of generative AI in actual political ads.

Indeed, we can only identify four instances of generative AI content used in a federal electoral campaign during this cycle. On June 5, 2023, DeSantis War Room, a communications arm of the former Ron DeSantis presidential campaign, posted a video on X that showed real clips and audio of President Trump explaining why he didn't fire Anthony Fauci.⁹ Interspersed in the video (between 00:24 and 00:30) was an image collage that included both real and purportedly AI-generated images of President Trump hugging and showing affection to Fauci. The video did not note it included AI-generated or manipulated images. According to NPR reporting, the fact-checking organization AFP detected the fake images two days after they were posted.¹⁰ The post received a community note on X. The story was widely covered by major news organizations.¹¹

9 DeSantis War Room (@DeSantisWarRoom), "Ron DeSantis answers the question: Would you pardon Donald Trump on day one?", X (June 5, 2023, 11:43 AM), <https://x.com/DeSantisWarRoom/status/1665799058303188992>.

10 Vanessa Romo, *DeSantis Campaign Shares Apparent AI-Generated Fake Images of Trump and Fauci*, NPR (June 8, 2023, 8:00 PM), <https://www.npr.org/2023/06/08/1181097435/desantis-campaign-shares-apparent-ai-generated-fake-images-of-trump-and-fauci>.

11 See, e.g., Shane Goldmacher & Maggie Haberman, *DeSantis Campaign Shares Apparent AI-Generated Fake Images of Trump and Fauci*, N.Y. Times (June 8, 2023), <https://www.nytimes.com/2023/06/08/us/politics/desantis-deepfakes-trump-fauci.html>; Andrew Kaczynski & Em Steck, *DeSantis Campaign Video Uses Fake AI Images of Trump*, CNN (June 8, 2023), <https://www.cnn.com/2023/06/08/politics/desantis-campaign-video-fake-ai-image/index.html>; *DeSantis Campaign Shares Apparent AI-Generated Fake Images of Trump and Fauci*, NPR (June 8, 2023), <https://www.npr.org/2023/06/08/1181097435/desantis-campaign-shares-apparent-ai-generated-fake-images-of-trump-and-fauci>; Steve Shepard, *DeSantis PAC Uses AI-Generated Trump Images in Ad*, Politico (July 17, 2023), <https://www.politico.com/news/2023/07/17/desantis-pac-ai-generated-trump-in-ad-00106695>; *Is Trump Kissing Fauci? Apparently Fake Photos Raise AI Ante*, Reuters (June 8, 2023), <https://www.reuters.com/world/us/is-trump-kissing-fauci-with-apparently-fake-photos-desantis-raises-ai-ante-2023-06-08/>.

The second instance was from DeSantis super-PAC Never Back Down. The one-minute ad, released on May 24, 2023, shows DeSantis speaking at an event in Port St. Lucie, Florida on November 5, 2022. The advertisement showed a group of fighter jets flying overhead during DeSantis's speech.¹² Another video of the same event does not show fighter jets flying overhead, suggesting that Never Back Down superimposed the fighter jets in the video with editing tools or generative AI.¹³

The third instance was also created by Never Back Down. They produced a 30-second ad faithfully reproducing the content of a Donald Trump post on Truth Social from July 10, 2023, using an AI-generated voice-over to read the post aloud.¹⁴ The original post by Trump criticized Iowa Governor Kim Reynolds, and the ad—which included Trump's voiceover—was released on July 18, 2023, and ran statewide in Iowa. On January 21, 2024, DeSantis dropped out of the presidential campaign.

The fourth instance was an ad published by the Republican National Committee on April 25, 2023.¹⁵ The 30-second ad painted a dark and scary version of the U.S. if President Joe Biden were reelected, including China invading Taiwan, migrants attempting to cross the U.S. border and soldiers lining the streets of San Francisco. The video includes a disclaimer in the top-left corner: "built entirely with AI imagery."

12 Ana Faguy, *New DeSantis Ad Superimposes Fighter Jets in AI-Altered Video of Speech*, Forbes (May 25, 2023), <https://www.forbes.com/sites/anafaguy/2023/05/25/new-desantis-ad-superimposes-fighter-jets-in-ai-altered-video-of-speech/>.

13 Governor Ron DeSantis, *Governor DeSantis Speaks at 'Don't Tread on Florida' Pit Stop in St. Lucie County*, Rumble (Nov. 5, 2022), <https://rumble.com/v1rsvw-governor-desantis-speaks-at-dont-tread-on-florida-pit-stop-in-st.-lucie-cou.html>.

14 Donald J. Trump (@realDonaldTrump), "... I opened up the Governor position for Kim Reynolds...", Truth Social (July 10, 2023), <https://truthsocial.com/@realDonaldTrump/posts/110690659780399869>.

15 GOP, "Beat Biden," YouTube (Apr. 25, 2023), <https://youtu.be/kLMMxgtxQ1Y>.

Other research shows limited use of AI imagery in political ads. Researchers at Purdue University identified 87 widely circulated deep fake or cheap fake pieces of political content in the U.S. since 2017.¹⁶ None of the identified pieces of content was a political advertisement; most was social media content. Only four pieces of such content were promoted by a politician's account.

B. AI Has Not Materially Affected International Elections

Despite initial fears, AI-generated content has not significantly impacted elections elsewhere in the world. This year has been called a “super-year” for elections because of the large number of important elections happening worldwide – “close to half the world’s population has the opportunity to participate in an election” in 2024.¹⁷ For many observers, the emergence of generative AI in conjunction with this large number of elections was a recipe for disaster.¹⁸ As researchers affiliated with Oxford and the University of Zurich conclude, however,

“With a substantial number of this year’s elections concluded, it is a good time to ask how accurate these assessments have been so far. The preliminary answer seems to be not very; early alarmist claims about AI and elections appear to have been blown out of proportion.”¹⁹

The authors go on to explain those who panicked were mistaken in part because “they ignored decades of research on the limited

16 In The News: Tracking Political Deepfakes: New Database Aims to Inform, Inspire Policy Solutions (last visited Sept. 19, 2024), <https://cla.purdue.edu/news/college/2024/itn-tracking-political-deepfakes.html>.

17 Felix M. Simon, Keegan McBride & Sacha Altay, *AI’s Impact on Elections Is Being Overblown*, MIT Tech. Rev. (Sept. 3, 2024), <https://www.technologyreview.com/2024/09/03/1103464/ai-impact-elections-overblown/>.

18 Valerie Wirtschafter, *The Impact of Generative AI in a Global Election Year*, Brookings (Jan. 30, 2024), <https://www.brookings.edu/articles/the-impact-of-generative-ai-in-a-global-election-year/>.

19 Simon, *et al.*, *supra* n.17.

influence of mass persuasion campaigns.”²⁰ They also note (as we have elsewhere²¹) that the primary bottleneck for misinformation or disinformation campaigns is not the cost of creating persuasive content, but the difficulty in delivering it to the intended audience.²²

The current evidence suggests that elections in the age of generative AI are no more and no less deceptive than before. This alone is a good reason for the FCC to pause this rulemaking effort and monitor developments further before singling out a particular form of content generation for new regulation.

C. Traditional Content Can Create Comparable Harm

Deceptive content created without AI can be equally harmful, questioning the need to single out AI-generated material. The focus on AI-generated content in the proposed rule overlooks the fact that misinformation and deceptive political advertising have long existed using traditional editing techniques. Misleading edits, out-of-context quotes, and manipulated images have been staples of negative political advertising for decades. These conventional methods can be just as effective at deceiving voters as AI-generated content, if not more so, due to their familiarity.

Indeed, the novelty of generative AI content has created a kind of “Streisand effect,” generating outsized coverage around an ad when use of AI is uncovered. (See the above discussion of two

²⁰ *Id.*

²¹ Neil Chilson, *The Integral Role of AI Tools in Modern Political Discourse* (Sept. 27, 2023), Testimony Before the United States Senate Committee on Rules and Administration on AI and the Future of our Elections, https://www.rules.senate.gov/download/06/02/2024/testimony_chilson1 (“One doesn’t need AI to create a deceptive text message or email. The real challenge often lies in distribution rather than content creation, and generative AI doesn’t significantly alter this cost dynamic.”).

²² Simon, *et al.*, *supra* n.17.

ads affiliated with the DeSantis campaign.) For this reason, even if any lone AI-powered ad might appear particularly convincing and deceptive, AI-powered ads overall may be less effective at deceiving voters than more traditional and subtle misleading techniques.

Furthermore, some claimed deep fakes have turned out to be "cheapfakes" - edits to content using traditional means such as slowing video or splicing content to remove context. This phenomenon highlights the difficulty in distinguishing between AI-generated content and skillfully edited traditional media. For example, a video that appears to show a candidate stumbling over words or making an inappropriate statement might be labeled as an AI-generated deepfake, when it could be a slowed-down or carefully edited version of real footage. The term "cheapfake" itself underscores that sophisticated AI technology is not necessary to create misleading content.²³

This blurred line between AI and traditional editing techniques raises several important points:

- 1. Effectiveness of Deception:** Traditional editing methods can be just as effective, if not more so, in creating misleading content. Voters may be more likely to believe slightly altered real footage than entirely AI-generated content.²⁴
- 2. Arbitrary Distinction:** By focusing solely on AI-generated content, the proposed rule creates an arbitrary distinction that may not effectively address the broader issue of deceptive political advertising.

23 Michael Hameleers, *Cheap Versus Deep Manipulation: The Effects of Cheapfakes Versus Deepfakes in a Political Setting*, 36 Int'l J. Pub. Op. Res. 1 (2024), <https://doi.org/10.1093/ijpor/edae004>.

24 *Id.*

3. Potential for Misdirection: The emphasis on AI could divert attention from more prevalent forms of misinformation created through conventional means, potentially leaving voters more vulnerable to these familiar tactics.

4. Enforcement Challenges: Given the difficulty in distinguishing between AI-generated content and skillfully edited traditional media, enforcement of the proposed rule could be problematic and inconsistent.

5. Unintended Consequences: The rule might inadvertently lend more credibility to deceptive content created through traditional means, as the absence of an AI disclosure could be misinterpreted as a sign of authenticity. This “liar’s dividend” is a well-known downside of certain mandatory disclosure regimes.²⁵

By singling out AI-generated content, the proposed rule fails to address the broader spectrum of deceptive practices in political advertising and may inadvertently create a false sense of security among viewers when encountering non-AI manipulated content.²⁶

D. This Proceeding Potentially Undermines Public Confidence in the Electoral Process

This proceeding and proposed rule could undermine its own purpose. The NPRM expresses concern about actions that “creat[e] confusion and distrust among potential voters.”²⁷ Yet this proceeding may

25 Schiff, K.J., Schiff, D.S. & Bueno, N.S., *The Liar’s Dividend: Can Politicians Claim Misinformation to Evade Accountability?*, *Am. Pol. Sci. Rev.* 1-20 (2024), <https://doi.org/10.1017/S0003055423001454>.

26 Note that a more comprehensive approach to promoting transparency and authenticity in political advertising is not an option here. Not only is such an effort well outside of the FCC’s authority, the First Amendment largely protects non-libelous false statements in political speech.

27 NPRM ¶ 9.

create more distrust than it resolves. Unsubstantiated claims about the effect of AI on elections erode public confidence in the process. The more the public hears that AI-manipulated content is being used to deceive them, the less they trust legitimate political messaging, fostering skepticism toward candidates, political institutions, and election outcomes.²⁸ Fears about AI deepfakes also generate a “liar’s dividend,” by strengthening the ability of those who want to disclaim genuine, actual evidence as AI-generated.²⁹

Because the FCC offers no evidence of widespread misuse or significant impact of AI on elections, this proceeding unjustifiably contributes to public anxiety about the integrity of political communications. Instead of enhancing transparency, the result may be to sow confusion and further erode trust in democratic institutions.

In short, the NPRM singles out AI-generated content despite no evidence that such content is being used in deceptive political advertising, no evidence that it poses any real risk to elections, and no evidence that the effects of generative AI content are worse than traditional methods of content creation. Worse, by singling out AI content, the FCC is fueling unwarranted concerns that could ultimately undermine public confidence in electoral processes. This alone provides good reason to abandon this rulemaking.

28 Chesney, Robert & Citron, Danielle Keats, *Deep Fakes: A Looming Challenge for Privacy, Democracy, and National Security*, 107 Calif. L. Rev. 1753, 1785-86 (2019), U. of Tex. Law, Pub. Law Research Paper No. 692, U. of Md. Legal Studies Research Paper No. 2018-21, <https://ssrn.com/abstract=3213954>.

29 *Id.*

III. The Proposed Rule is Overbroad and Vague

The proposed rule reaches far beyond its intended boundaries, due primarily to the overbreadth of the definition of the core term, “AI-generated content.” The overbreadth of that definition creates a rule that likely encompasses most digitally created content, will result in disclosures that are at best useless and at worse deceptive, and will be complicated and burdensome to apply.

This overbreadth is not easily fixable. It is a direct and predictable result of attempting to regulate a vast and diverse suite of technologies known as AI, rather than taking a technology neutral approach to the effects with which the agency is concerned.

A. The Definition of “AI-Generated Content” is Overly Broad

The primary source of the rule’s overbreadth is the definition of “AI-generated content.” which NRPM proposes to define as:

“[A]n image, audio, or video that has been generated using computational technology or other machine-based system that depicts an individual’s appearance, speech, or conduct, or an event, circumstance, or situation, including, in particular, AI-generated voices that sound like human voices, and AI-generated actors that appear to be human actors.”³⁰

This proposed definition of “AI-generated content” suffers from ambiguity and overbreadth, creating uncertainty in its application and enforcement. Several key issues emerge from the definition's language:

³⁰ NRPM ¶ 11.

1. The term "generated" lacks precision. It could include any content merely edited or processed by humans using computational tools, rather than content solely or primarily created by AI. Without specification, "generated" content can be interpreted to include any content that has been manipulated, enhanced, or even simply stored using computational technology. This would sweep in a vast range of digital content.
2. "Computational technology" and "machine-based system" are exceptionally broad, overlapping terms that could each include nearly all modern electronic devices and software, from simple calculators to advanced AI algorithms. (Indeed, a "machine-based system" arguably includes technology as ancient as the Gutenberg press.) The definition fails to specify which types of technology or systems qualify, potentially capturing content created with basic editing software, smartphones, or traditional recording equipment.
3. The phrase "depicts an individual's appearance, speech, or conduct, or an event, circumstance, or situation" covers any visual or auditory representation of people or events. Without further limitations, this could include diverse forms of content from photographs to animations, regardless of realism, intent, or degree of manipulation.
4. By using "including, in particular" when referencing AI-generated voices and actors that mimic humans, the definition expands rather than limits its scope. This phrasing implies these are merely examples within a broader category, indicating that the definition encompasses far more than highly realistic imitations or modifications of content.
5. The definition lacks exclusions or limitations. It does not exempt content created through traditional means or with minimal

computational assistance. Nor does it establish a threshold for AI involvement necessary to qualify as "AI-generated content."

6. Relatedly, the definition ignores the role of human involvement in content creation. It suggests that content primarily created by humans with even a *de minimus* amount of computer manipulation falls under this definition.

These ambiguities and broad interpretations could lead to overinclusive application of the proposed rule, potentially affecting a wide range of content not intended for regulation. The definition's lack of precision raises significant concerns about the practical implementation and enforcement of the proposed regulations.

B. The Result is a Rule That Encompasses Most Digital Content

Modern communications use sophisticated computational techniques everywhere, making it challenging to distinguish between AI-generated content and traditionally produced media under the FCC's proposed definition. The pervasive use of AI in everyday devices means that much of the content captured or created with these tools could inadvertently fall under the rule's purview.

For instance, consider video and images captured by modern cameras and smartphones. These devices employ AI algorithms to enhance image quality and the user experience. Features such as automatic focusing rely on AI to identify and track subjects within a frame.³¹ Scene recognition algorithms adjust camera settings based on the detected environment—be it a landscape, portrait, or night scene—

31 Matthew Saville, *Autofocus Technology is Changing, Here's Why It's Not Just Bells & Whistles Anymore*," SLR Lounge, <https://www.slrlounge.com/autofocus-technology-is-changing-heres-why-its-not-bells-whistles-anymore/>. Some recent cameras even have animal and bird tracking. See, "Everything You Wanted to Know about Autofocus (AF)," Canon, <https://www.canon-europe.com/pro/infobank/autofocus/>.

to optimize image capture.³² Noise reduction techniques use AI to improve low-light photography, and facial recognition features help organize photo libraries.³³ Even simple actions like applying filters or editing images often involve AI-powered software to adjust colors, sharpness, and other attributes.

Similarly, video content is increasingly enhanced using AI-driven stabilization, color grading, and special effects.³⁴ Live video streams might use AI to blur backgrounds or enhance resolution. Audio tracks are cleaned up using AI algorithms that remove background noise or balance sound levels. Under the proposed definition, any political advertisement incorporating such commonplace enhancements could be considered “AI-generated content,” necessitating disclosure.

Because such sophisticated techniques are commonplace, virtually any digitally created or edited content used in a political ad could require an AI disclosure. The rule’s overbreadth thus risks encompassing a vast array of content that poses no threat of misleading viewers, diluting the intended focus on genuinely deceptive practices.

32 Understanding Color Interpolation,” Teledyne Flir, May 11, 2017, <https://www.flir.com/support-center/iis/machine-vision/application-note/understanding-color-interpolation/>; Ron Lowman, “How Cameras Use AI and Neural Network Image Processing,” Synopsys, June 29, 2022, <https://www.synopsys.com/blogs/chip-design/how-cameras-use-ai-neural-network-image-processing.html>.

33 iStock Staff, “How iStock Search Helps You Find the Best Possible Image,” iStock, February 19, 2020, <https://marketing.istockphoto.com/blog/how-istock-search-helps-you-find-the-best-possible-image/>; “Visual Search Powered by Shutterstock.AI,” Shutterstock, <https://www.shutterstock.com/developers/solutions/computer-vision>.

34 TDK, “Electronic Image Stabilization,” <https://invensense.tdk.com/solutions/electronic-image-stabilization/>. Apple’s recent iPhone 15 announcement takes computational photography to a whole new level: the iPhone 15 uses multiple custom neural nets to process images. It’s not an exaggeration to say that every photo taken by an iPhone 15 will be AI generated in part. Jaron Schneider, *Apple Explains What the iPhone 15 Camera Can and Can’t Do – and Why*, PetaPixel, <https://petapixel.com/2023/09/18/apple-explains-what-the-iphone-can-and-cant-do-and-why/>.

C. The Required Disclosure Will Confuse, Exhaust, or Mislead Viewers

Such disclosures will not benefit viewers. Consider the implications for live news coverage of events used in political ads. Modern cameras and microphones used by broadcasters incorporate AI features in challenging environments. If a news outlet covers a political protest or rally occurring at night or in a noisy setting, the equipment's AI functionalities—such as noise reduction, low-light enhancement, and image stabilization—are actively processing the footage in real-time.

Under the proposed rule's broad definition, this live coverage could be deemed as containing "AI-generated content" because the devices employ computational technology to enhance the depiction of individuals and events. Consequently, broadcasters might be required to include an AI disclosure during the live broadcast of such events. This requirement could be both impractical and confusing to the viewers, who might question the authenticity of the live footage solely due to the standard technological enhancements meant to improve the video quality. The overly broad standard may also lead to overuse of disclosures, causing "disclosure fatigue" among viewers.

D. The Proposed Rule is Not Simple to Apply, Contrary to the NPRM's Claims

The NPRM claims that "[t]he proposed definition of AI-generated content is straightforward and simple to apply. Thus, the administrative burden would be modest."³⁵ However, given the deep integration of AI into model technology, political advertisers

35 NPRM ¶ 34.

will be hard pressed to determine when their content contains AI generated content and when it doesn't. Their safest choice will be to always say that it does.

The rule is also made complicated by the need to cover news. For example, what if a broadcast news show wants to cover an online-only ad that contains AI-generated content? Perhaps even to debunk the ad's false claims? If they play a clip of the ad during a news segment, will they be required to report to the broadcaster? Would the broadcaster have to put the disclaimer before the news segment? The ambiguity here could drive broadcasters to avoid airing news segments that include AI-generated content - even to debunk that content - out of concern of running afoul of the disclosure requirements.

E. Resolving the Definitional Issue is Not Easy

Given the problems with the proposed definition, an obvious mitigation would be to refine the definition to better target the specific behavior with which the Commission is concerned. But this is not a straightforward or simple task. There is no consensus definition of "AI." Other statutory definitions are also flawed, at least for the purposes of this proceeding. Narrowing the definition to "generative AI" or "deepfakes" would be linguistically clearer, but legally more vulnerable, because it demonstrates the dilemma at the core of this proceeding: the Commission wants to regulate deceptive speech but cannot.

1. AI is an Evolving Concept With Indistinct Boundaries³⁶

The Commission will not be able to rely on an industry standard for AI, because there is no industry standard. As the NPRM notes, “AI can encompass a wide range of technologies and functions...”³⁷ Indeed, experts have debated the term “artificial intelligence” for decades. The widely used AI textbook by Peter Norvig and Stuart Russell begins by discussing the history of attempts to define AI.³⁸ They describe how the scope of AI is fluid and has included different types of software and algorithms over the decades. Indeed, John McCarthy, who coined the term, remarked that once an AI algorithm works, “we stop calling it AI.”³⁹

Norvig and Russell identify four historical approaches to defining AI: acting humanly, thinking humanly, thinking rationally, and acting rationally. They emphasize “acting rationally” as the prevailing model, defining AI as the study and construction of agents that “do the right thing.”⁴⁰ In other words, after much discussion, Norvig and Russell advocate a functional approach: categorize something as AI or not based, not on its design or nature, but on its uses and actions.

So, what counts as AI today? There is no definitive answer, but it is an expansive category. Norvig and Russell provide an incomplete

36 Significant portions of this section are based on my testimony to the U.S. Senate Rules Committee, *The Integral Role of AI Tools in Modern Political Discourse* (Sept. 27, 2023), Testimony Before the United States Senate Committee on Rules and Administration on AI and the Future of our Elections, https://www.rules.senate.gov/download/06/02/2024/testimony_chilson1.

37 NPRM ¶ 10.

38 Stuart Russell and Peter Norvig, *Artificial Intelligence: A Modern Approach*, 4th ed. (London: Pearson 2021).

39 Moshe Vardi, “Artificial Intelligence: Past and Future,” *Communications of the ACM* 55, no.1 (Jan. 2012): 5.

40 Russell and Norvig, *Artificial Intelligence*, 3-4. The authors further note that this “right action” should align with human benefits; Vardi, “Artificial Intelligence: Past and Future,” 5.

catalog of example AI applications, including such recognition, recommendation algorithms, image understanding, game playing, and medical diagnosis.⁴¹

2. The NPRM Cannot Rely on the Biden Executive Order’s Definition of AI, Because That Definition is Also Overly Expansive

The NPRM points to the following definition of AI set forth in President Biden’s Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence (“Biden EO”):

The term “artificial intelligence” or “AI” has the meaning set forth in 15 U.S.C. 9401(3): a machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations, or decisions influencing real or virtual environments. Artificial intelligence systems use machine- and human-based inputs to perceive real and virtual environments; abstract such perceptions into models through analysis in an automated manner; and use model inference to formulate options for information or action.⁴²

This broad definition includes far more software than the large language models or “generative AI” tools that have generated massive attention since the launch of ChatGPT.⁴³ Indeed, as former Microsoft executive Steven Sinofsky has pointed out, the EO’s AI definition likely covers 1980’s era financial software.⁴⁴ It also appears to

41 Russell and Norvig, *Artificial Intelligence*, 28-30.

42 Exec. Order No. 14,173, § 3(b), 88 Fed. Reg. 74,177 (Oct. 30, 2023), <https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/>.

43 Testimony of Neil Chilson, Hearing on Oversight of Artificial Intelligence, H. Comm. on Oversight & Accountability, 118th Cong. (Mar. 21, 2024, 2024), <https://oversight.house.gov/wp-content/uploads/2024/03/Chilson-Testimony.pdf>.

44 Steven Sinofsky, “211. Regulating AI by Executive Order is the Real AI Risk” (Nov. 1, 2023), <https://hardcoresoftware.learningbyshipping.com/p/211-regulating-ai-by-executive-order>.

cover algorithms used for: social media content moderation and feeds; targeted advertising; search engine algorithms; in-game bots; insurance models; any number of financial tools; and more. In short, the EO defines AI to include a wide range of software.

The EO borrows this definition from the National AI Initiative Act of 2020, a bill intended to boost government agency spending on AI technology.⁴⁵ There, a vague and over-inclusive definition received little attention because it posed little risk: no legal consequences followed for software developers that entirely ignored the National AI Initiative Act.

Here, broadcasters will not have the luxury of ignorance. They and their advertisers will struggle to distinguish between AI and not-AI in the technology they use to create ad content; and the EO definition offers no more help than does the proposed, flawed definition.

3. Limiting “AI-Generated Content” to “Generative AI” or “Deepfakes” Would be Narrower, but Raises New and Serious Issues

One way to revise the definition would be to focus on content from “generative AI.” Of course, this is a somewhat circular definition; what does it mean that “AI-generated content” is content created by generative AI? It shifts the definitional problem to what exactly is “generative AI.” Still, generative AI is a narrower category of AI, and this perhaps could clarify the rule for broadcasters and advertisers. Some state statutes have defined generative AI as AI models that are designed to generate new data resembling human-created content, such as text, images, audio, or video.

⁴⁵ 15 U.S.C. § 9411.

However, this does not solve the difficulty of hybrid content: what degree of generative AI contribution triggers inclusion into the “AI-generated content”? The draft rule appears to apply even for de minimus uses of AI. Does an ad trigger the threshold if the script was written or polished by ChatGPT? If a graphic designer used context-aware fill to erase bystanders from the background of a video?⁴⁶ If the team created a Hmong overdub using a voice translation model?⁴⁷ If the rule incorporates a threshold, how will broadcasters and advertisers apply that threshold? What would it mean for an ad to be “majority” or “substantially” generative AI? These types of thresholds are non-administrable to the point of being arbitrary.

Given this difficulty, one could imagine the FCC taking a different path. The concern driving this matter is deceptive political ads, specifically deepfakes.⁴⁸ Yet the terms “deepfake” or “deceptive” do not appear in the draft rule. For practical reasons, the FCC cannot require broadcasters to ask their advertisers if their content intends to deceive and issue the disclosure if it is. Those trying to deceive people cannot be expected to disclose their intentions. Those who aren’t trying to deceive people shouldn’t

46 *Content-Aware Fill*, Adobe Photoshop Help, <https://helpx.adobe.com/photoshop/using/content-aware-fill.html> (last visited Sept. 18, 2024).

47 Transcript, *AI and the Future of Our Elections*, Hearing Before the S. Comm. on Rules & Admin., 118th Cong. (Sept. 27, 2023), available at <https://www.techpolicy.press/transcript-senate-rules-committee-hearing-on-ai-and-elections/> (Neil Chilson discussing the use of AI to automate language translation in campaign ads in response to a question from Senator Padilla).

48 Federal Communications Commission, Press Release, *FCC Proposes Disclosure of AI-Generated Content in Political Ads* (July 25, 2024), <https://www.fcc.gov/document/fcc-proposes-disclosure-ai-generated-content-political-ads>; NPRM ¶ 9 (“Of particular concern is the use of AI-generated ‘deepfakes’—altered images, videos, or audio recordings that depict people doing or saying things they did not actually do or say, or events that did not actually occur. Such manipulated media could mislead the public about candidates’ assertions or positions on particular issues or about whether certain events actually happened, creating confusion and distrust among potential voters.”)

be hindered. More importantly, the First Amendment prohibits the agency from conditioning broadcaster obligations based on the content of a political communications without satisfying strict scrutiny. A definition of “deepfake” or “AI-generated content” that depends on its deceptive content would trigger strict scrutiny.

The FCC thus faces a definitional dilemma. The agency cannot prohibit broadcasters from carrying deceptive political speech, or even require them to disclose it. Such judgements would clearly be content-based and subject to strict scrutiny. Yet how does one mandate disclosure of certain content without a content-based criteria for what would be covered? “How the content was made” seems to be the only other alternative. And yet, as we have seen, there are no good ways to characterize the content generation process to cover all the desired behavior while excluding other, non-problematic behavior. Thus, the agency is stuck between a rock (violating the Constitution) and a hard place (adopting an arbitrary and capricious rule).

IV. Consequences of an Overly Broad Definition

The proposed rule raises significant problems, including constitutional concerns, unanticipated impact on small entities, and unfavorable benefit-cost analysis. Rather than comprehensively address these, we will focus on how the overly broad definition of “AI-generated content” exacerbates the problems in each of these areas.

A. The Definition Exacerbates First Amendment Issues

The overly broad definition of “AI-generated content” makes the rule vulnerable to a First Amendment challenge that it is not narrowly tailored to target the supposed problem. The regulation

applies to all political ads with AI-generated content, not just deceptive ones. As a result, a wide range of ads likely contain such content, meaning the disclosure will be ineffective in identifying potentially problematic content for viewers or listeners. In addition, the regulation only covers certain FCC-governed broadcast media. It does not apply to online platforms, streaming services, or other digital outlets. Consequently, consumers could see the exact same advertisements, but one medium would include the required disclosure while another would not. Such discrepancies would increase consumer confusion rather than reduce it.⁴⁹

B. The Definition Undercuts the NPRM's IRFA Analysis

The broad definition also undermines the NPRM's Initial Regulatory Flexibility Analysis (IRFA). The Commission expects that "the proposed rules would impose only a modest burden on the affected entities ... because the candidates or entities requesting airtime should be aware of whether the ad which they seek to have aired contains AI-generated content."⁵⁰

However, as discussed above, it may be quite difficult to assess definitively whether a particular ad includes "AI-generated content" under the Commission's broad definition. As a result, the Commission wrongly assumes that the burden on small entities will be "modest."⁵¹

49 See, Statement of Commissioner Brendan Carr, Dissenting at 5 (July 25, 2024), <https://www.fcc.gov/document/fcc-proposes-disclosure-rules-use-ai-political-ads/carr-statement>.

50 NPRM ¶ 56.

51 NPRM ¶ 56.

C. The Definition Undercuts the NPRM's Benefit-Cost Analysis

As the Commission notes, "The benefits and costs of our rules for disclosing AI-generated content depend on the share of political advertisements for which such disclosure would plausibly be required."⁵² As discussed above, the broad definition of AI-generated content means that a large share of political ads will require disclosure. Thus, compliance costs will likely exceed the Commission's current estimates. In addition, the rule imposes other substantial and tangible costs, such as the administrative burden on broadcasters, the potential chilling effect on free speech, and voter confusion due to inconsistent disclosure across different media platforms.

The anticipated benefits may also be significantly smaller than the FCC projects. Deceptive AI-generated content in political ads remains rare, suggesting a low need for disclosures. Furthermore, if the rule results in ubiquitous "disclosure" on every political ad, it will provide little meaningful information to viewers and listeners.

This imbalance suggests that the rule's costs outweigh its purported benefits, based solely on the likely effects of the broad "AI-generated content" definition.

V. Conclusion

The rule, in its current form, poses significant legal and practical challenges that outweigh its intended benefits. There is no evidence that AI is increasing deceptive political advertising. Singling out AI-generated content without clear justification not only overlooks the comparable harm from traditional content

52 NPRM ¶ 35

creation methods but also risks undermining public confidence in our elections. The overly broad definition of “AI-generated content” encompasses a vast array of digitally created content, and the required disclosures would only confuse viewers while imposing undue burdens on broadcasters and advertisers.

Moreover, the proposed rule raises serious constitutional concerns, infringing upon First Amendment rights by compelling speech despite failing to be narrowly tailored.

Considering the concerns outlined above, we urge the FCC to abandon this rulemaking.

Sincerely,

Neil Chilson

Head of AI Policy

Abundance Institute