WWW.IJRPR.COM

International Journal of Research Publication and Reviews

Journal Homepage: www.ijrpr.com ISSN 2582-7421

Transparency Consent and Responsibility in AI Generated Political Speech

MU YIZE, Dr. Mohd Haizra Hashim

Faculty of Business and Communication, INTI International University, Nilai, Malaysia
Email: i25030568@student.newinti.edu.my; mhaizra.hashim@newinti.edu.my

ABSTRACT

AI-generated political speech raises a difficult question for media ethics: when does making a politician appear to speak support public participation, and when does it undermine the public's ability to judge? This paper examines that boundary through a normative comparison of three documented cases: the Biden voice-cloning robocalls before the 2024 New Hampshire primary, Imran Khan's AI-generated address from prison in December 2023, and the synthetic Suharto campaign video circulated before Indonesia's 2024 election. Deontological ethics and social responsibility theory guide the analysis, supported by the TARES principles of ethical persuasion. The cases are assessed through four questions concerning truthful representation, authorisation, meaningful disclosure, and responsibility for foreseeable harm. The comparison supports a qualified argument: disclosure is necessary when realistic synthetic speech could be mistaken for an authentic recording, but disclosure alone cannot justify false endorsements or misleading voting information. Conversely, a blanket prohibition can restrict legitimate expression when an identifiable speaker approves a message and its synthetic delivery is made clear. The paper proposes proportionate safeguards that connect disclosure to verifiable authorship, accountable publication, timely correction, and protection for satire. Its contribution is a practical ethical distinction between using AI to communicate someone's position and borrowing someone's identity to manufacture political authority.

Keywords: Media ethics; AI-generated political content; synthetic speech; political communication; informed judgement; accountability

1 Introduction

A familiar voice can make a political message feel trustworthy before a listener has considered its argument. When that voice is generated by AI, the listener must judge two things at once: whether the message is credible and whether the person apparently speaking actually stands behind it. This creates an ethical problem that cannot be settled simply by asking whether AI was used.

The contrast between recent cases makes the problem clear. A cloned voice resembling Joe Biden's was used to discourage participation in the 2024 New Hampshire primary. In Pakistan, Imran Khan's party used a synthetic voice to deliver a speech that party officials said he had approved from prison. In Indonesia, a simulated version of the late President Suharto urged support for Golkar candidates (Federal Communications Commission [FCC], 2024; Shahzad, 2023; Firdaus, 2024). Each involved manufactured speech associated with a real politician, yet the relationships between speaker, producer, and audience differed substantially.

Media ethics provides a way to examine those relationships. Political persuasion can appeal to emotion, simplify a policy, and invite identification with a leader. However, audiences should still be able to understand whose position is being presented and assess the reasons offered. Synthetic speech becomes ethically troubling when it supplies a false answer to that basic question of authorship. It can also harm political discussion without producing complete belief in a falsehood: uncertainty about authentic recordings may itself weaken trust (Vaccari & Chadwick, 2020).

This paper asks: under what conditions is realistic AI-generated political speech ethically defensible? Two further questions guide the discussion. What can disclosure accomplish, and what remains unresolved after a label is added? How should responsibilities be shared among political communicators, production services, distributors, and platforms? The argument is that ethical acceptability depends on truthful representation, appropriate authorisation, meaningful disclosure, and accountable handling of foreseeable harm. A label informs an audience about production; it does not by itself make a message honest or its publisher responsible.

2 Literature review and conceptual boundaries

2.1 Synthetic speech and political deception

Synthetic political speech refers here to AI-generated audio or audiovisual material that presents an identifiable political figure as speaking. The focus is realistic impersonation or reconstruction, rather than spelling correction, ordinary editing, or an obviously fictional cartoon. This scope matters because different uses make different claims about reality. An AI translation of an approved speech can preserve the speaker's position. A fabricated endorsement can assign the speaker a position they never held.

Wardle and Derakhshan (2017) distinguish misinformation from disinformation partly through the intention behind its production and circulation. This distinction helps avoid treating every person who forwards a false clip as its deliberate author. It also shows why technical origin is insufficient for ethical classification. Synthetic material may be honestly presented, while a genuine recording can mislead when edited or placed in a false context. Paris and Donovan (2019) similarly situate deepfakes within a longer history of audiovisual manipulation. The relevant ethical questions concern how evidence is represented, who controls it, and whose interests are served.

Research also cautions against assuming that all realistic fakes successfully persuade audiences. Vaccari and Chadwick's (2020) experiment found that political deepfakes were more likely to generate uncertainty than straightforward deception, with that uncertainty reducing trust in news on social media. Pawelec (2022) broadens the democratic concern to participation, public deliberation, and collective decision-making. These accounts suggest that the ethical stakes extend beyond whether a particular viewer believes a particular sentence. They include the conditions under which citizens can question public figures and discuss shared evidence.

2.2 What disclosure can and cannot achieve

Disclosure has a clear ethical purpose: it gives audiences information relevant to their judgement. Its practical effects, however, depend on what the label says and what outcome is measured. In two preregistered experiments involving misleading AI-generated images, Wittenberg et al. (2025) found that labels reduced belief in the claims presented. Labels identifying AI production alone tended to have little effect on reported engagement intentions. The distinction between belief and sharing is important: understanding that an image is synthetic does not necessarily stop someone from circulating it.

Gallegos et al. (2026) examined a different setting: AI-generated arguments about public policies. Authorship labels did not significantly change persuasion, perceived accuracy, or sharing intentions. Their study concerned text about selected policy issues, rather than voice impersonation or false voting instructions. It therefore cannot establish that labels are ineffective in every political setting. Taken together, these studies support a narrower conclusion: disclosure can improve information about production without reliably removing a message's influence.

This paper builds on that distinction rather than claiming that disclosure has been neglected. It brings audience information into the same ethical assessment as authorisation and responsibility. The three cases permit comparison between deceptive impersonation, reported authorised representation, and an endorsement attributed to a deceased figure. Their value lies in exposing different ethical boundaries, not in estimating how frequently such practices occur.

3 Method and analytical approach

The study uses normative comparative case analysis based on publicly available documents. The cases were selected purposively because they involve recognisable political identities but differ in the source of authority claimed for the speech. The New Hampshire case is documented through the FCC's 2024 enforcement record and reporting on the subsequent criminal verdict. The Khan case is examined through Reuters reporting, including statements attributed to party officials. The Suharto case draws on BenarNews reporting and Reuters' account of AI use in Indonesia's campaign (FCC, 2024; Ramer, 2025; Shahzad, 2023; Firdaus, 2024; Lamb et al., 2024).

For each case, the analysis first identifies the documented production and publication context. It then separates what sources establish from what remains uncertain, applies the ethical criteria below, considers a plausible defence, and reaches a qualified judgement. Reported party approval is treated as a claim about authorisation, rather than independent proof of every production decision. Public controversy is treated as evidence of an ethical dispute, rather than proof of electoral harm.

The study does not conduct interviews, inspect original production files, or measure changes in voting behaviour. It is not a systematic review or a representative survey of political deepfakes. Its findings are reasoned ethical assessments of documented practices. They should not be read as causal conclusions about election results or as a statement of legal rules across different countries.

4 Ethical framework

4.1 Deontological ethics and truthful representation

Deontological ethics directs attention to duties that cannot be dismissed because a campaign expects a useful result. Kant's account of persons as ends supports a duty to respect an audience's capacity for independent judgement (Kant, 1785/2012). Applied here, a communicator should not secure agreement by hiding a material falsehood about who is speaking. This application concerns the audience's agency as well as the apparent speaker's identity.

The paper distinguishes the truth of a political claim from the truth of its representation. A policy argument might be factually defensible while the claim that a particular leader endorsed it is false. Equally, permission to clone a voice would not excuse false information about voting. These are separate duties. The proposed ethical test therefore asks both whether the message misrepresents relevant facts and whether it misrepresents a person's participation or approval.

4.2 Social responsibility and the power to distribute

Social responsibility theory connects media freedom with obligations to public understanding and access to meaningful discussion (Siebert et al., 1956). Its extension to digital platforms requires care. Platforms, telecommunications providers, and campaign teams have different functions, and a theory developed around the press does not automatically assign them identical editorial duties. Nerone (2018) also highlights the historical and cultural limitations of the original press typology.

This paper uses the theory as a normative starting point, rather than a universal description of media systems. Where an organisation controls the production, paid distribution, or visibility of realistic political impersonation, it has a corresponding responsibility to address predictable deception within that control. This does not require guaranteeing that every political opinion is correct. It requires reasonable precautions, a route for challenges, and a response when serious misrepresentation is established.

4.3 Four questions for ethical assessment

Baker and Martinson's (2001) TARES test identifies truthfulness, authenticity, respect, equity, and social responsibility as principles of ethical persuasion. These principles help translate the two theoretical perspectives into four practical questions. First, does the content truthfully represent the message and the apparent speaker's involvement? Second, what supports the producer's authority to speak in that person's identity? Third, can an ordinary recipient understand the synthetic nature, publisher, and approval status of the message? Fourth, who can prevent, explain, and remedy foreseeable harm?

These questions form an interpretive guide rather than a validated measurement scale. They also require context. Authorisation is especially important when a publisher claims to deliver a politician's own position; it should not become a general permission requirement for political criticism or recognisable satire. Respect and equity require considering both the audience's ability to assess a message and the risk that restrictions silence less powerful speakers.

5 Comparative ethical analysis

5.1 Biden voice cloning and misleading voting information

Before the January 2024 New Hampshire primary, voters received robocalls using an AI-generated imitation of Biden's voice. The FCC documented a campaign involving misleading voting messages and caller ID spoofing, and adopted a six-million-dollar forfeiture against consultant Steve Kramer in September 2024 (FCC, 2024). A separate criminal proceeding ended in June 2025 with Kramer's acquittal on voter-suppression and candidate-impersonation charges (Ramer, 2025). The enforcement action and criminal verdict concern different proceedings; neither should be substituted for a careful ethical judgement.

The ethical problem begins with the relationship created between the listener and the apparent speaker. A recognisable presidential voice lent authority to advice about electoral participation. The message did not merely invite support for a policy or criticise an opponent. It encouraged a decision about voting through a representation that recipients had reason to misunderstand. Under deontological ethics, this interferes with informed judgement even without evidence that a particular recipient changed their behaviour.

Ramer (2025) reports Kramer's defence that he intended to draw attention to AI's dangers. Taken seriously, this is a public-interest argument: a disturbing demonstration might prompt stronger protections. Yet it asks the audience to bear an undisclosed risk for the producer's educational objective. A demonstration directed at consenting participants, or clearly introduced as synthetic, could expose the same vulnerability without using voters' immediate electoral decisions as its testing ground. The claimed purpose does not resolve the means used.

Social responsibility also extends the assessment beyond the person commissioning the recording. A provider's obligations should depend on its knowledge, capacity, and role, rather than the mere fact that its service was used. Nevertheless, organisations that distribute high-risk impersonation should have proportionate checks and a workable escalation route. The recommendation is precaution where foreseeable harm and practical control meet, rather than automatic blame for every intermediary.

Adding an AI label would address only part of this case. Clear disclosure could reduce mistaken attribution, but inaccurate voting advice would still require correction. The ethical judgement is therefore adverse on two grounds: the misuse of another person's apparent authority and the misleading treatment of electoral participation. This analysis does not claim that the calls altered the primary result, nor does it treat criminal acquittal as an ethical endorsement.

5.2 Imran Khan and reported authorised representation

In December 2023, Khan's party used a synthetic version of his voice in an online rally while he was imprisoned. Reuters reported party officials' statement that the speech came from a written version he had approved, alongside restrictions on physical campaigning and his media access (Shahzad, 2023). This account supports considering an alternative use of synthetic speech: enabling a person to communicate a position when direct delivery is constrained. The public reporting does not independently verify the full approval process.

Assuming the reported authorisation was genuine, the central ethical relationship differs from impersonation intended to conceal the actual source. A political figure may approve a speech drafted by someone else and delivered through mediated channels. AI does not automatically destroy that relationship. On a deontological reading, the crucial duties are to represent the approved position faithfully and avoid giving the audience a false impression of the event. Agreement to a script should not be stretched into permission for unlimited future synthetic statements.

The strongest objection concerns the emotional force of a familiar voice. Supporters may experience the recording as personal presence, even when they know it is generated. Approval of words also does not necessarily establish approval of tone, emphasis, or the surrounding images. These concerns call for specific disclosure and review of the finished recording. They do not establish that an authorised synthetic address is equivalent to deceptive impersonation.

Social responsibility introduces a competing concern: access to political expression. An absolute ban could remove a channel of participation for a speaker who cannot record a conventional address. That possibility does not validate every claim made by the party or settle the surrounding political dispute. It does show that preventing manipulation and protecting participation must be considered together. A defensible approach would disclose that the voice is synthetic, identify the publisher, explain the claimed approval, and provide the approved text for comparison.

The judgement is conditionally permissive. The speech is more ethically defensible if the speaker approved both the meaning and its presentation, the audience can recognise the synthetic delivery, and an accountable publisher handles disputes. Where approval cannot be established, the material should not be presented as a verified personal address. Unlike the New Hampshire case, the medium itself does not supply sufficient grounds for condemnation.

5.3 Suharto and the manufacture of posthumous endorsement

Before Indonesia's 2024 election, a synthetic video depicted Suharto, who died in 2008, encouraging support for Golkar candidates. BenarNews documented its circulation and the resulting controversy. Reuters subsequently reported that the campaign clips were marked as AI-generated (Firdaus, 2024; Lamb et al., 2024). This case is particularly useful because disclosure was reported, yet a substantial ethical question remained: what authority did the synthetic figure have to endorse candidates in a later election?

The problem is not resolved by observing that informed viewers know Suharto is dead. A viewer can understand that an image is synthetic while still receiving its endorsement as an emotionally powerful expression of historical authority. The paper's interpretation is that the video borrows political standing from a person who cannot approve the new statement. A production label answers how the appearance was made; it does not establish that the endorsement reflects an authorised position.

There is a reasonable defence. Political parties frequently use historical symbols, inherited identities, and references to former leaders. A clearly disclosed reconstruction might be understood as an imaginative statement about a party's tradition. This defence is strongest where the material unmistakably presents the living campaign's interpretation. It is weaker when a realistic first-person address asks the audience to treat a newly composed endorsement as the former leader's own voice.

Under deontological ethics, historical affiliation cannot simply be converted into approval of present candidates. This paper does not establish that there was no prior permission or that no family member supported the production. Rather, the reviewed reporting does not establish a prior mandate for the specific endorsement. Even family approval would answer a different question from whether the deceased person held that current political position. The ethical issue concerns the attribution being made to the audience.

Social responsibility adds the need to protect historical discussion without allowing campaign publicity to masquerade as historical testimony. A more defensible presentation would identify the video as a fictional reconstruction produced by the current campaign and explicitly distinguish it from an authentic statement. Documentary reconstruction and obvious satire deserve different treatment from a realistic electoral endorsement. The adverse judgement here concerns the unsupported attribution of present political authority, rather than every creative use of a deceased public figure.

Table 1 Comparative ethical assessment of synthetic political speech

Case	Authority claimed	Main ethical concern	Qualified judgement
Biden robocalls	Apparent personal advice from a living politician	Misattribution and misleading voting information	Ethically unacceptable; a production label would not cure false voting advice
Khan virtual address	Party reported approval of a written speech	Verification of approval and faithful synthetic delivery	Conditionally defensible with meaningful disclosure and accountable publication
Suharto campaign video	A deceased leader apparently endorses current candidates	New endorsement borrows historical authority without an established mandate	Ethically problematic as personal endorsement; explicit fictional framing changes the assessment

The table summarises this paper's interpretations of the sources discussed above. It is not a ranking of national political systems or a measurement of effects on voters.

6 Discussion

6.1 Transparency and consent address different wrongs

The comparison shows why transparency should be treated as a duty within a wider ethical judgement. A useful disclosure allows recipients to distinguish an authentic recording from synthetic delivery. It should also identify the organisation responsible for the message. However, an audience can be fully informed that a voice is generated while the underlying endorsement remains invented. Transparency about production and honesty about representation are related but distinct.

Consent also has limits. A politician's permission can support the legitimate delivery of their own position, but cannot excuse a false claim about electoral procedures or a harmful fabrication about someone else. Conversely, requiring a politician's consent for every depiction would give public figures an unacceptable power over criticism. The relevant distinction is between claiming to communicate on someone's behalf and commenting on that person. Ethical policy must preserve that difference.

The TARES principles help clarify the boundary. Truthfulness concerns both the assertion and its attribution. Authenticity asks whether the persuasive relationship conceals its source. Respect protects the audience's capacity to judge, while equity asks whether people with fewer resources can access and challenge political messages. Social responsibility directs attention to the wider conditions of public discussion (Baker & Martinson, 2001). No single disclosure substitutes for those obligations.

6.2 Shared responsibility needs identifiable decisions

Calling responsibility shared can make it so vague that no organisation has to act. A more useful allocation follows decisions. The commissioning campaign should establish approval and check factual claims. The producer should preserve the agreed scope of identity use. A distributor should be able to identify the customer responsible for high-risk publication. Platforms should provide clear rules and a prompt response to credible impersonation complaints. These are recommendations drawn from the framework, rather than findings that every actor in the three cases failed each duty.

Technical provenance can support this process, but should not be confused with truth verification. Chandra et al. (2024) survey approaches including provenance records, watermarking, and detection. Such tools can help establish an asset's production history; they cannot decide whether a policy claim is sound or an endorsement is morally justified. Floridi and Cowls (2019) likewise connect beneficial AI use with autonomy, justice, and explicability. Responsible publication therefore requires an answerable human decision as well as information about the technology.

6.3 Protecting trust without encouraging indiscriminate suspicion

Warnings about synthetic media can themselves create a difficulty. If citizens assume that any inconvenient recording might be fake, public figures gain another way to challenge authentic evidence. Chesney and Citron (2019) describe this risk as the liar's dividend. However, the effect should not be exaggerated. Schiff et al. (2025) found that false claims of misinformation could protect politician support in their experimental settings, but such claims were largely ineffective against video evidence. Their findings caution against assuming that public awareness inevitably makes all evidence useless.

The ethical goal should be better grounds for confidence. Journalists should explain how an important recording was verified, and platforms should distinguish confirmed manipulation from unresolved authenticity. Political communicators should preserve approved texts and source records so that disputes can be checked. These practices respect audiences more than asking them either to trust every convincing voice or to distrust all recordings. Their value lies in making judgement possible, rather than promising certainty in every case.

7 Recommendations for responsible political communication

Campaigns should adopt an approval process proportionate to the proposed use. Realistic synthetic statements presented as a living politician's own position should require documented approval of the script and final presentation. Records should name the organisation authorising publication and specify the intended context. A previous voice-cloning agreement should not be assumed to authorise new endorsements indefinitely. For deceased figures, publishers should distinguish documented historical statements from newly imagined speech.

Disclosure should explain the alteration that matters to the audience. For a voice clone, a useful statement would identify the synthetic voice, the publisher, and whether the message was approved by the apparent speaker. Audio requires audible disclosure; video requires legible disclosure during the relevant material. The exact design should be tested with the intended audience, including local languages and accessibility needs. These proposals establish a practical standard to evaluate, rather than assuming that any particular label has already proved effective.

Platform responses should distinguish production transparency from substantive deception. An accurately presented synthetic translation may need explanation without removal. A fabricated endorsement needs correction of attribution. A false instruction about voting may justify urgent restriction under clear, consistently applied rules, regardless of whether AI produced it. Recognisable satire and political criticism need protection, with reasons and an appeal route when moderation is contested. This approach limits the opportunity to use anti-deception measures against legitimate opposition.

Correction should reach people through the channels in which the misleading claim travelled, as far as reasonably possible. Deleting an original upload alone may leave redistributed copies or mistaken impressions untouched. Platforms and publishers should identify what was wrong, provide an authoritative correction, and make the publisher responsible for responding. For disputed voting instructions, correction should direct audiences to the relevant election authority rather than a competing campaign's account.

Finally, evaluation should examine whether people understand the message's origin, speaker approval, and factual status. Counting labels or removals provides little information about that understanding. Audience testing, independent scrutiny, and reporting on correction times could help assess whether safeguards work. Media literacy can support these measures, but ordinary citizens should not carry the whole burden of investigating political impersonation. Organisations with greater control over production and distribution have stronger duties to make responsible judgement feasible.

8 Limitations and future research

The comparison relies on public reporting and official documentation, which provide uneven access to production decisions. Authorisation in the Khan case is reported rather than independently audited, and the reviewed material does not establish a prior mandate for Suharto's specific synthetic endorsement. The paper assesses those evidential limits openly. It also cannot show how representative the cases are or establish their effects on turnout, trust, or electoral results.

Theories developed in Western philosophical and press traditions may overlook important local understandings of authority, representation, and collective responsibility. Applying the same questions across cases improves analytical consistency, but does not remove those contextual differences. Future studies could examine how audiences in different languages interpret approval disclosures, compare responses to authorised and unauthorised clones, and investigate whether correction reaches people who encountered the original material. Such work would test the practical value of the safeguards proposed here.

9 Conclusion

AI-generated political speech should be judged through the relationship it creates among the apparent speaker, publisher, and audience. The Biden robocalls illustrate the ethical harm of lending a false identity to misleading voting advice. Khan's reported authorised address shows why synthetic delivery can sometimes support political participation. The Suharto video demonstrates that an AI label can coexist with an unresolved claim to political authority.

Deontological ethics protects audiences from being manipulated through false attribution, while social responsibility directs attention to the organisations able to prevent and remedy foreseeable harm. Together, they support a proportionate standard: realistic synthetic political speech needs truthful representation, appropriate authority, meaningful disclosure, and identifiable responsibility. Permission does not excuse falsehood, and disclosure does not create an endorsement. Ethical communication gives citizens a fair opportunity to understand who is addressing them, on what basis, and who will answer when that representation is challenged.

References

Baker, S., & Martinson, D. L. (2001). The TARES test: Five principles for ethical persuasion. *Journal of Mass Media Ethics, 16*(2-3), 148-175. https://doi.org/10.1080/08900523.2001.9679610

Chandra, B., Dunietz, J., Roberts, K., Lee, Y., Fontana, P., & Awad, G. (2024). *Reducing risks posed by synthetic content: An overview of technical approaches to digital content transparency* (NIST AI 100-4). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.AI.100-4

Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review, 107*(6), 1753-1820. https://lawcat.berkeley.edu/record/1136469

Federal Communications Commission. (2024, September 26). *In the matter of Steve Kramer* (FCC 24-104, Forfeiture Order). https://docs.fcc.gov/public/attachments/FCC-24-104A1.pdf

Firdaus, A. (2024, January 12). Fake Suharto video fuels debate on AI use in Indonesian election campaign. *BenarNews*. https://www.benarnews.org/english/news/indonesian/suharto-deepfake-used-in-election-campaign-01122024135217.html

Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review, 1*(1). https://doi.org/10.1162/99608f92.8cd550d1

Gallegos, I. O., Shani, C., Shi, W., Bianchi, F., Gainsburg, I., Jurafsky, D., & Willer, R. (2026). Labeling messages as AI-generated does not reduce their persuasive effects. *PNAS Nexus, 5*(2), pgag008. https://doi.org/10.1093/pnasnexus/pgag008

Kant, I. (2012). *Groundwork of the metaphysics of morals* (M. Gregor & J. Timmermann, Trans.; Rev. ed.). Cambridge University Press. (Original work published 1785). https://doi.org/10.1017/CBO9780511919978

Lamb, K., Potkin, F., & Teresia, A. (2024, February 8). Generative AI faces major test as Indonesia holds largest election since boom. *Reuters*. https://uk.investing.com/news/stock-market-news/generative-ai-faces-major-test-as-indonesia-holds-largest-election-since-boom-3328611

Nerone, J. (2018). Four theories of the press. In *Oxford Research Encyclopedia of Communication*. Oxford University Press. https://doi.org/10.1093/acrefore/9780190228613.013.815

Paris, B., & Donovan, J. (2019). *Deepfakes and cheap fakes: The manipulation of audio and visual evidence*. Data & Society Research Institute. https://datasociety.net/research-library/deepfakes-and-cheap-fakes-the-manipulation-of-audio-and-visual-evidence/

Pawelec, M. (2022). Deepfakes and democracy (theory): How synthetic audio-visual media for disinformation and hate speech threaten core democratic functions. *Digital Society, 1*, Article 19. https://doi.org/10.1007/s44206-022-00010-6

Ramer, H. (2025, June 13). NH jury acquits consultant behind AI robocalls mimicking Biden. *Associated Press*. https://www.boston.com/news/local-news/2025/06/13/nh-jury-acquits-consultant-behind-ai-robocalls-mimicking-biden/

Schiff, K. J., Schiff, D. S., & Bueno, N. S. (2025). The liar's dividend: Can politicians claim misinformation to evade accountability? *American Political Science Review, 119*(1), 71-90. https://doi.org/10.1017/S0003055423001454

Shahzad, A. (2023, December 18). Pakistan's jailed Imran Khan uses AI-crafted speech to lure votes. *Reuters*. https://www.investing.com/news/world-news/pakistans-jailed-imran-khan-uses-aicrafted-speech-to-call-for-votes-3258449

Siebert, F. S., Peterson, T., & Schramm, W. (1956). *Four theories of the press: The authoritarian, libertarian, social responsibility, and Soviet communist concepts of what the press should be and do*. University of Illinois Press. https://www.bibliovault.org/BV.book.epl?ISBN=9780252051050

Vaccari, C., & Chadwick, A. (2020). Deepfakes and disinformation: Exploring the impact of synthetic political video on deception, uncertainty, and trust in news. *Social Media + Society, 6*(1), 1-13. https://doi.org/10.1177/2056305120903408

Wardle, C., & Derakhshan, H. (2017). *Information disorder: Toward an interdisciplinary framework for research and policy making* (DGI(2017)09). Council of Europe. https://rm.coe.int/information-disorder-report-november-2017/1680764666

Wittenberg, C., Epstein, Z., Péloquin-Skulski, G., Berinsky, A. J., & Rand, D. G. (2025). Labeling AI-generated media online. *PNAS Nexus, 4*(6), pgaf170. https://doi.org/10.1093/pnasnexus/pgaf170