Political Slopaganda

AI-generated content in the 2026 U.S. election

Scott Babwah Brennen
Anna Viñals Musquera
Center on Technology Policy
New York University

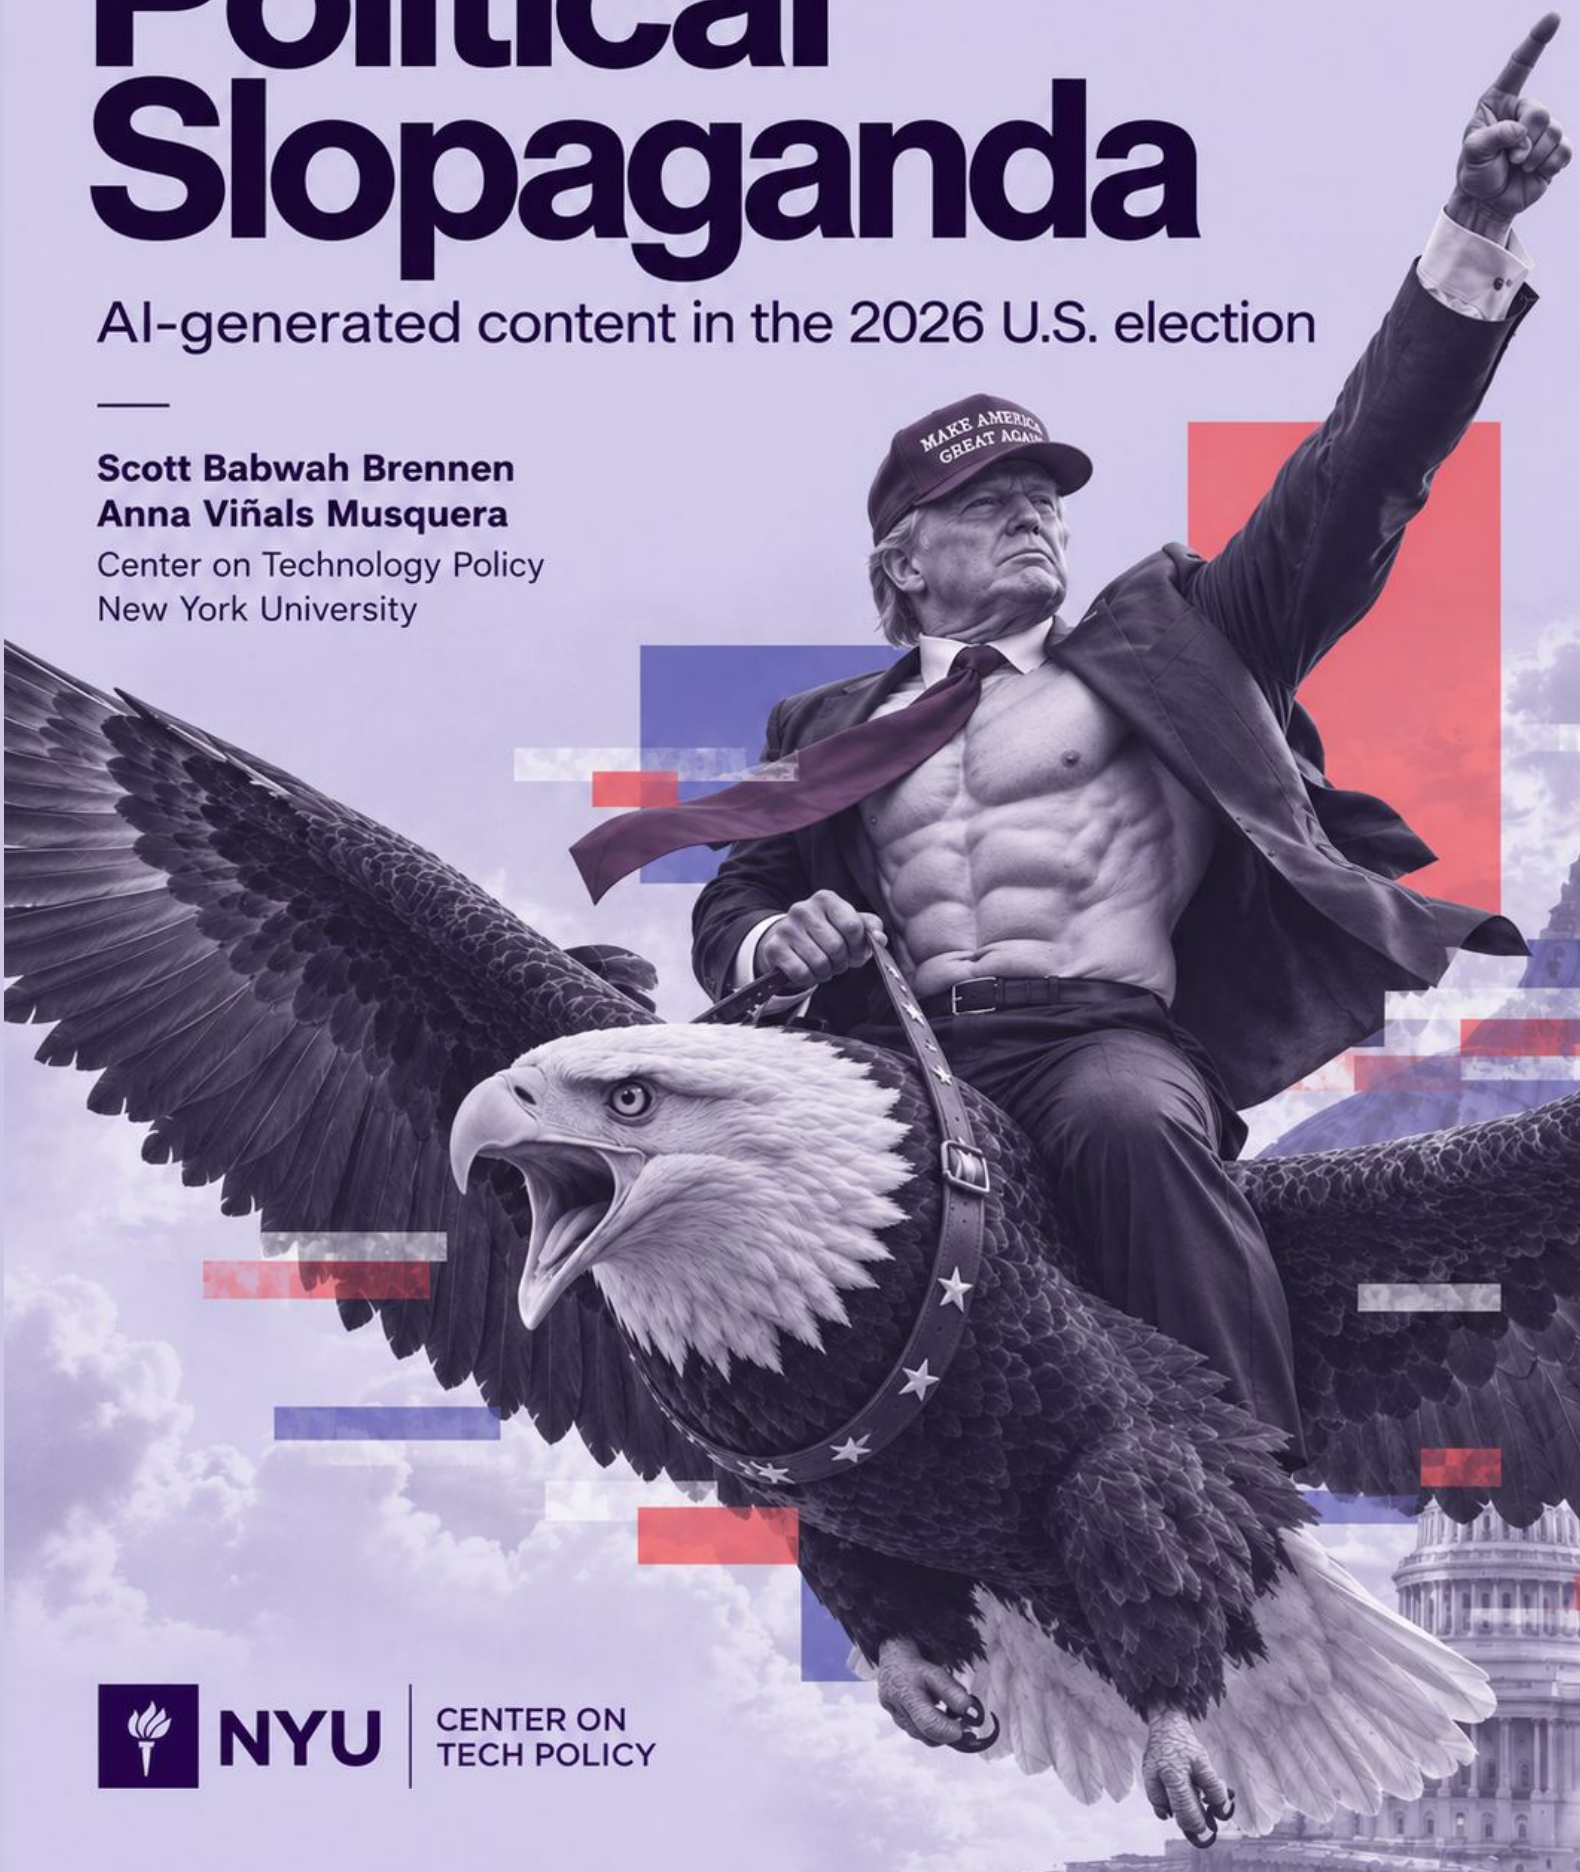

NYU | CENTER ON TECH POLICY

Executive Summary

As they looked forward to the U.S. 2024 presidential election, commentators dubbed it the "AI election," and anticipated a "torrent," a "tidal wave" of deceptive AI-generated content. After the election had ended, Politifact took stock and concluded simply that it was, in fact, "not the AI election." While there were a handful of notable examples of deceptive deepfakes, that threat did not materialize at scale.

Two years later, generative AI has only gotten more advanced and easier to use. Many are again sounding the alarm, predicting a torrent of deceptive deepfakes, and again calling it the "AI election."

However, looking at the 2026 primaries, it is becoming clear that, once again, although there have been very worrisome incidents, deceptive deepfakes are not emerging as a wide-scale, pernicious threat.

Rather, a *different* AI-related threat is developing: political slopaganda. Rather than AI-generated content that is created to look authentic and meant to deceive voters, we are seeing an explosion of AI-generated content that does not appear to be hiding its AI origins. While it may deceive some voters in some instances, this content appears designed to achieve other ends. Rather than deceptive, political slopaganda appears more directed at being humorous, satirical, or attention-grabbing.

Political slopaganda is coming both from the top: from political campaigns or political groups, and from the bottom: from users on social media. It is showing up both in paid political ad spots and in organic content. Some content is meant to criticize candidates, often using extreme ad hominem attacks; other examples are meant to frame candidates in a positive light.

Political slopaganda represents something new; but it also combines three long-running media formats or styles: political ads, AI memes, and satire. Like political ads, political slopaganda is made to support or attack political candidates. Like digital memes, it is designed to be spread virally online and to grab attention. And like satire, despite being factually false, it aims to tell a deeper "inherent truth."

Yet, also like political ads, political slopaganda is likely more effective at mobilizing voters or raising money than it is in impacting candidate selection. Like digital memes, there may be limits to its attentional staying power, especially as we become more and more inundated with AI slop. And rather than shape the deep stories through which people make sense of their lives, it more likely reflects or amplifies them. This does not mean, however, that political slopaganda has no impact: like some repeated misinformation it may further undermine epistemic trust, our trust in knowledge-producing institutions. And like negative campaign ads it may exacerbate pernicious affective polarization, enmity or animosity between political parties.

Against these potential impacts, however, our toolkit of existing interventions is weak. Bans, labels, and likeness protections not only face significant constitutional hurdles, but are designed more for deceptive than satirical content. What good are labels declaring something is AI, when that fact is intentionally obvious?

In contrast, we need interventions that recognize both that political slopaganda is protected political speech and that acknowledge it is less a cause and more a symptom of our faltering political information system. We need interventions that address the infrastructures that breed system-level disorder. Specifically, we need to close loopholes in our campaign finance system to create true transparency into who is funding political content. Similarly, we need to improve the health of our information system by supporting local journalism, prohibiting election disinformation, and building the capacity of audiences of all ages to recognize and make sense of AI-generated content.

There will be no simple solutions to the problems of political slopaganda; instead, we need a concerted effort to rebuild the foundations of our public lives.

What is Political Slopaganda

Many fact-checkers, commentators, and scholars came to the same conclusion that deceptive deepfakes were neither widespread nor particularly impactful in the 2024 U.S. election. One working paper, based on a sample of AI-generated political messages from the 2024 presidential election, found only 10.5% of images were both realistic and negative, what they call "dark fakes."

In contrast, the research found that 78.3% of AI images were *not* realistic. This work highlights that an alternative use of generative AI has been emerging. Rather than realistic deepfakes, we are increasingly seeing inexpensive, rapidly produced, and obviously artificial images: candidates reimagined as superheroes, saints, cartoon villains, or other exaggerated characters. The purpose of this content appears less to deceive, and more to entertain, to attack, and to reinforce political identities. This is not to suggest that such content cannot deceive some audiences; there will always be a wide range of interpretations and understandings that audiences have of content. Yet, what defines this content is less an effort to hide that it is AI-generated, and more to embrace that it is.

At the same time, this obviously AI-generated content derives both from the top and from the bottom: spread directly by politicians or political campaigns and by users on social media. Similarly, we see examples that are negative: vicious, ad hominem attack ads; and examples that are positive and are meant to frame a candidate in the best possible light.

Following existing scholarship highlighting its similarity to AI slop, we refer to this content as political slopaganda. Yet, that term obscures that this content is better described as the combination of three distinct media styles or formats: political ads, digital memes, and political satire or cartoons. Below we highlight examples of current political slopaganda to demonstrate this lineage and use this to identify some of the main goals that producers likely have for this content. Importantly, we do not mean to draw strong lines between political slopaganda that functions as political ads, digital memes, or satire. There is a good argument that many of the examples have characteristics of each of these forms -- indeed that is the point. We break these up only to highlight different aspects of real examples of political slopaganda.

With little systematic empirical analysis, there have been few opportunities to hear from producers or distributors themselves about their intentions for this content. However, recognizing that political slopaganda combines at least three different lineages: political advertising, digital memes, and political satire helps surface three likely intentions or goals producers have. Similarly, we do not yet have a strong evidence base to understand what impact this content may have on elections. However, here we draw on our understanding of the impact of related content to help us anticipate the potential success in achieving what creators intend and anticipate the potential impact this content may have on our elections.

What is Political Slopaganda Trying to Do? And How Successful might it be?

Political Slopaganda supports candidates like political ads

As political messaging, political slopaganda is produced or distributed in order to support a candidate or issue's electoral success. That is, like any other political ad, the primary goal is to help that candidate be elected.

Another ad from Michigan Senate candidate Mike Rogers depicts opponent Abdul El-Sayed hiding in a Trojan horse with Bernie Sanders, Alexandria Ocasio-Cortez and Zohran Mamdani, as it is dragged into the state.

Figure 1: Three AI-generated images from Michigan Senate Candidate Mike Rogers. Source. Source.

Republican candidate for Governor of New York, Bruce Blakeman, has released a series of AI-generated videos attacking Gov. Kathy Hochul. One ad takes the form of a pharmaceutical ad for "Kathy the Political Delusion Pill."

In another, an AI-generated Gov. Hochul describes her accomplishments, including "We provided luxury accommodations for illegal migrants in beautiful hotels, just like this."

The long literature on campaign advertising distinguishes a series of functions that political ads can serve. One common typology recognizes four main goals or purposes for political ads:

- persuasion (persuading voters),
- acquisition (finding new contacts),
- mobilization (convincing supporters),
- fundraising (raising money).

Simply put, there is strong evidence that political ads have little impact on candidate selection. One review of 49 field experiments on political ads concluded that "the best estimate of the effects of campaign contact and advertising on Americans' candidate choices in general elections is zero."

This finding is largely in line with a century of research on media effects, which has broadly concluded that media messages rarely have strong, persistent, persuasive effects. As Altay, Berriche, and Acerbi explain, "Information is not passed from brain to brain like a virus is passed from body to body. When humans communicate, they constantly reinterpret the messages they receive, and modify the ones they send."

Figure 2. AI-generated ad created by NY Gov. candidate Bruce Blakeman. Source.

This is not to say that political ads cannot have meaningful effects. Jessica Baldwin-Philippi concluded that "ad effects are conditional, and generally small," and that ads may be more impactful in "ballot measure and down-ballot elections, earlier in the campaign, and under conditions where persuadable voters are more numerous and information is low." Even if political ads rarely change who we vote for, they "have been much more productive at mobilising action, like getting out the vote and improving donation rates, than at persuasive goals of getting someone to support a candidate."

Similarly, while media messages generally have only conditional persuasive effects, there is strong evidence that they can have more indirect impact, such as by priming viewers to temporarily consider certain views, or setting the agenda for discussion. As one scholar famously put it, even if media do not tell us what to think, they can influence what we think *about*.

The same is likely true for political slopaganda. Even if it is unlikely to convince a voter to support or reject a candidate, it may still do important work introducing or amplifying a frame or topic, or mobilizing supporters to vote or to donate money.

In contrast, even as political ads generally have small persuasive effects, there is evidence that negative political ads can, in some very narrow contexts, have a "backfire effect," where they *hurt* the candidate producing the ad. This backfire effect is more likely when negative ads rely on personal attacks, as opposed to policy attacks, and are seen as less civil or less fair. The effect is also stronger amongst voters who report having less tolerance for political conflict.

While the backfire effect is generally uncommon, it is possible that amid a growing AI backlash, mean, ad hominem attacks that are clearly AI-generated may be more likely to backfire. We saw evidence of this in an experiment we ran on state-mandated labels on AI-generated political ads. We observed a backfire effect on ads that included labels, where those that saw AI-generated ads with labels gave, on average, a lower trust and favorability rating to the producer of ads, than to the candidate attacked.

We may be seeing evidence of this in the 2026 primaries. California state senator Scott Wiener created a “satirical” AI chatbot of his opponent in the primary for Nancy Pelosi’s seat, Connie Chan.

There was widespread criticism of the chatbot, including by Pelosi herself, arguing that the chatbot was sexist, racist, and was in poor taste -- especially given Weiner’s experience sponsoring new AI regulation.

A local news outlet shared an image of graffiti on a billboard Wiener purchased to advertise the chatbot.

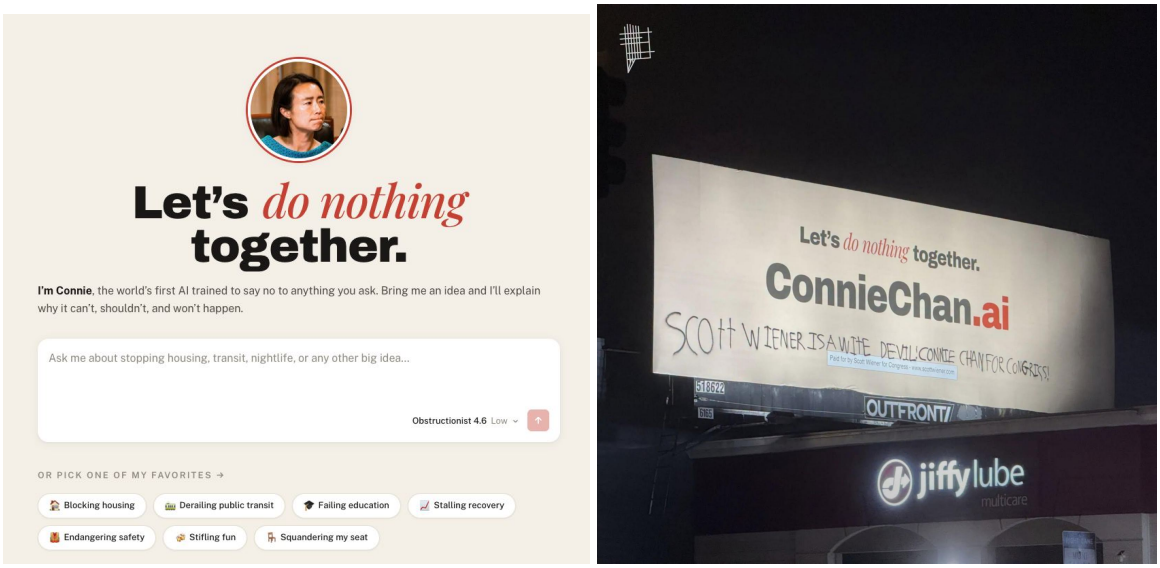

Figure 3. AI chatbot of Connie Chen created by her opponent, Scott Wiener. Source. And fraffiti on a billboard advertising Scott Wiener's chatbot representing his opponent, Connie Chan. Source.

Political Slopaganda seeks attention like digital memes.

Like digital memes, political slopaganda is designed to gain attention within a crowded media environment.

Perhaps the most famous example of political slopaganda was created in support of Spencer Pratt, a reality TV star, running in the Republican primary for Los Angeles Mayor. A filmmaker, Charles Curran, created a series of AI-generated videos. One depicts Pratt as Batman, Joe Rogan as Commissioner Gordon, and prominent Democrats as Batman villains, including current LA Mayor Karen Bass as the Joker.

Figure 4. AI-generated video supporting LA Mayoral primary candidate, Spencer Pratt. Source.

Charles Curran, the filmmaker behind the viral videos, did an interview for The Clay Travis & Buck Sexton Show, a right-wing podcast, in which he explained his thinking behind the ads.

"I knew that it would be popular, I knew that it would be exciting for people because, as I was looking around and kind of getting a gauge of political messaging today I was like, 'oh my gosh, this is all so formulaic. This is all so stale, you know, like it's just a paint by numbers.' I was like, 'let's shake it up a little bit.' I think that the technology is there to really be able to cut through the noise. And I know that it would be effective. I just did not appreciate how effective it would be."

Using extreme, flashy, or sensational content to drive attention in political ads is by no means new. And indeed, much has been written about the role that digital memes play in contemporary life and the contemporary attention economy. Over the last 20 years, we have learned much about the ability of disinformation and other low-quality information to drive engagement online.

At the same time, we've seen hundreds, if not thousands of AI-generated memes depicting prominent politicians, from those showing Donald Trump as a Jedi, to a Reel shared by Oregon primary candidate Jonathan Lockwood showing his opponents as demons attacking children.

Figure 5. AI-generated images. The first shows Donald Trump as a Jedi. The second was created by Oregon primary candidate, Jonathan Lockwood, showing his opponents as demons. Source.

Others have extended this idea to politics, observing that the ability to command public attention is increasingly necessary to gaining and holding political power. Some explain Donald Trump's unique hold on the Republican Party as his mastery of attentional dynamics -- his ability to hold the public's attention, for good or bad. Others have argued that politicians such as Marjorie Taylor Greene, Alexandria Ocasio-Cortez or Zohran Mamdani owe their success, to some degree, to their ability to gain and hold public attention.

For now, some of the extreme political slopaganda seems to be successfully drawing attention--for good or bad, whether it is the Spencer Pratt Batman video or the infamous image of Trump as either Jesus or a doctor (See Figure 6). What remains unclear is the *staying power* that AI-generated political slopaganda will have to continue to drive attention to candidates. As political slopaganda becomes more common, it is very likely that its ability to shape public discussion will diminish. When matched with a growing

AI-backlash, it is possible that AI-generated ads, especially outrageous ones, will have less capacity to break through.

Figure 6: Image shared by President Trump. Source

Political Slopaganda tells an inherent "truth" like satire.

Many examples of political slopaganda include some element of satire or parody; some explicitly adopt the form of political cartoons. Political cartoons often exaggerate certain characteristics to make their subject seem ridiculous in order to make a political point.

Figure 9 compares a famous 1871 Thomas Nast cartoon of Boss Tweed with an image from an AI video showing a grotesque version of Donald Trump kissing Speaker of the House Mike Johnson.

Figure 7: Thomas Nast 1871 cartoon of Boss Tweed and 2026 image of Trump from the TikTok and Instagram account "Late Stage Living". Source.

Continuing the long tradition of political satire or cartooning, political slopaganda can be used to express a "truth" beyond its obvious factual falsity. This has long been the justification for satire or political humor, which often presents obviously extreme caricatures or situations to make a deeper political point. Writing about the long history of political cartoons, Hess and Northrop observe: "A good cartoon is so simple, so straightforward, it seems to speak an inherent truth." But the "truth" is not necessarily an informational one, but rather something deeper and more about what something means. They note, "A cartoonist's job is not to present the news but to interpret it."

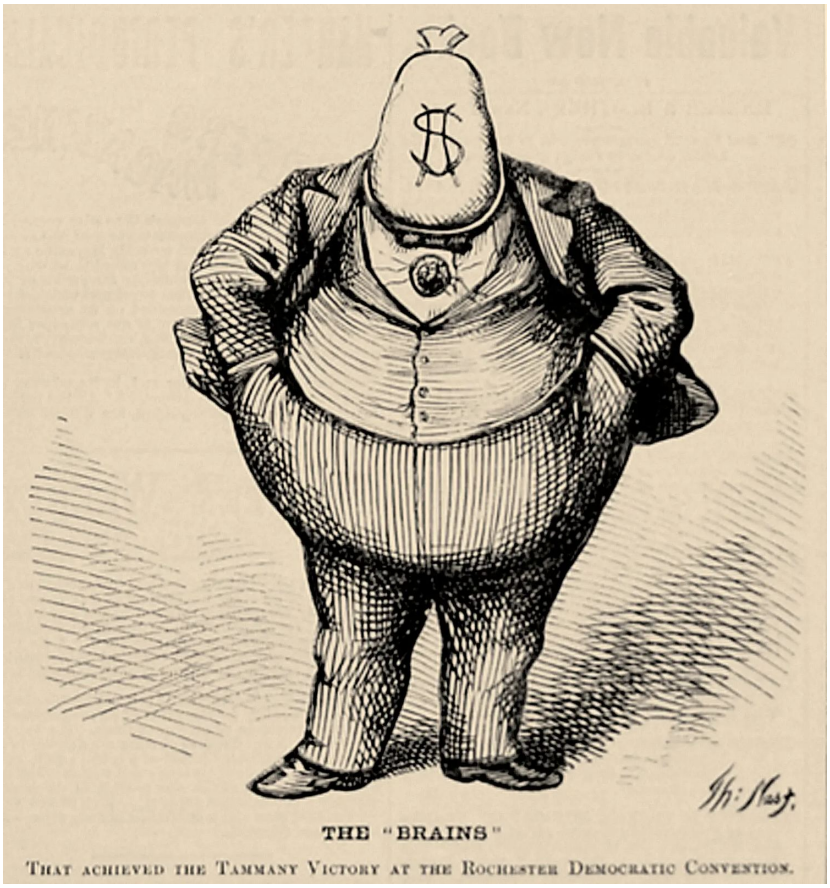

Figure 8. 1871 Thomas Nast cartoon depicting a member of Tammany Hall. Source.

Charles Curran, the filmmaker behind the Spencer Pratt Batman videos, expressed something similar in a podcast interview, speaking about his AI-generated ads: “They’re effective because there’s truth in them. And it doesn’t matter whether it’s AI that communicates the truth. They resonate with people because they’re speaking to something true.”

Similarly, New York gubernatoria candidate Bruce Blakeman’s ad described above begins with the following text:

Figure 9. Opening screen of Bruce Blakeman ad. Source.

It is notable, however, that this invocation of an inherent truth comes through content that is *obviously false*. One way of conceptualizing this contradiction and what political slopaganda is trying to accomplish is Arlie Hochschild's concept of "deep stories" from her book *Strangers in Their Own Land*. For Hochschild, "A deep story is a feels-as-if story -- it's the story feelings tell, in the language of symbols. It removes judgement. It removes fact. It tells us how things feel." Deep stories are the narratives that structure how people make sense of their world. In their book *You are Here*, Whitney Phillips and Ryan Milner extend this idea to "deep memetic frames" that "reflect fundamental ideals and can cohere an entire lifetime of choice, beliefs, and experiences," and ultimately, establish "ethical and ideological ways of being in the world."

For example, the Republican-aligned PAC Citizens for Sanity released an ad that showed Senate Candidate James Talarico singing a version of "My Favorite Things" where the lyrics are: "Boys in white dresses with blue satin sashes, Girls dosed with hormones till they grow moustaches, changing the gender of all your offspring, these are a few of my favorite things."

Figure 10. Citizens for Sanity released an AI-generated ad of James Talarico. Source.

This ad is not true. Talarico did not produce this video; he did not sing this song; he has not said he supports any of what the song lists. And at the risk of being too generous to what is above all an attack ad, the ad captures what some on the right believe many Democrats think but won't admit regarding LGBTQ+ rights. But more, this ad plays on common Republican claims about the masculinity of those on the left, and on what they see as a persistent cultural decline in traditional gender roles, and masculinity in particular. That is, the point isn't likely to make people believe that Talarico sang the song, but instead might be to make viewers feel a certain way about Democratic priorities and about how Democrats articulate masculinity.

It is no coincidence that right-wing political slopaganda routinely displays Republicans as muscular, whether that is Spencer Pratt, Mike Rogers, or Donald Trump.

Figure 11: Three AI-generated images of Republican politicians.

Deep stories or deep memetic frames have real force in public life. They give people cultural tools to make sense of the world and their place in it. These concepts, however, are all broader and bigger than a single piece of communication. But the point is not that a single ad could reshape a deep story, but rather that the ad is powerful *because* it resonates with an existing deep story. Deep stories are persistent and deep-rooted. While, of course, ads can amplify or reinforce deep stories, the far more immediate value is helping us understand that this content may draw power from its ability to align or resonate with these persistent narratives, tropes, or ideas, especially amongst politically aligned groups.

Political Slopaganda's Other Impacts

Regardless of political slopaganda's ability to successfully support candidates, gain attention, or tell the "truth," it may have other, potentially unintended, effects or consequences.

Political Slopaganda May Lower Epistemic Trust

Scholars and commentators from across fields have observed and lamented a broad "crisis in trust," in which the public reports worsening trust in most institutions. Rather than the result of a single force, it is likely a confluence of partisanship, institutional performance, economic or political inequality, and media trends.

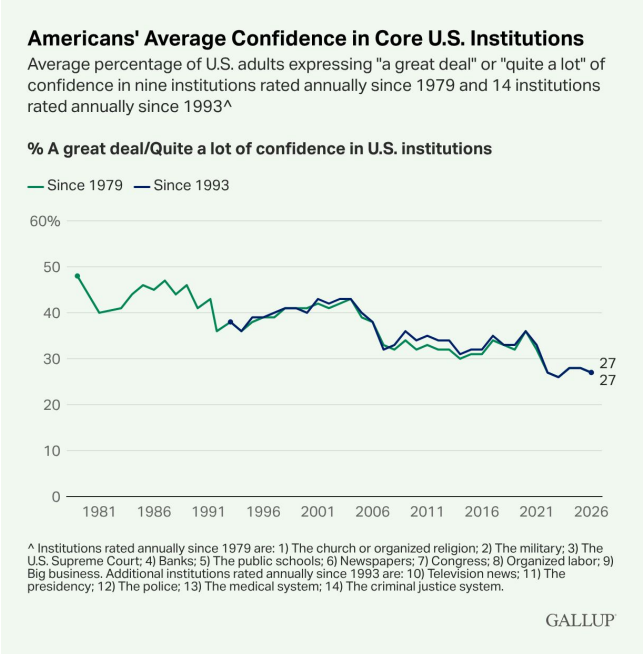

Figure 12: Gallup survey data on trust in U.S. institutions between 1981 and 2026. Source

While we do not have empirical evaluations of the impact of political slopaganda on public trust, scholars have examined whether satire, political advertising, and digital memes or other low-quality information might contribute to the decline of institutional trust.

Over the last decade, there has been a great deal of work exploring the impact that misinformation exposure has on trust. To simplify somewhat, there is little evidence that exposure to pieces of misinformation has strong negative impacts on institutional trust, even when that piece of misinformation has some (small) effect on factual belief. In contrast, there is stronger evidence that persistent or repeated narratives or conspiracy claims -- especially coming from political leaders -- can have a larger impact both on factual belief and on trust in institutions.

In particular, repeated falsehood can negatively impact what scholars call epistemic trust, or trust in knowledge-producing institutions, such as media or science. The effect, however, is generally small. One explanation for this observed effect is that even if viewers don't necessarily *believe* falsehoods, the uncertainty generated by falsehoods contributes to a generalized confusion about our ability to tell the truth from falsehoods. This uncertainty created by persistent lies has been a central strategy of fascist and authoritarian governments. As Hannah Arendt famously described, in a landscape defined by uncertainty and distrust, authoritarian leaders work to position themselves as the only trustworthy source of truth.

Supporting this, there is growing evidence supporting what some are calling the Liar's Dividend. The Liar's Dividend refers to how epistemic uncertainty, especially around AI-generated deepfakes, permits public figures, such as politicians, to claim that (real) evidence of their wrongdoing is fake.

In general, the relationship between misinformation and epistemic trust often works as part of a self-reinforcing cycle, where people with existing low levels of trust consume more falsehoods, and see a further decline in epistemic trust. Indeed, there is far more support for the proposition that having low trust predicts misinformation exposure than misinformation exposure predicts low trust.

Political Slopaganda may Worsen Affective Polarization

Similar to the broad decline in trust, many have observed a notable rise in affective polarization over the last several decades.

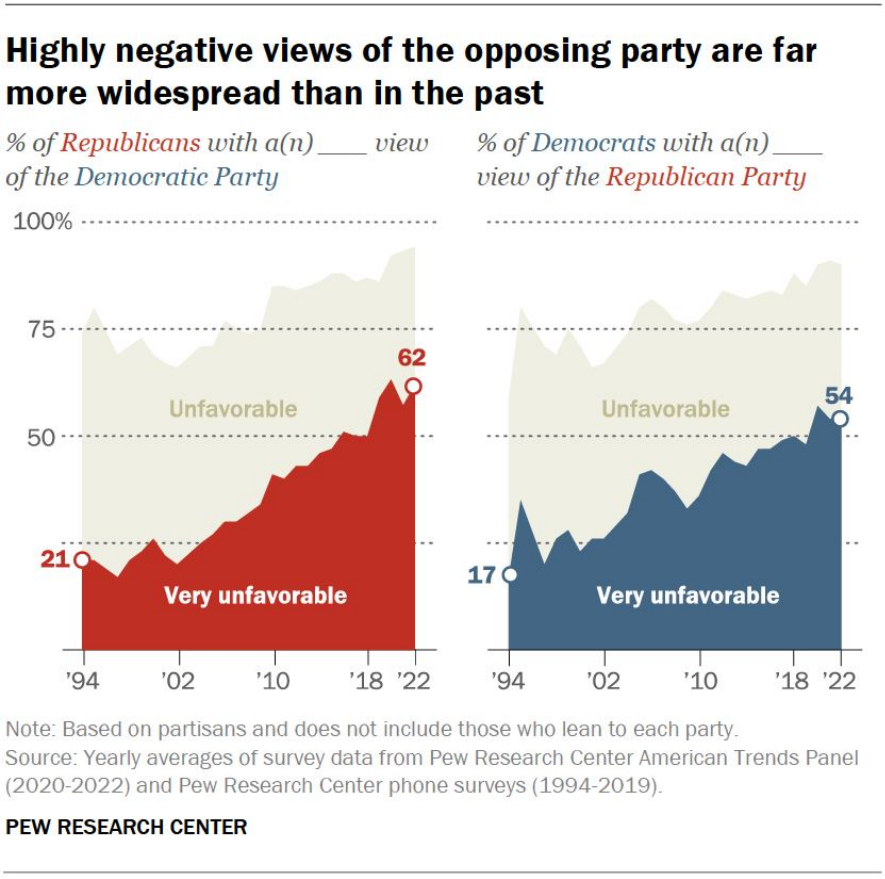

Figure 13: Pew Research on negative views of opposition political parties. Source.

Affective polarization is characterised by animosity between the political parties. It is described as the tendency to feel good about your own political party, while feeling dislike, hostility, or enmity for the opposing party.

Like the decline in trust, there is no one simple explanation for the rise in affective polarization, and it likely has grown out of party realignment, social sorting, gerrymandering, and the growth of outrage as a media style.

However, early research on affective polarization drew a link between exposure to negative political campaigning materials and affective polarization. Notably, there is evidence that both negative attacks from aligned politicians against opponents and unaligned politicians attacking your aligned politicians *both* increase polarization.

Given that much political slopaganda effectively functions as negative campaign advertising, there is reason to suspect it may also exacerbate affective polarization. However, that effect is likely very small. Even during the height of campaign season, negative ads constitute a tiny portion of the content each of us experiences.

Efforts to draw a causal link between viewing other low-quality content and affective polarization have mostly come up short. Despite popular belief, there is little empirical evidence that algorithmic curation of social media or exposure to misinformation has been major drivers of affective polarization. In contrast, there is better evidence that affective polarization can contribute to misinformation exposure and that misinformation consumption is concentrated among strong partisans.

How Can We Address Political Slopaganda?

Although we still lack a strong evidence base to understand the specific ways that political slopaganda may shape our elections, above we saw some indication that it may have limited negative impacts. Given this, what can and what should we do to address those impacts?

Currently, there are three main policy interventions meant to address *deceptive* AI-generated content. We begin by discussing each, showing that each approach is largely insufficient to address the potential harms this content poses to our elections.

Bans

Several states have passed laws that prohibit deceptive AI content in a narrow window before elections. For example, in 2024, the California legislature approved AB 2839, which prohibited "materially deceptive" AI-generated ads depicting someone else doing something they did not do in the 3 months before an election.

However, there are three things to note about these laws.

First, it is unclear if much of the political slopaganda described here would count as materially deceptive deepfakes, which the CA bill defines as "falsely appear[ing] to a reasonable person to be an authentic record of the actual speech or conduct of the individual depicted in the media."

Second, the California law, like others, has an explicit exemption for satire or parody, as long as it includes a label claiming it is satire or parody. While there will continue to be debates about the definition of satire and parody, much of the political slopaganda discussed above would likely pass the exemption.

Third, and most importantly, several of the laws outright banning this content have been overturned by the courts. The courts have not yet considered constitutional challenges to the laws that remain in effect in Maryland, Texas, and Minnesota. Simply put, there is strong reason to believe that political slopaganda is constitutionally protected speech. In the U.S., political speech receives the strongest First Amendment protection. This includes political parody and satire. This baseline dates back to *New York Times v. Sullivan* (1964), where the Supreme Court held that even false statements about public officials can be protected speech under the First Amendment.

Ruling on AB 2839, in *Kohls v. Bonta* (2024), Judge John Mendez described the statute as "a hammer instead of a scalpel" that "hinders humorous expression and unconstitutionally stifles the free and unfettered exchange of ideas." If the courts have not permitted restrictions on *deceptive* AI-generated content, there is little reason to believe that courts will permit restrictions on non-deceptive content.

Labels

More than 30 states have enacted requirements that political advertisers include labels on deceptive political ads containing AI. Many of these laws only apply to materially deceptive content.

However, again, as with the laws above, it is unclear whether political slopaganda would count as "materially deceptive." More directly, many of these laws have explicit exemptions for parody and satire.

But perhaps most importantly, it is unclear if labeling would address the potential negative impacts on trust and affective polarization that might result from political slopaganda. That is, if most viewers already can tell that a piece of content is AI-generated, what value will come from a label announcing it? Labeling is ultimately designed to target deception rather than spectacle. If the harms of this content extend beyond deception, labels are unlikely to have a positive impact.

Likeness protections

Several U.S. jurisdictions have passed or considered likeness protections. The laws enacted at the state level have all only addressed *commercial* likeness rights. For example, Tennessee's Elvis Act grants recording artists a right over their vocal likeness. The No Fakes Act, federal Legislation that remains under consideration, would grant everyone digital likeness rights. However, all of these laws have explicit exemptions for parody or satire. Again, while courts may still adjudicate if a piece of political slopaganda does in fact qualify as parody or satire, these laws likely leave candidates depicted in political slopaganda little recourse.

What Should We Do?

Bans, labeling requirements, and likeness protections are all unlikely to address the potential harms discussed above. So what should we do?

Whatever we do must contend with two realities.

First, political slopaganda is and should be constitutionally protected speech. The courts have long held that lying is protected. There is no reason that AI-generated lies should be treated differently.

Second, political slopaganda is ultimately a problem emerging from the political and informational systems themselves. Like misinformation, it is better seen as a symptom rather than the problem.

Together, these realities mean that interventions addressing the content of political slopaganda are unlikely to be constitutional; moreover, they are unlikely to be effective. Systemic problems require systemic solutions.

In particular, we see two areas where practical interventions could address some of the structural forces behind political slopaganda's potential role in worsening epistemic trust and affective polarization: government transparency and information system resilience.

Campaign finance transparency. Congress should close loopholes to increase transparency into our campaign finance system.

1 The FEC should revise subvendor reporting requirements. Currently, campaigns must disclose disbursements to vendors. However, they are under no obligation to disclose how those vendors disburse money on their behalf. This means that while most campaign advertising is handled by campaign consultants, FEC records only include payments from campaigns to consultants, not data about how those consultants then disburse money on behalf of campaigns.

As we have argued for years, the lack of subvendor reporting requirements is obscuring much about the vast majority of campaign advertising spending. These rules hinge on a 1983 FEC ruling, and could be amended through a new ruling.

2 Congress should alter donor disclosure requirements for political spending by 501(c)4 organizations. While requirements that 501(c)4 organizations reveal all donors may be overbroad (see for example *NAACP v. Alabama*, Buckley v. Valeo), there should be rules that require disclosure for money explicitly earmarked for political spending. There is strong Supreme Court precedent that the government has a compelling interest in helping voters understand the sources of election spending. Even while political slopaganda may remain constitutionally protected, it is in the public interest that voters can understand who exactly is funding it.

Protect and build the health of the information system. Our media and information systems face threats from all fronts. There is much we could do to increase both individual and system resilience.

1 Prohibit misinformation about the time, manner, and place of elections. In general, misinformation is and should be protected speech. However, as we have written before, we believe that intentional lies about the time, place, and manner of elections should not be permitted. While some states have narrow provisions that prohibit some of this content, there should be federal prohibitions on this content. There is no compelling reason that we should protect lies specifically meant to disenfranchise voters; in contrast, prohibiting such content would protect voting rights.

2 Both the federal and state governments should fund media and AI literacy programs for both students and adults. AI and media literacy can play an important role in building resilience to low-quality content. Most AI and media-literacy have focused on training young people in schools. This is important, and should be continued and expanded. However, most misinformation is shared by *older* viewers. It is essential that we are working to make sure all people have the tools and understanding to engage critically with media content.

3 Support local journalism. Local journalism is key to protecting the informational health of communities. Unfortunately, as local journalism collapses in the U.S., there are no easy silver bullet solutions to address it. However, we need to invest in a range of options, from supporting non-profit outlets to considering tax subsidies for newsrooms to cover hiring, to helping broker relationships between news outlets and digital media companies, to other structural changes, in order to help support a vibrant local news ecosystem.