

Should Candidates Campaign Like Influencers? Evidence From 16,572 Instagram Reels and Two Generative AI Video Experiments

Cecilia Y. Sui*

September 1, 2026

Abstract

Political candidates are increasingly urged to communicate like social media influencers on the theory that platform-native intimacy builds parasocial bonds and support. However, evidence that candidates follow the advice or benefit from it remains thin. Voters judge candidates as prospective officeholders and informality that signals authenticity in influencers may signal incompetence in candidates. I test these claims using 16,572 Instagram Reels from 200 accounts of state legislative candidates and two pre-registered experiments using AI-generated audiovisual stimuli that vary visual style and relational script while holding a candidate's face, voice, and message fixed. Candidates adopt the style selectively, with widespread use of engagement-predicting informal visuals and rare use of relational cues. The experiments show that neither dimension increases perceived closeness while both impose competence costs. Platforms thus reward what voters penalize. More broadly, the study contributes to an enduring debate of political communication and self-presentation, the need to appear fit to govern while remaining relatable to the governed, and introduces a reusable design for audiovisual experiments.

Keywords: digital campaigning, parasocial relationship, influencer-style communication, generative AI, audiovisual data

Word Count: 9,562

*PhD Candidate in Political Science, Washington University in St. Louis. Mailing Address: One Brookings Drive, St. Louis, MO 63130, United States. Email: c.sui@wustl.edu

1 Introduction

“I know it’s important. I use it. I need to use it. I have to be on it ... but being popular on social media does not convert to votes as much as some people like to think it does ... Likes aren’t votes.”

— *Candidate for the South Carolina House of Representatives*¹

To win office, politicians must appear qualified to govern while seeming enough like the people they represent to earn their trust. Fenno (1978) called the management of this tension “home style.” Although the tension itself is as old as electoral politics, new communication technologies periodically change how politicians navigate it. Radio brought politicians’ voices into the living room and rewarded those who could speak conversationally (Jamieson, 1990; Loviglio, 2005). Television made how candidates looked and behaved electorally consequential (Druckman, 2003; Mutz, 2015), and entertainment media later provided new ways to reach voters outside traditional political news (Baum, 2005; Kim and Patterson, 2025). Social media further reduced the role of intermediaries, allowing politicians to communicate with voters directly and continuously (Barberá et al., 2019; Castanho Silva and Proksch, 2022; Enli, 2017). Each time a new medium has changed how voters encounter politicians, politicians have had to adjust to the medium’s native style or risk losing the audience. However, adaptation also carries risks. The informality that humanizes some politicians may make others look unserious (Young, 2020), and candidates rarely know in advance which side of that line they will fall on. How politicians navigate this tension, and at what cost, is a recurring question in the study of political communication and representation.

Social media, especially short-form video, has produced a sharp version of this dilemma: the growing pressure to communicate like social media influencers.² Strategists in both

¹The quote is from an interview conducted by the author on July 7, 2026, under IRB protocol #202606070.

²Social media scholars distinguish among influencers, content creators, micro-celebrities, and related

parties now urge candidates to film themselves talking directly to the camera and to post like successful influencers. One strategist told reporters that “folks don’t want blow-dried politicians anymore” (Parnes, 2026). A handful of high-profile campaigns have made the influencer approach viable and electorally powerful. Alexandria Ocasio-Cortez drew widespread praise for livestreaming herself cooking dinner on Instagram while discussing pandemic relief (Smith, 2020). Zohran Mamdani’s rapid, intimate short-form videos became a celebrated feature of his successful mayoral campaign (Gee, 2025). Jon Ossoff’s popular TikTok presence during his Georgia Senate runoff reinforced that an internet-native style could attract attention and mobilize young voters (Perrett, 2020).

Text-based social media had already transformed political communication before short-form video arrived. Instead of relying on occasional appearances in programs controlled by others, candidates could produce and distribute their own content continuously, in spaces where informality and intimacy increasingly became the norm. Every post also generated immediate, quantified feedback on attention and engagement, even as its effects on voter judgments remained invisible (Enli, 2017; Graham et al., 2013; Jungherr, 2023; Kreiss, Lawrence and McGregor, 2018). Short-form video pushes this dynamic further in two ways. First, the candidate becomes the content: their face, voice, body, and delivery become part of the message in ways campaign staff cannot reproduce. Second, recommendation-based feeds expose candidates far beyond their follower networks, making platform-native style increasingly consequential for how widely their content travels (Meta, 2026; TikTok, 2026).

The advice to campaign like influencers follows naturally from research on parasocial theory and platform affordances. Influencer-style communication is built on direct forms of social media labor. While these concepts differ in important ways, this paper is not primarily concerned with occupational distinctions among social media actors. Instead, I use the term “influencer-style communication” as a shorthand for a bundle of communicative practices commonly associated with successful social media personalities, including informal visual aesthetics and relational audience engagement.

address, self-disclosure, and performed authenticity. Communication scholars have long argued that these cues can foster parasocial relationships, defined as the felt intimacy audiences develop with media figures they have never met (Abidin, 2015; Horton and Wohl, 1956; Khamis, Ang and Welling, 2017; Marwick and boyd, 2011). Parasocial bonds, in turn, rank among the strongest predictors of candidate support, rivaling partisanship and issue agreement (Cohen and Holbert, 2021; Tsfati et al., 2022). Platform recommendation systems add another reason for candidates to adopt this style by rewarding content that looks native to the feed and invites engagement (Meta, 2026; Peng, 2021; TikTok, 2026). If influencer cues build parasocial bonds, parasocial bonds build support, and platforms amplify candidates who deploy the cues, then candidates who resist the advice are leaving votes on the table.

However, systematic evidence that candidates actually adopt this style, or that they benefit from doing so, remains surprisingly limited. The high-profile examples like Ocasio-Cortez, Mamdani, or Ossoff, are often young, digitally fluent, mostly progressive, and running in contexts favorable to their style. Their success reflects a mix of ideology, charisma, district characteristics, and campaign resources that is difficult to disentangle. Existing research shows that parasocial bonds predict candidate support (Cohen and Holbert, 2021; Tsfati et al., 2022), and that influencers themselves can move political attitudes (Cheng et al., 2024; Chmel et al., 2025). Yet it does not show that a candidate’s own adoption of the style improves how voters judge the candidates. The behavior of politicians themselves may suggest some quiet skepticism. Politicians have widely adopted social media, but most still use it for conventional, one-way campaign communication rather than embracing the informal, personal, and interactive style associated with influencers (Graham et al., 2013; Grusell and Nord, 2023). When mainstream officeholders reach for the full influencer playbook, the attempts are often dismissed as “cringe” (Miller, 2025), “inviting ridicule” (Tamasco, 2025), and failing the “teenager eye-roll test” (Vega, 2022).

There are good theoretical reasons to be skeptical, since voters evaluate candidates differently from how followers evaluate influencers. Research on candidate evaluation

shows that voters judge prospective officeholders on perceived competence inferred from remarkably thin cues such as faces, voices, attire, and delivery (Funk, 1997; Klofstad, 2017; Todorov et al., 2005), precisely the cues that text withheld and short-form video makes unavoidable. The same informality that signals authenticity for a lifestyle influencer can read as unprofessional for a prospective lawmaker, because violations of communication expectations are evaluated in light of the role and context in which they occur (Bullock and Hubner, 2020; Burgoon, 1993). In short, parasocial theory predicts that influencer style buys closeness, while research on candidate evaluation warns that it costs competence. The candidate quoted at the beginning captures the central uncertainty: “likes aren’t votes.” The signals platforms return to campaigns may not reflect the judgments that ultimately determine electoral support.

Therefore, this gap in the literature leaves several important questions unanswered. To what extent do candidates actually use influencer-style communication? Does it help them gain exposure and engagement on social media platforms? Most importantly, what effects does it have on candidate evaluation?

This paper provides the first systematic test of these questions through three studies. I begin with an observational analysis of a near-complete archive of 16,572 Instagram Reels,³ the short-form video platform most widely adopted by state legislative candidates,⁴ posted by 200 candidates from New Jersey and Virginia during the 2025 election cycle, the most recent legislative elections in the United States.⁵ The analysis decomposes

³These data cover candidates for state legislative office. I have secured funding to extend the same coding and analysis to congressional candidates to assess whether these patterns hold in higher-profile races.

⁴Appendix A demonstrates that state legislative candidates in the 2025 election cycle were roughly four times as likely to maintain an Instagram account as a TikTok one.

⁵State legislative candidates are a useful population because they are often less well-known and have fewer resources than national-level candidates, making them especially dependent on exactly the low-cost visibility the platforms promise.

influencer style communication into two dimensions. The visual component, which I call *aesthetic proximity*, captures the informal and platform-native presentation that makes a candidate look close to an ordinary citizen. The relational component, which I call *relational signaling*, captures communicative cues such as self-disclosure and interaction invitations that make a candidate communicate closeness. The analysis examines how widely candidates adopt each dimension and which dimensions the platform rewards.

The second and third studies test the causal effects of each dimension on voter evaluations through two large-sample pre-registered survey experiments. Until recently, isolating these effects has been infeasible, because real-world campaign videos bundle the candidate’s identity, face, voice, message, ideology, and production polish into a single stimulus. I address this problem with a modular generative AI pipeline that creates controlled audiovisual experimental stimuli. Specifically, I hold the candidate’s face, voice, and message constant, while varying the visual and relational components independently. Study 2 ($N = 1,952$) isolates aesthetic proximity and Study 3 ($N = 1,575$) manipulates aesthetic and relational cues independently in a 2×2 factorial design.

The results challenge the prominent recent calls for candidates to campaign like influencers. In the observational analysis, candidates widely adopt the informal visual cues but rarely use the relational cues at the heart of parasocial theory. Platform rewards reinforce the division. Informal visuals are associated with more views, likes, and comments, while relational cues produce weak and inconsistent returns. However, observational data cannot reveal whether any of this serves the candidates electorally. In the experiments, neither looking like an influencer nor communicating like one made respondents feel closer to the candidate, while both independently lowered perceived competence without delivering support. Combining them compounded the costs without producing the parasocial closeness or voter support either promised.

The paper makes three contributions. Theoretically, the findings suggest that traditional models of candidate evaluation centered on competence and office-holding capacity remain the dominant framework through which voters assess candidates, even on plat-

forms where influencer culture is the norm. Cues that demonstrably generate parasocial responses in non-political settings (Breves, Amrehn, Heidenreich, Liebers and Schramm, 2021; Breves, Liebers, Motschenbacher and Reus, 2021; Hartmann, 2016; Nah, 2022) do not transfer to politicians. Empirically, the paper offers the first systematic account of how candidates actually use influencer style communication cues through a novel dataset of Instagram Reels. Methodologically, the paper introduces a modular generative AI pipeline that builds controlled audiovisual stimuli, isolating the dimensions of interest while holding candidate identity and policy content fixed. To my knowledge, these are the first experiments to separate audiovisual style from candidate identity using fully synthetic videos. More broadly, the findings reveal that what platforms reward and what voters reward have come apart, and candidates can observe only one side in their account dashboards.

2 Candidates’ Dilemma in the Influencer Era

Politicians now face growing pressure to treat campaigning as content creation and produce platform-native videos in the style of social media influencers. David Plouffe, Barack Obama’s former campaign manager, argues that “a successful campaign in 2026 must operate like a full-time production studio,” with candidates spending hours each day filming content tailored to TikTok and Instagram (Plouffe, 2026). Some Republicans give the same advice, urging candidates to take an “authenticity risk” and go native on the platforms (Vale, 2026). Party organizations now directly distribute digital guides that specify posting frequency and timing (Labour, 2019), while party members worry about needing to “crack the code on reaching new audiences” (Wells, 2025).

Such pressure to adopt the influencer style may come from the fact that 43% of U.S. adults under thirty now regularly get news on TikTok (Tomasik and Matsa, 2025). Both presidential campaigns in 2024 treated short-form video as a primary channel for reaching voters who no longer watch cable news or read newspapers (Shah, 2024). Candidates can

appear in that feed without posting anything themselves, through opposition footage or fan edits that could shape voter evaluation (Munger and Li, 2025). However, that visibility arrives on terms the campaign does not set, and a candidate who does not post is present only as others choose to portray them. How candidates respond is therefore not a question of style alone. Social media has rapidly reshaped who reaches voters and how citizens experience political representation.

A small number of unusually fluent politicians have made influencer-style campaigning appear politically effective. Ocasio-Cortez turned casual Instagram Live sessions into a recognizable series of accessible political communication, where she answered constituent questions while zesting lemons and cooking dinner (Smith, 2020). Mamdani carried that approach into short-form video through fast and intimate clips, announcing his candidacy while jogging and promoting a rent freeze by plunging into the Atlantic Ocean in a suit (Gee, 2025; Wells and Hunter, 2025).

Yet even the most celebrated adopters use the influencer repertoire selectively. Ocasio-Cortez moves between informal platform-native communication and formal congressional hearings and floor speeches (De Fina, 2025), while Mamdani pairs the fast-paced clips with policy proposals (Wells and Hunter, 2025). Politicians who imitate the style without the same fluency of the exemplars often find the attempt turned against them. Videos by mainstream politicians have been dismissed as “cringe” (Haley and Hannah, 2025), “cheesy” (Vega, 2022), and “inviting ridicule” (Tamasco, 2025). The advice to campaign like an influencer is widely issued, unevenly followed, and dangerous to get wrong.

2.1 A Recurring Dilemma in Its Newest Form

Pressure to adapt to a new medium’s native style is a recurring political dilemma that predates social media. Candidates must appear capable of governing while remaining accessible to the people they represent, and those goals do not always align. Fenno (1978) described the management of this tension as “home style,” the effort to cultivate trust with constituents while preserving the latitude and legitimacy needed to govern.

Those efforts do not always succeed, and Kaslovsky (2022) finds that local visits may lower approval among ideologically opposed constituents rather than build the goodwill home style is meant to produce. New communication technologies repeatedly change how voters encounter politicians, forcing candidates to re-navigate this balance.

Radio carried officeholders' voices into the living room, and Franklin Roosevelt's fire-side chats were deliberately crafted to create a sense of conversational intimacy (Loviglio, 2005; Ryfe, 1999). Television then made candidates' appearance and delivery more consequential to voter evaluations (Mutz, 2015), with viewers of the Kennedy–Nixon debate more likely than audio listeners to judge Kennedy the winner (Druckman, 2003). Television also pushed political rhetoric away from formal public oratory toward personal narrative and conversational delivery (Jamieson, 1990). Ronald Reagan mastered this televisual style, translating abstract political themes into memorable stories told as if to someone sitting across the room (Jamieson, 1990). Later, the soft-news and talk-show circuit drew candidates into entertainment formats that reached inattentive voters often missed by traditional political news (Baum, 2005).

Social media removed the intermediary altogether, and lowered the cost of reaching voters, especially for lesser-known candidates. Candidates could communicate with voters continuously and on their own schedule (Castanho Silva and Proksch, 2022). Opening a Twitter account increased campaign contributions for congressional newcomers, and social media helped voters encounter candidates they might otherwise overlook (Petrova, Sen and Yildirim, 2021). These platforms also encouraged a more personal and informal style, as politicians increasingly blended private self-presentation with political messaging (Enli and Skogerbø, 2013; Metz, Kruikemeier and Lecheler, 2020). Donald Trump pushed this logic further on Twitter, turning an unpolished, self-authored style into a political asset (Enli, 2017). Campaigns could also observe audience reactions in real time, and political professionals increasingly treated platform metrics as a running signal of public response (McGregor, 2020).

Adapting to a new medium also carried risks, since campaigns could not know in

advance how voters would respond. A performance that made one politician appear more human could make another seem contrived or unserious, depending on the political context and the audience expectations (Young, 2020). Bill Clinton’s saxophone performance on late-night television became an iconic moment of the 1992 campaign (Lauter, 1992), while Gerald Ford’s attempt to eat an unshucked tamale became a lasting cautionary tale about performed relatability gone wrong (Renteria, 2025). Such episodes show that politicians cannot fully control the meaning audiences attach to their attempts at informality.

The influencer trend is the newest version of this recurring demand, but short-form video changes the structure of the problem in two ways. First, the candidate becomes the content. On text-based platforms, a staffer can write an informal post in the candidate’s voice without the candidate ever appearing. In short-form video, the candidate’s face, voice, body, and delivery are part of the message itself. A casual tweet depends largely on language, while a casual Reel also depends on what the candidate wears, where they film, how they hold the camera, and how they speak. Candidates also enter feeds shaped by influencer conventions, where personal disclosure, direct audience engagement, and invitations to interact are the norms (Abidin, 2015; boyd, 2026; Marwick and boyd, 2011). Adapting to these conventions requires candidates to perform on camera in a way that cannot simply be delegated to campaign staff. Second, distribution has shifted from follower networks toward algorithmic recommendation, making reach more contingent on content of candidates’ videos. Early social media mostly showed users content from accounts they had chosen to follow, while short-form platforms increasingly surface videos from unfamiliar accounts (Gerbaudo, 2026). TikTok’s For You feed is driven almost entirely by algorithmic curation (Bhandari and Bimo, 2020), and Meta reports that most content users now see on Instagram is recommended rather than posted by accounts they follow (Etherington, 2023). As a result, a candidate’s reach depends less on the audience they have already built and more on whether their videos fit the platform’s style.

Candidates receive direct feedback on content performance from the social media plat-

forms. Every post returns immediate and quantified measures of views, watch time, likes, comments, and shares, updated continuously and attributable to a specific video. The effects on candidate evaluation are much harder to observe. Whether a video left viewers more or less willing to see the candidate as fit for office may surface only indirectly in later polling or election results and is difficult to trace to any particular post. Campaigns can therefore identify which videos attract engagement far more easily than which strengthen or weaken voter evaluations. Whether candidates should follow the influencer advice ultimately depends on what the style does to the voters who see it. On that question, existing research suggests two different possibilities.

2.2 Why Influencer Style Could Help (and Hurt)

The appeal of influencer-style campaigning rests on a plausible promise, but the same cues that make a candidate seem authentic and relatable may also make that candidate appear less serious. Research on parasocial communication explains how influencer-style cues can foster feelings of closeness, while research on candidate evaluation shows that voters infer competence from many of the same presentational cues. These literatures have largely developed in different settings and rarely consider how these processes operate together.

2.2.1 Parasocial Closeness and the Appeal of Influencer Style

The appeal of influencer-style campaigning rests largely on parasocial theory. Horton and Wohl (1956) argued that media figures can create a relationship that feels social despite lacking reciprocity, an “intimacy at a distance” built when a media figure acknowledges the audience and simulates face-to-face interactions. Viewers understand that the relationship is one-sided, yet the experience of being personally addressed can persist and deepen across encounters (Horton and Wohl, 1956). Later work distinguished the momentary sense of parasocial interaction during a particular exposure from the more enduring parasocial relationship that develops through repeated contact (Dibble, Hartmann and Rosaen, 2016). Influencers cultivate this sense of parasocial relationship through a rec-

ognizable set of communicative practices. They disclose aspects of their personal lives, acknowledge viewers directly, invite responses and reply to them, and present carefully produced content as casual or unpolished on purpose, a practice Abidin (2017) calls “calibrated amateurism.” Audiences come to describe successful influencers less as distant celebrities than as friends or older siblings (Berryman and Kavka, 2017).

Research on social media influencers shows how particular cues build parasocial closeness. Direct verbal and visual address can intensify viewers’ immediate sense of connection with a media figure (Hartmann and Goldhoorn, 2011), and a single self-disclosing message can make a performer seem more authentic and likable and increase message receptivity (Nah, 2022). Greater self-disclosure is also associated with stronger parasocial relationships (Leite and Baptista, 2022). These bonds are stronger when audiences see media figures as similar and attractive, and when followers view their interactions with influencers as fair and respectful (Tukachinsky, Walter and Saucier, 2020; Yuan and Lou, 2020). In non-political settings, stronger parasocial relationships are associated with greater credibility, more favorable attitudes, stronger purchase intentions, and greater willingness to pass a message along (Balaban, Szabolcs and Chirică, 2022; Breves, Amrehn, Heidenreich, Liebers and Schramm, 2021; Yuan and Lou, 2020). These are precisely the communicative practices that candidates are encouraged to borrow from influencers. These cues can produce immediate relational responses in the context where influencer culture originates, and established parasocial relationships carry persuasive effects.

Parasocial relationships can also carry political consequences. In surveys of U.S. and Israeli voters, the intensity of parasocial relationships with political leaders was among the strongest predictors of candidate support, rivaling partisanship and issue agreement (Cohen and Holbert, 2021; Tsfati et al., 2022). Such attachment can also begin before a candidate enters politics and later carry electoral benefits. Exposure to Ronald Reagan as a television host improved perceptions of his leadership and increased his subsequent votes and campaign contributions (Xiong, 2021). More recently, familiarity with Donald Trump through *The Apprentice* was associated with trust and parasocial connection (Gabriel

et al., 2018), and exposure to that program increased his vote share in the 2016 Republican primary (Kim and Patterson, 2025). Social media influencers can also shape political attitudes. Parasocial relationships with political influencers are associated with greater receptivity to their messages and stronger intentions to share their advocacy (Cheng et al., 2024). A field experiment during the 2024 U.S. election further shows that encouraging adults to follow selected influencers shifted policy positions, partisan evaluations, and under some conditions, vote choice (Chmel et al., 2025).

Platform design offers another reason to adopt influencer-style communication. Recommendation systems rank content partly through signals such as likes, comments, shares, and watch time (Meta, 2026; TikTok, 2026). Personalized political content also performs well on social media. In an analysis of 59,020 Instagram posts from 172 U.S. politicians, Peng (2021) finds that images showing politicians’ faces, emotions, and private settings generally received more likes and comments than conventional professional content. Both algorithmic distribution and platform engagement tend to favor the cues that build parasocial closeness.

Together, these findings explain why influencer-style campaigning appears promising. If influencer-style cues foster parasocial relationships, which strongly predicts political support, and platforms reward such cues, candidates should benefit from adopting the style themselves. However, the evidence for each link comes from different actors and contexts. Influencer studies examine how commercial influencers cultivate parasocial bonds, while candidate surveys measure pre-existing parasocial relationships with prominent politicians without identifying how they are formed. None of these findings establishes that candidates themselves benefit when they adopt the full influencer repertoire.

2.2.2 Candidate Evaluation and the Importance of Competence

Voters do not evaluate candidates purely as social media influencers, but also as prospective officeholders. In Fenno’s account of “home style,” representatives build trust by remaining accessible to constituents, but that accessibility must coexist with confidence

in their ability to govern (Fenno, 1978). Practices that work for influencers may take on a different meaning when performed by someone seeking political power. Voters can infer candidates' capacity to govern from surprisingly limited information. Character traits such as competence, leadership, integrity, and warmth shape broader candidate evaluations, and these impressions are often filtered through partisan expectations (Funk, 1997, 1999; Hayes, 2005; Laustsen, 2021).

These judgments can form from remarkably thin presentational cues, making them especially relevant to short-form video, where viewers encounter many visual and vocal signals in a short time. Competence inferred from facial appearance alone predicts electoral outcomes and vote shares (Bilson et al., 2024; Mattes et al., 2010; Todorov et al., 2005), while impressions drawn jointly from candidates' faces and voices also affect vote choice (Klofstad, 2017). Nonverbal displays of confidence affect perceived qualifications and electability (Dumitrescu, Gidengil and Stolle, 2015), and even attire can alter assessments of a politician's personality and qualifications (Kaczmarek and Stencel, 2021). These impressions form quickly and draw on many of the same features associated with influencer-style communication. A candidate who delivers a policy message from a kitchen in a T-shirt may be doing more than adopting a platform-native casual aesthetic. They may simultaneously be changing the cues voters use to judge whether they appear capable of governing.

These presentational cues matter because voters interpret them against expectations about how candidates should behave. Expectancy violations theory holds that communication is evaluated relative to what audiences consider appropriate in a particular setting and relationship, and departures from those expectations can be received in different ways (Burgoon, 1993; Burgoon and Jones, 1976). Applied to candidates, practices that signal authenticity for influencers, such as casual presentation, personal disclosure, and unpolished delivery, may make politicians appear unserious and unprofessional. Consistent with this account, Bullock and Hubner (2020) find that humorous, informal posts by political candidates violated audience expectations and reduced candidates' perceived

credibility and voter support. The consequences of such informality therefore depend on whether voters interpret it as appealing accessibility or failure to meet the expectations of political office.

Expectancy violations theory also suggests that the same style need not affect every candidate in the same way. Audiences give more latitude to unexpected behavior from communicators they already view favorably, while the same departure may hurt someone regarded less positively (Burgoon, 1993). In politics, reactions may therefore depend on the candidate, office, context, and audience. A casual video may fit the image of a young insurgent more easily than of a senior institutional figure, and voters may differ in how much informality they consider appropriate from an office-seeker.

Neither the parasocial nor the candidate-evaluation literature directly tests whether effects found among influencers carry over to political candidates. Research on social media influencers shows that direct address, self-disclosure, and simulated interactions can build parasocial closeness (Abidin, 2015; Hartmann and Goldhoorn, 2011; Leite and Baptista, 2022). Research on candidate evaluation shows that voters infer competence from appearance, voice, and other thin presentational cues, and may penalize communication they consider inappropriate for an office-seeker (Bullock and Hubner, 2020; Klostad, 2017; Todorov et al., 2005). Both findings may be correct within the settings in which they were established. What remains unclear is whether practices that produce relational gains when used by influencers retain those benefits when performed by a candidate seeking power, or whether voters instead interpret them through the standards they apply to prospective officeholders. Answering that question requires holding the candidate and message constant while changing the style of presentation, a comparison difficult to achieve with real-world candidate videos.

2.3 Two Problems With Existing Evidence

Existing research has not resolved this tension in part because influencer-style communication has not been conceptualized consistently. Relevant studies have examined a

wide range of practices, from professional, private, and emotional forms of visual self-personalization (Farkas and Bene, 2021; Metz, Kruikemeier and Lecheler, 2020; Peng, 2021), to platform norms of liveness and authenticity (Lindholm, Carlson and Högväg, 2021; Zulli and Towner, 2021), to affordances and vernaculars specific to Instagram and TikTok (Cervi, 2023; Cervi, Tejedor and Blesa, 2023; Towner and Muñoz, 2024). Case studies focus on other practices, including Matteo Salvini’s recurring Instagram morning selfies (Starita and Trillò, 2022), Ada Colau’s adaptation to TikTok culture (Cervi, 2023), and Alexandria Ocasio-Cortez’s movement between institutional settings and the informality of everyday life (De Fina, 2025).⁶ These practices could all be justified as elements of influencer-style communication, but treating them as manifestations of the same concept makes findings difficult to compare and accumulate. When two studies reach different conclusions, it is often unclear whether the evidence genuinely conflicts or whether the studies examined different components.

A second problem is causal identification. Candidates choose how to present themselves strategically, so influencer-style videos may differ from conventional content in many ways besides style. Each post bundles together topic, emotional tone, background setting, production quality, and the particular politician delivering it. A successful influencer-style video may perform well because of its style, policy position, or candidate prominence. High-profile cases further illustrate the difficulty of isolating the effect of style. Mamdani and Ocasio-Cortez are unusually visible figures with distinctive ideologies and charisma that cannot be separated from their communication styles (Ceccobelli and Colombo, 2025; De Fina, 2025). Their success shows that influencer-style can coexist with political prominence, while its contribution to such prominence remains unclear.

Experiments provide greater causal leverage, though existing designs often preserve some of the same bundles found in observational research. Realistic stimuli preserve the audiovisual form, but allow treatments to vary in setting, activity, emotional expression,

⁶Salvini is the leader of Italy’s right-wing Lega party; Colau was then the left-wing incumbent mayor of Barcelona; Ocasio-Cortez is a Democratic Member of U.S. Congress.

and candidate identity simultaneously (Lindholm, Carlson and Högväg, 2021; Zulli and Towner, 2021). Among the few studies examining short-form political videos, Munger and Li (2025) show that stylized TikTok edits can change evaluations of prominent politicians, though these edits were produced by fans rather than the candidates themselves and varied visual, audio, and editing features together. Existing experiments therefore reveal that platform-specific images and videos can potentially shape candidate evaluations, while leaving uncertainty about which elements of the presentation produced the effects.

Addressing both problems requires a way of separating the components of influencer-style communication while preserving the audiovisual form in which voters encounter them.

2.4 Definitions and Research Questions

To address the first problem, I decompose influencer-style communication into two dimensions. *Aesthetic proximity* refers to visual cues that make a candidate look close to an ordinary citizen, including selfie framing, casual attire, everyday settings, and low production polish. These cues change how a candidate appears while leaving the substance of the message unchanged. *Relational signaling* refers to communicative cues that simulate interpersonal exchange, including direct address, personal disclosure, audience acknowledgment, identity bridging, and interaction invitations. These cues lie closest to the parasocial mechanism because they make the audience feel recognized and personally addressed rather than spoken to as part of a mass audience (Horton and Wohl, 1956; Hartmann and Goldhoorn, 2011). Though the two often appear together in influencer content, they are analytically distinct and can vary independently.

To address the second problem, I propose a modular generative pipeline that allows me change the aesthetic proximity and relational signaling cues while holding candidate identity, voice, and message constant. Section 3 describes how each dimension is measured in candidate videos and manipulated experimentally. This decomposition makes it possible to ask whether candidates adopt the two dimensions differently, whether plat-

forms reward them differently, and whether they have different effects on voters.

RQ1 (Adoption). To what extent do candidates use aesthetic proximity and relational signaling cues in their short-form videos, and do they adopt the two dimensions to the same degree?

RQ2 (Platform rewards). How is each dimension associated with reach and engagement, and does the platform reward the two dimensions equally?

RQ3a (Aesthetic proximity). How does looking like an influencer by adopting the aesthetic proximity cues affect parasocial closeness, perceived competence, and candidate support?

RQ3b (Relational signaling). How does communicating like an influencer by adopting the relational signaling cues affect parasocial closeness, perceived competence, and candidate support?

RQ4 (Interaction). Do the two dimensions interact, so that combining them differs from the sum of their separate effects?

The observational analysis addresses RQ1 and RQ2, describing how candidates communicate and how each dimension is associated with platform reach and engagement. The two preregistered experiments address voter response, with Study 2 estimating the effects of aesthetic proximity cues (RQ3a) and Study 3 estimating the effects of both dimensions and their interaction (RQ3a, RQ3b, and RQ4). The next section describes the data, stimuli, and research designs used in each study.

3 Data and Research Design

The two problems identified in Section 2 call for a design that measures influencer-style communication consistently and separates the stylistic cues from everything that normally travels with them. I address them with three studies that build on each other. Study 1 uses a novel and near-complete archive of candidate Instagram Reels to establish how widely candidates adopt each dimension and which dimensions the platform rewards.

Studies 2 and 3 take up the causal question observational data leaves open, using AI-generated treatment videos that hold the candidate and the message fixed while varying aesthetic proximity and relational signaling cues independently.

3.1 Study 1: Candidate Communication on Instagram Reels

Study 1 examines how candidates use aesthetic proximity and relational signaling cues in their own short-form videos (RQ1) and how each dimension is associated with platform reach and engagement (RQ2). The unit of analysis is the individual video.

3.1.1 Sample and Data Collection

The candidate universe consists of everyone who advanced to the November 2025 general elections for the New Jersey General Assembly and Virginia House of Delegates, the only two state legislative elections held that year.⁷ This includes 159 candidates in New Jersey and 196 in Virginia, two of whom later withdrew and excluded from analyses. These races are recent enough to reflect current platform norms and small enough to make near-complete collection of the candidate data feasible. Because they took place in an off-year cycle, candidates were also not competing for attention with a presidential race, making engagement more likely to reflect responses to the candidates and their content instead of a national campaign. State legislative candidates typically have fewer professional resources and lower name recognition than national candidates, and often manage their own accounts or work with small teams. Therefore, their posts provide a useful view of how candidates communicate when platform decisions remain largely under their control and when low-cost visibility may be especially valuable.

I focus on Instagram Reels because Instagram provided substantially broader coverage of this candidate population than TikTok while retaining the short-form vertical-video

⁷In ongoing work, I extend the same analysis to congressional candidates to assess whether these patterns hold in higher-profile offices.

format. An undergraduate research assistant searched for each candidate on both platforms and verified accounts against campaign websites, party listings, and other official sources. In New Jersey, 111 candidates (69.8%) had public Instagram accounts compared with 28 (17.6%) on TikTok. In Virginia, the corresponding numbers were 131 (67.5%) and 23 (11.9%). Appendix A reports the full platform-coverage and account-identification procedure.

Not every Instagram account contained video. As of February 22, 2026, publicly available Reels were identified for 99 New Jersey and 101 Virginia accounts. I collected the complete available Reel history for each of these accounts using the open-source Python library *instaloader*, producing an archive of 16,572 videos from 200 accounts. One New Jersey account is jointly maintained by two Assembly candidates.⁸ Candidate-level descriptives therefore refer to the remaining 199 individual candidate accounts. Among these accounts, 139 represent Democrats and 55 Republicans, with the remainder representing minor-party or independent candidates. 110 are incumbents and 89 challengers. Appendix A reports the complete candidate and account descriptives.

The archive captures how candidates running in the 2025 elections used Reels across their available account histories.⁹ The earliest Reel was posted on November 22, 2019. Posting activity is highly concentrated by more active accounts: accounts contain a median of 31 Reels (mean= 82.9, IQR= 66.5), with a range from 1 to 1,270.

3.1.2 Measuring Influencer-Style Communication

I measure aesthetic proximity and relational signaling using seven binary indicators, each coded independently as present or absent in a video. Two indicators capture aesthetic

⁸The Instagram account *njdistrict38* is the official Instagram page for New Jersey’s 38th Legislative District, and jointly maintained by Senator Joe Lagana, Assemblywoman Lisa Swain, and Assemblyman Chris Tully.

⁹In ongoing work, I include data on posting date, days relative to the general election, and video duration.

proximity and build on research distinguishing professional, private, and informal forms of politicians’ visual self-presentation (Farkas and Bene, 2021; Metz, Kruikemeier and Lecheler, 2020; Peng, 2021). I distinguish how a video is filmed from how the candidate appears in it, since the two do not always align.

Relational signaling combines five indicators derived from the parasocial literature: direct address, audience acknowledgment, personal self-disclosure, interaction invitations, and identity bridging (Hartmann and Goldhoorn, 2011; Horton and Wohl, 1956; Leite and Baptista, 2022). I first hand-coded a random sample of 100 Reels using this codebook, then applied the same coding instructions to Gemini-3.1-Pro, which classified the full audiovisual content of each video on all seven indicators. The aesthetic-proximity and relational-signaling indices are normalized sums of their component indicators and range from 0 to 1. Appendix B provides the complete codebook, examples, edge-case rules, human validation and alternative index specifications. Human and model labels achieved 95.1% overall agreement, and agreement levels were high across every indicator.¹⁰

3.1.3 Platform Outcomes and Statistical Approach

RQ2 examines four measures of platform performance: views, likes, comments, and reposts. Views capture video reach, while likes, comments, and reposts capture different forms of audience engagement. These measures are highly right-skewed and contain varying numbers of zero, so I transform each as $\log(y + 1)$. Appendix A reports their untransformed distributions.

Engagement metadata were available for 10,723 Reels from 160 candidate accounts. Accounts with and without usable engagement data are similar in party, race, perceived gender, incumbency, and election outcome. Appendix A reports the missing-data comparison and additional diagnostics. Engagement metrics were collected over five consecutive

¹⁰88% for informal production, 98% for informal candidate presentation, 90% for direct address, 96% for audience acknowledgment, 90% for self-disclosure, 100% for interaction invitation, and 98% for identity bridging.

days, February 23-27, 2026.

For each engagement outcome, I estimate ordinary-least-square models relating the aesthetic-proximity and relational-signaling indices to logged engagement with candidate-account fixed effects. Standard errors are clustered by account. The account fixed effects compare videos posted by the same candidate, absorbing stable differences in candidate prominence, party, district, and baseline follower size.

These estimates are descriptive rather than causal. Candidates choose how to present themselves, and videos that differ in communication style may also differ in other features that affect engagement. Study 1 therefore establishes patterns of adoption and within-candidate associations with observable platform feedback. Studies 2 and 3 address the causal effects of the same communicative dimensions on voter evaluations.

3.2 Studies 2 and 3: Experimental Designs

Studies 2 and 3 turn from how candidates communicate to how voters respond to the same two dimensions of influencer-style communication. Both use controlled audiovisual stimuli that hold the candidate and policy content constant while varying the communication style. Study 2 provides a focused test of aesthetic proximity cues, comparing high- and low-aesthetic presentations across multiple candidate videos. Study 3 then extends the design by independently manipulating aesthetic proximity and relational signaling cues in a 2×2 factorial experiment. Together, the two studies test whether looking like an influencer, communicating like one, or combining the two changes voter evaluation.

3.2.1 Samples and Experimental Designs

Study 2 was fielded through Lucid from February 13 - March 1, 2026, and recruitment did not use quota sampling.¹¹ Of the 2002 respondents who entered the study, those who failed to finish the survey were excluded from analysis, yielding a final analytic sample

¹¹The survey was fielded as part of the Graduate Weidenbaum Center Survey at Washington University in St. Louis.

of 1,952 respondents. Respondents were randomly assigned to either a gun-policy or reproductive-policy issue domain. Within their assigned domain, each respondent viewed four videos crossing two ideological positions with the high- and low-aesthetic treatments. The gun-policy videos used a male candidate avatar, while the reproductive-policy videos used a female candidate avatar. Because implied candidate sex is confounded with issue domain by design, I do not make causal comparisons across these features. The resulting data contain 7,808 respondent-video evaluations. Study 2 uses a mixed design, with each respondent evaluating four treatment videos, and standard errors are clustered by respondent in the treatment-effect models described below.

Study 3 was fielded on Prolific on May 15, 2026. Respondents were eligible if they were above 18 years old, resided in the United States, and used English as their primary language. Recruitment used quotas to balance respondents by sex. Of the 1,601 respondents recruited, those who failed to complete the survey were excluded from analysis, leaving a final analytic sample of 1,575. Study 3 uses a between-subjects 2×2 factorial design in which each respondent viewed a single video delivering the same cost-of-living policy message. Respondents were randomly assigned to one of four conditions: traditional broadcast (low aesthetic, low relational), aesthetic only (high aesthetic, low relational), relational only (low aesthetic, high relational), or the full influencer package (high aesthetic, high relational). Across conditions, the videos depict the same fictional female candidate with the same voice, while the aesthetic treatment changes the visual presentation and the relational treatment changes the communicative cues in the script. The relational scripts contain the same policy substance, matched in spoken length. This design allows the independent effects of aesthetic proximity and relational signaling, as well as their interaction, to be estimated.

3.2.2 Treatments and Stimulus Construction

The experiments manipulate the same two dimensions of influencer-style communication measured in Study 1. The aesthetic-proximity treatment changes how the candidate

appears on screen while leaving the substance of the message unchanged. In the low-aesthetic condition, the candidate is presented in a more conventional campaign style, with formal attire, horizontal framing, a professional setting, and a polished presentation. In the high-aesthetic condition, the same candidate appears in a more platform-native style, with vertical or selfie-style framing, casual attire, a more relaxed presentation, and an everyday setting. These features correspond to the production and candidate-presentation cues used to measure aesthetic proximity in Study 1.

The relational-signaling treatment changes how the candidate communicates with the audience. The low-relational script uses a conventional broadcast-style policy register and does not include the relational cues identified in Study 1. The high-relational script incorporates all five: direct address, audience acknowledgment, personal self-disclosure, identity bridging, and an invitation to interact. For example, the candidate acknowledges what viewers have been telling her, shares a personal experience related to the policy problem, explicitly places herself and viewers within a shared group, and invites viewers to respond. The underlying policy position and substantive proposals remain the same across the low- and high-relational scripts. In Study 2, only the aesthetic bundle is manipulated. Study 3 manipulates the aesthetic and relational bundles independently, making it possible to estimate their separate effects and their interaction.

The central design goal is to change these communication cues while holding other features of the candidate and message as constant as possible. Within each matched treatment comparison, the candidate’s facial identity and voice are held constant, as are the issue domain, ideological position, and policy substance. Video length and spoken duration are also matched as closely as possible across conditions. The result is a set of treatment contrasts in which the intended visual and relational cues change while many features that ordinarily vary across real campaign videos do not.

Creating these controlled comparisons requires separating elements that are normally bundled together in audiovisual campaign communication. I therefore use a modular generative-AI pipeline to construct fictional candidate videos in six stages. As Figure 1




Figure 1: Generation Pipeline.

shows, I first construct the treatment scripts and generate a consistent synthetic voice. I then create matched candidate images for the aesthetic conditions, animate the still images into talking-head videos, synchronize the candidate’s mouth movements to the audio, and validate the completed videos with human realism screening before fielding. Because each component is generated separately, the pipeline allows the script and visual presentation to vary without simultaneously changing the candidate’s identity, voice, or policy position.¹²

Appendix C documents the complete generation procedure, models and prompts, treatment scripts, and stimulus-selection process. The following sections describe the experimental outcomes and statistical approach, followed by the validation procedures used to assess whether the generated videos provided sufficiently comparable treatment stimuli.

3.2.3 Outcomes, Estimands, and Statistical Analysis

Studies 2 and 3 use the same four core outcomes, summarized in Table 1. The outcomes capture three theoretically distinct stages of candidate evaluation. Parasocial closeness

¹²All stimuli depict fictional candidates. No real politician’s face or voice was used or cloned in the generation process. Additional details on stimulus provenance, screening, and research safeguards are reported in Appendix C.

measures whether the candidate feels personally familiar, authentic, and socially close to the viewer. Because respondents encounter a fictional candidate during a single experimental session, I interpret this outcome as immediate parasocial closeness rather than an enduring parasocial relationship formed through repeated exposure. Perceived competence captures whether the candidate appears experienced, serious, professional, and electorally competitive. Finally, favorability and candidate support measure downstream evaluations of the candidate.

Table 1: Experimental Outcomes and Theoretical Roles

Outcome	Operationalization	Theoretical Role
Parasocial closeness	Four-item index capturing group membership, interpersonal comfort, naturalness, and authenticity	Relational benefit
Perceived competence	Four-item index capturing experience, seriousness, professionalism, and electoral competitiveness	Candidate-evaluation cost
Favorability	Single-item overall candidate evaluation	Downstream evaluation
Candidate support	Single-item willingness to vote, donate, or volunteer	Downstream evaluation

Note: Exact item wording, response categories, scale construction, and measurement diagnostics are reported in Appendix E.

For the two multi-item outcomes, I sum the four constituent items after coding them so that higher values indicate greater parasocial closeness or perceived competence. I then standardize the resulting indices within each study to have mean zero and standard deviation one. The two single-item outcomes are standardized in the same way, so all treatment effects can be interpreted in standard-deviation units. The experiments also include manipulation checks assessing whether respondents perceived the intended communication-style differences. Appendix E reports the complete survey instruments, exact manipulation-check wording, scale construction, and reliability diagnostics.

For Study 2, the primary estimand is the average effect of assignment to the high- rather than low-aesthetic-proximity treatment cue bundle on each outcome. Let $A_{ij} = 1$ indicate that respondent i viewed the high-aesthetic version of video j , and $A_{ij} = 0$ the

corresponding low-aesthetic version. For each standardized outcome, I estimate

$$Y_{ij} = \alpha + \beta A_{ij} + \epsilon_{ij}, \quad (1)$$

where β is the average treatment effect of the high-aesthetic-proximity cue bundle. Because each respondent evaluates four videos, the respondent-video evaluation is the unit of analysis and standard errors are clustered by respondent to account for dependence across evaluations contributed by the same participant. The primary specification includes no additional covariates.

Study 3 estimates the separate and joint effects of aesthetic proximity and relational signaling in the 2×2 factorial design. I use conventional 0/1 treatment coding, with $A_i = 1$ denoting assignment to high aesthetic proximity and $R_i = 1$ denoting assignment to high relational signaling. For each standardized outcome, I estimate

$$Y_i = \alpha + \beta_A A_i + \beta_R R_i + \beta_{AR}(A_i \times R_i) + \epsilon_i. \quad (2)$$

where α is the mean outcome in the traditional-broadcast condition, in which both aesthetic proximity and relational signaling cues are low. β_A is the effect of high aesthetic proximity cue bundle when relational signaling cues are low, while β_R is the effect of high relational signaling cues when aesthetic proximity cues are low. The interaction coefficient, β_{AR} , tests whether the effect of either dimension changes when the other is present. I estimate these models using OLS with heteroskedasticity-robust standard errors. The main results report these treatment contrasts together with estimated means for the four experimental conditions. Appendix H additionally reports average treatment effects for each dimension across both levels of the other treatment, along with the equivalent specification using centered $-0.5/+0.5$ treatment coding.

Both experiments were preregistered before fielding. Appendix E provides the preregistration links, complete survey instruments, and detailed scale construction.

3.2.4 Stimulus Validation and Realism

Because the experiments rely on synthetic audiovisual stimuli, I evaluated the completed videos on perceived realism before fielding. For Study 2, three trained research assistants independently screened 32 candidate videos for naturalness, facial consistency, movement, and overall plausibility through pairwise comparison, and the top-ranked stimuli were selected that best preserved the intended aesthetic contrast without salient generation artifacts. For Study 3, I conducted a separate realism pilot with 108 respondents recruited through Prolific. Appendix D reports the complete Study 2 screening procedure, the design and results of the Study 3 validation.

The Study 3 pilot evaluated the experimental videos against both authentic and synthetic benchmarks. Each respondent viewed nine videos in randomized order: one of the four experimental stimuli, four authentic candidate videos drawn from the Study 1 corpus, and four existing AI-generated political videos already circulating online. After each video, respondents rated it on a four-point scale ranging from “definitely AI-generated”(= 1) to “definitely real”(= 4). The experimental stimuli received an average realism rating of 1.94, compared with 3.04 for authentic candidate videos and 1.50 for the existing AI benchmark. Experimental videos were therefore rated 1.09 points less realistic than authentic candidate videos (95% CI [-1.27,-0.91]), but 0.45 points more realistic than the existing AI-generated political videos in-the-wild (95% CI [0.27,0.63]). Thus, although the experimental stimuli did not reproduce the perceived realism of human-recorded campaign videos, they were perceived as more realistic than the synthetic political content used as an external benchmark.

More importantly for interpreting the experimental treatment effects, perceived realism differed slightly across the aesthetic manipulation. High-aesthetic videos received an average realism rating of 1.74, compared with 2.15 for low-aesthetic videos, a difference of -0.41 points on the four-point scale (95% CI [-0.73,-0.09]). The corresponding difference across the relational-signaling manipulation was much smaller. High-relational videos averaged 1.98 compared with 1.91 for low-relational videos, a difference of 0.07 points

(95% CI [-0.25,0.40]). The point estimate provides no evidence that relational signaling systematically reduced perceived realism. By contrast, the aesthetic treatment slightly changed perceived artificiality in addition to the intended visual presentation. I therefore interpret the aesthetic estimate as the effect of the complete visual treatment bundle rather than as a perfectly isolated effect of informal presentation. Appendix D reports stimulus-level distributions and pairwise comparisons in more detail.

4 Results

Study 1 examines which dimensions of influencer-style communication candidates actually adopt (RQ1) and how these dimensions are associated with reach and engagement on the platform (RQ2). Studies 2 and 3 then turn to the voter side, testing whether the same aesthetic and relational cues affect parasocial closeness, perceived competence, and downstream candidate evaluations (RQ3a–RQ4).

4.1 Candidate Adoption and Platform Rewards

Study 1 reveals an asymmetry in how candidates adopt the two dimensions of influencer-style communication. Although 72.5% of the 16,572 Reels contain at least one of the seven cues, candidates rely much more heavily on aesthetic proximity than on relational signaling. Figure 2 (a), reports the prevalence of each cue. Informal production appears in 53.4% of videos, and informal candidate presentation appears in 47.8%. Relational cues are much less common, with direct address being the most frequent, appearing in 36.2% of videos, followed by self-disclosure at 22.7%. The cues that most explicitly simulate interpersonal exchange are even rarer: audience acknowledgment appears in only 11.4% of videos, identity bridging in 10.9%, and invitations to interact in 10.3%.

The two composite indices show the same divide. The aesthetic-proximity index averages 0.506, meaning that the typical Reel contains roughly one of its two component cues. The relational-signaling index averages only 0.183, corresponding to fewer than




Figure 2: Influencer-Style Communication in Candidate Instagram Reels

Note: Panel A reports the percentage of Instagram Reels containing each influencer-style communication cue. The archive contains 16,572 Reels posted by 200 accounts of state legislative candidates running in the 2025 off-year elections in New Jersey and Virginia. 72.5% of videos contain at least one of the seven cues. Panel B reports associations between the aesthetic proximity and relational signaling indices and log-transformed measures of views, likes, comments, and reposts. Models include account fixed effects with standard errors clustered at the account level and use the 10,723 videos successfully matched to post-level engagement metadata.

one of its five component cues. Candidates therefore appear more willing to look like influencers than to communicate with audiences in the more relational manner associated with influencer culture.

The joint distribution makes this selective adoption especially clear. Using 0.5 as the threshold for high use of each dimension, 50.2% of Reels combine high aesthetic proximity with low relational signaling cues, making this by far the most common influencer-style configuration. Another 40.0% are low on both dimensions. By contrast, only 8.4% combine high levels of both aesthetic and relational cues, while just 1.4% pair high relational signaling with low aesthetic proximity. The typical adaptation to short-form video is therefore not the full influencer repertoire. Candidates frequently adopt the informal visual conventions of the medium while leaving the more explicitly relational cues behind. Appendix C reports the complete joint distribution and shows that this descriptive asymmetry persists under alternative thresholds for defining high use of each dimension.

The platform feedback candidates receive points in the same direction. Figure 2 (b) reports associations between each communication cue and four measures of platform performance: views, likes, comments, and reposts. Each model includes candidate-account

fixed effects, so the estimates compare differently styled videos posted by the same candidate, and standard errors are clustered by account. The outcomes are modeled as $\log(y + 1)$ to account for zeros in the platform feedback data. These estimates remain descriptive, since candidates choose how to present themselves, and differently styled posts may differ in topic, timing, or other features that also affect engagement.

Aesthetic proximity is consistently associated with greater platform performance. Moving from zero to one on the aesthetic-proximity index is associated with approximately 18% more views ($\beta = 0.164$, $SE = 0.070$, $p = .018$), 26% more likes ($\beta = 0.233$, $SE = 0.054$, $p < .001$), and 21% more comments ($\beta = 0.194$, $SE = 0.041$, $p < .001$). The association with reposts is statistically indistinguishable from zero ($\beta = -0.002$, $SE = 0.027$, $p = .953$). Across three of the most visible forms of platform feedback, videos more resemble the informal visual style of short-form social media perform better than conventional videos posted by the same candidate.

Relational signaling receives much weaker and less consistent returns. A full-range shift across the index is negatively associated with views, though the estimate does not reach conventional levels of statistical significance ($\beta = -0.400$, $SE = 0.220$, $p = .070$). The corresponding associations with likes ($\beta = -0.055$, $SE = 0.155$, $p = .723$) and comments ($\beta = 0.018$, $SE = 0.175$, $p = .920$) are small relative to their uncertainty. Reposts provide the one exception, where greater relational signaling is associated with approximately 9% more reposts ($\beta = 0.086$, $SE = 0.043$, $p = .046$).

Taken together, the adoption and engagement results reveal a striking alignment on the platform side of the candidate’s decision. Candidates make greater use of aesthetic proximity cues, the dimension most consistently associated with higher platform rewards, while relational signaling cues are less common and have weaker, less consistent associations with engagement. While the observational patterns may capture the incentive environment campaigns face on the platforms, they cannot show whether such styles improve voter evaluation of candidates. Studies 2 and 3 turn to the consequence that platform metrics cannot reveal: how voters evaluate candidates when those same communication

styles are experimentally varied.

4.2 Study 2: Effects of Aesthetic Proximity

Study 2 tests whether adopting the visual cues of influencer-style communication changes how voters evaluate a candidate (RQ3a). Figure 3 reports standardized treatment effects comparing the high- and low-aesthetic-proximity conditions. The manipulation worked as intended: respondents rated candidates in the high-aesthetic condition as substantially more influencer-like ($\beta = 0.230$, $SE = 0.020$, $p < .001$). Full regression results and item-level estimates are reported in Appendix G.




Figure 3: Effects of Aesthetic Proximity on Candidate Evaluations. Points report standardized OLS estimates comparing the high-aesthetic-proximity condition with the low-aesthetic-proximity condition. Bars show 95% confidence intervals, with standard errors clustered by respondent. All outcomes are standardized within Study 2. Full regression results are reported in Appendix G.

The aesthetic treatment produced no corresponding gain in parasocial closeness. Instead, closeness declined slightly by 0.050 standard deviations ($SE = 0.015$, $p < .001$). The much larger effect appeared in perceived competence, where the high-aesthetic condition reduced the competence index by 0.366 standard deviations ($SE = 0.018$, $p < .001$). Thus, making the same candidate look more native to the short-form social media environment did not make respondents feel closer to the candidate, but substantially weakened perceptions of competence.

These competence costs were accompanied by smaller declines in downstream evalu-

ations. High aesthetic proximity cues reduced favorability by 0.092 standard deviations ($SE = 0.016$, $p < .001$) and candidate support by 0.087 standard deviations ($SE = 0.015$, $p < .001$). Study 2 therefore provides little evidence for the relational benefit commonly attributed to influencer-style communication and clear evidence of an evaluative cost. The visual cues that make a candidate appear more influencer-like lower perceived competence and also reduce favorability and political support.

4.3 Study 3: Aesthetic Proximity and Relational Signaling

Study 3 separates the two dimensions of influencer-style communication to test whether aesthetic proximity (RQ3a), relational signaling (RQ3b), or their combination (RQ4) changes how voters evaluate a candidate. I organize the results around the central theoretical tradeoff: whether these cues produce the expected gain in parasocial closeness, whether they instead carry a cost to perceived competence, and whether either effect extends to favorability and political support. Figure 4 reports estimated means across the four experimental conditions. With 0/1 treatment coding, the lower-order coefficients compare the aesthetic-only and relational-only conditions, respectively, with the low-aesthetic-low-relational baseline. Full regression results are reported in Appendix H.

The experiment provides little evidence that either dimension generates the relational benefit assumed by the influencer advice. When relational signaling is low, aesthetic proximity cues increased the parasocial-closeness index by 0.122 standard deviations, but the estimate is imprecise and does not reach conventional levels of statistical significance ($SE = 0.068$, $p = .074$). When aesthetic proximity is low, adding relational signaling cues produces an estimate close to zero ($\beta = 0.034$, $SE = 0.071$, $p = .627$). The interaction is likewise small and statistically indistinguishable from zero ($\beta = -0.055$, $SE = 0.101$, $p = .587$). As Figure 4 shows, estimated closeness remains similar across the four conditions. Thus, neither looking like an influencer nor adopting the relational language most closely associated with parasocial communication provides clear evidence of greater parasocial closeness with the candidate.




Figure 4: Estimated Outcomes Across the Four Experimental Conditions. Points show estimated condition means for standardized outcomes, with 95% confidence intervals based on heteroskedasticity-robust standard errors. Outcomes are standardized to have a mean of zero and a standard deviation of one. Low and high on the x-axis refer to the aesthetic proximity condition. Full regression tables and item-level results are reported in Appendix H.

The clearest effects instead appear in perceived competence. Relative to the baseline condition, adopting aesthetic-proximity cues while holding relational signaling low reduced perceived competence by 0.251 standard deviations ($SE = 0.064$, $p < .001$). Relational signaling produced an even larger penalty when aesthetic proximity was low, reducing competence by 0.363 standard deviations ($SE = 0.066$, $p < .001$). The interaction between the two dimensions is close to zero ($\beta = -0.028$, $SE = 0.098$, $p = .778$), providing little evidence that the effects of either dimension depends on the presence of the other. The fitted condition means make the cumulative pattern particularly clear, where the full influencer condition is approximately 0.64 standard deviations below the baseline condition in perceived competence.

The downstream effects on favorability and support are less definitive. Aesthetic proximity cues have little effect on either outcome when relational signaling is low, with estimates of 0.018 standard deviations for favorability ($SE = 0.066$, $p = .780$) and -0.022 for support ($SE = 0.068$, $p = .750$). Relational signaling produces larger negative estimates when aesthetic proximity is low, reducing favorability by 0.129 standard deviations ($SE = 0.070$, $p = .067$) and support by 0.123 standard deviations ($SE = 0.071$, $p = .082$), although neither estimate reaches the conventional threshold for statistical

significance. The interactions are again small and statistically insignificant. Compared with the baseline condition, the full influencer package reduces both favorability and support by approximately 0.18 standard deviations ($p = .007$ and $p = .008$, respectively; see Appendix H).

Study 2 found statistically significant reductions in both favorability and support from aesthetic proximity cues, whereas Study 3 does not reproduce those downstream aesthetic effects. The two experiments agree more clearly on perceived competence. Aesthetic proximity reduces competence across both designs, while Study 3 additionally shows that relational signaling carries a substantial competence penalty. Whether these competence costs translate into lower favorability and support is less consistent across the two studies.

Taken together, Study 3 provides little evidence for the expected relational payoff from influencer-style communication or for an interaction that makes the two dimensions more effective when combined. Relational signaling does not clearly increase parasocial closeness and instead reduces perceived competence (RQ3b), while aesthetic proximity again produces a competence penalty (RQ3a). The small interaction terms provide little evidence that combining the dimensions changes the effect of either one (RQ4). The full influencer package therefore accumulates substantial competence costs without producing the expected increase in closeness.

5 Discussion and Conclusion

Candidates are increasingly encouraged to communicate like social media influencers, but the three studies in this paper provide little evidence that adopting the full influencer repertoire improves voter evaluation or voter support. In Study 1, candidates adapt to influencer culture selectively. Informal and platform-native visual cues are commonly used, while relational cues remain comparatively rare. Platform feedback shows the same division. Aesthetic proximity is associated with more views, likes, and comments, whereas returns to relational signaling are weaker and less consistent. In the experiments, neither

dimension reliably increases parasocial closeness, while both reduce perceived competence. Downstream effects on favorability and support are less consistent across Studies 2 and 3, but the competence penalty appears large and significant across both designs. In other words, communication practices that are platform native do not necessarily improve how the politician is evaluated.

The central implication is that the effects of influencer-style communication depend on the role of the person using it. Influencer-style practices appear attractive because direct address, self-disclosure, informality, and simulated interaction can make media figures feel more accessible and closer (Horton and Wohl, 1956; Hartmann and Goldhoorn, 2011; Nah, 2022). Yet candidates are also evaluated as prospective officeholders, and voters infer competence from thin presentational cues (Todorov et al., 2005; Klofstad, 2017). The results instead suggest that evidence of these cues fostering parasocial closeness with influencers and other media figures does not imply that they will have the same effects when used by politicians. Practices that signal authenticity for an influencer may signal unseriousness for someone seeking political power. Short-form video changes the conventions through which politicians reach voters without eliminating the standards voters use to judge whether they appear fit to govern.

This places influencer-style campaigning within a much older tension in politics: candidates need to seem approachable without making themselves look less capable of governing (Fenno, 1978). New media change the tools available for managing that balance without dissolving the underlying tension. The influencer era is distinctive because candidates are encouraged not only to communicate more directly with the audience, but to borrow practices mainly developed for a different role. The success of these communicative cues among influencers may not translate directly into electoral politics. Candidates may still benefit from appearing native to short-form video without adopting the full influencer package. The results may help explain why previous research on political personalization has produced mixed results. Looking informal and communicating closeness are distinct cues, and they need not produce the same consequences for candidate

evaluation.

The combined evidence also reveals a broader problem of political campaigning on social media platforms. Social media gives campaigns unusually precise feedback about what attracts attention and engagement through visible measures of views, likes, comments, and shares. Effects on perceived competence, favorability, or willingness to vote are much harder to observe and attribute to a particular video. While Study 1 cannot establish that aesthetic cues causally increases engagement, the descriptive pattern is still consequential. The aesthetic cues most consistently associated with observable platform rewards are also the cues that consistently reduce competence when experimentally manipulated. Campaigns therefore operate under a feedback asymmetry. They can learn quickly which communication choices perform well according to platform metrics while receiving much weaker signals about whether those choices improve the political evaluations they ultimately care about. Attention and persuasion should not be treated as interchangeable outcomes.

Methodologically, the paper demonstrates how generative tools can make difficult audiovisual features experimentally manipulable. Real candidate videos often bundle communication style with appearance, voice, setting, ideology, message, and production quality. The modular pipeline proposed in this paper allows several of these features to be held constant while visual presentation and relational language are varied independently. This pipeline provides a reusable strategy for studying political communication in the audiovisual form in which voters increasingly encounter it. The same architecture can extend to other questions that vary presentation, candidate characteristics, or delivery while preserving identity and message.

Several limitations define the scope of these conclusions. The experiments capture immediate reactions to unfamiliar fictional candidates after limited exposure, while parasocial relationships may develop through repeated encounters. The treatments also manipulate cue bundles rather than individual components, leaving future work to isolate which specific features drive these effects. In addition, the synthetic videos were perceived as less

realistic than authentic campaign footage, and realism was somewhat lower in the high-aesthetic conditions in Study 3. The aesthetic estimate should therefore be interpreted as the effect of the broader visual treatment bundle rather than a perfectly isolated effect of aesthetic proximity. Finally, the studies cannot determine whether greater algorithmic reach could compensate for a worse impression among those reached. A lesser-known candidate might benefit from reaching more voters even if each exposure modestly reduces perceived competence. Testing that possibility requires designs connecting communication style to repeated exposure, distribution, and ultimately electoral behavior. These questions remain important directions for future work in political communication.

Each new communication medium, from radio and television to talk shows and social media, has required candidates to adapt how they present themselves to voters. Short-form video and the influencer trend create new opportunities for politicians to reach audiences who might otherwise never encounter them. The challenge is that platforms clearly show campaigns which communication styles attract attention and engagement, but provide much less information about how those same choices affect voters' evaluations of the candidate. The central question is therefore not whether politicians should adapt to social media, but how they can become native to the medium without mistaking the incentives of the platform for the judgments of the electorate. What looks successful in the feed need not make a candidate look fit for office. Likes are not votes.

References

- Abidin, Crystal. 2015. “Communicative Intimacies: Influencers and Percieved Interconnectedness.”
- Abidin, Crystal. 2017. “#familygoals: Family Influencers, Calibrated Amateurism, and Justifying Young Digital Labor.” *Social Media + Society* 3(2):2056305117707191.
- Balaban, Delia Cristina, Julia Szambolics and Mihai Chirică. 2022. “Parasocial relations and social media influencers’ persuasive power. Exploring the moderating role of product involvement.” *Acta Psychologica* 230:103731.
- Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost and Joshua A. Tucker. 2019. “Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data.” *American Political Science Review* 113(4):883–901.
- Baum, Matthew A. 2005. “Talking the Vote: Why Presidential Candidates Hit the Talk Show Circuit.” *American Journal of Political Science* 49(2):213–234.
- Berryman, Rachel and Misha Kavka. 2017. “‘I Guess A Lot of People See Me as a Big Sister or a Friend’: the role of intimacy in the celebrification of beauty vloggers.” *Journal of Gender Studies* 26(3):307–320.
- Bhandari, Aparajita and Sara Bimo. 2020. “TikTok and the “Algorithmized Self”: A New Model of Online Interaction.” *AoIR Selected Papers of Internet Research* .
- Bilson, Thomas, Birendra Rai, Russell Smyth and Liang Choon Wang. 2024. “Competence in the Eye of the Electorate: Appearance, Incumbency, and Vote Shares.” *Legislative Studies Quarterly* 49(1):187–202.
- boyd, danah. 2026. “Social Media Is Now Parasocial Media.” *Social Media + Society* 12(2):20563051261437487.

- Breves, Priska, Jana Amrehn, Anna Heidenreich, Nicole Liebers and Holger Schramm. 2021. "Blind trust? The importance and interplay of parasocial relationships and advertising disclosures in explaining influencers' persuasive effects on their followers." *International Journal of Advertising* 40(7):1209–1229.
- Breves, Priska, Nicole Liebers, Bernadette Motschenbacher and Leonie Reus. 2021. "Reducing Resistance: The Impact of Nonfollowers' and Followers' Parasocial Relationships with Social Media Influencers on Persuasive Resistance and Advertising Effectiveness." *Human Communication Research* 47(4):418–443.
- Bullock, Olivia M. and Austin Y. Hubner. 2020. "Candidates' use of informal communication on social media reduces credibility and support: Examining the consequences of expectancy violations." *Communication Research Reports* 37(3):87–98.
- Burgoon, Judee K. 1993. "Interpersonal Expectations, Expectancy Violations, and Emotional Communication." *Journal of Language and Social Psychology* 12(1-2):30–48.
- Burgoon, Judee K. and Stephen B. Jones. 1976. "Toward a Theory of Personal Space Expectations and Their Violations." *Human Communication Research* 2(2):131–146.
- Castanho Silva, Bruno and Sven-Oliver Proksch. 2022. "Politicians unleashed? Political communication on Twitter and in parliament in Western Europe." *Political Science Research and Methods* 10(4):776–792.
- Ceccobelli, Diego and Arianna Colombo. 2025. "Political Leaders or Social Media Influencers? The Cases of Alexandria Ocasio-Cortez and Chiara Ferragni." *Polity* 57(2):253–281.
- Cervi, Laura. 2023. "TikTok Use in Municipal Elections: From Candidate-Majors to Influencer-Politicians." *Más Poder Local* (53):8–29.
- Cervi, Laura, Santiago Tejedor and Fernando García Blesa. 2023. "TikTok and Political

- Communication: The Latest Frontier of Politainment? A Case Study.” *Media and Communication* 11(2):203–217.
- Cheng, Zicheng, Jin Chen, Rachel X. Peng and Heather Shoenberger. 2024. “Social media influencers talk about politics: Investigating the role of source factors and PSR in Gen-Z followers’ perceived information quality, receptivity and sharing intention.” *Journal of Information Technology & Politics* 21(2):117–131.
- Chmel, Kirill, Eunji Kim, John Marshall, Tiffany Fisher-Love and Nathaniel Lubin. 2025. “How Social Media Creators Shape Mass Politics: A Field Experiment during the 2024 US Elections.”. Working Paper.
URL: <https://papers.ssrn.com/abstract=5932534>
- Cohen, Jonathan and R. Lance Holbert. 2021. “Assessing the Predictive Value of Parasocial Relationship Intensity in a Political Context.” *Communication Research* 48(4):501–526.
- De Fina, Anna. 2025. “Bonding with followers: Chronotopes and scales in political communication on Instagram.” *Discourse Studies* 27(2):192–213.
- Dibble, Jayson L., Tilo Hartmann and Sarah F. Rosaen. 2016. “Parasocial Interaction and Parasocial Relationship: Conceptual Clarification and a Critical Assessment of Measures.” *Human Communication Research* 42(1):21–44.
- Druckman, James N. 2003. “The Power of Television Images: The First Kennedy-Nixon Debate Revisited.” *The Journal of Politics* 65(2):559–571.
- Dumitrescu, Delia, Elisabeth Gidengil and Dietlind Stolle. 2015. “Candidate Confidence and Electoral Appeal: An Experimental Study of the Effect of Nonverbal Confidence on Voter Evaluations.” *Political Science Research and Methods* 3(1):43–52.
- Enli, Gunn. 2017. “Twitter as arena for the authentic outsider: exploring the social

- media campaigns of Trump and Clinton in the 2016 US presidential election.” *European Journal of Communication* 32(1):50–61.
- Enli, Gunn Sara and Eli Skogerbo. 2013. “Personalized Campaigns in Party-Centred Politics: Twitter and Facebook as arenas for political communication.” *Information, Communication & Society* 16(5):757–774.
- Etherington, Darrell. 2023. “Meta says time spent on Instagram grew 24% thanks to TikTok-style AI Reel recommendations | TechCrunch.”
URL: <https://techcrunch.com/2023/04/26/meta-says-time-spent-on-instagram-grew-24-thanks-to-tiktok-style-ai-reel-recommendations/>
- Farkas, Xénia and Márton Bene. 2021. “Images, Politicians, and Social Media: Patterns and Effects of Politicians’ Image-Based Political Communication Strategies on Social Media.” *The International Journal of Press/Politics* 26(1):119–142.
- Fenno, Richard F. 1978. *Home Style: House Members in Their Districts*. New York: Scott Foresman & Co.
- Funk, Carolyn L. 1997. “Implications of Political Expertise in Candidate Trait Evaluations.” *Political Research Quarterly* 50(3):675–697.
- Funk, Carolyn L. 1999. “Bringing the Candidate into Models of Candidate Evaluation.” *The Journal of Politics* 61(3):700–720.
- Gabriel, Shira, Elaine Paravati, Melanie C. Green and Jason Flomsbee. 2018. “From Apprentice to President: The Role of Parasocial Connection in the Election of Donald Trump.” *Social Psychological and Personality Science* 9(3):299–307.
- Gee, Andre. 2025. “People Are Discovering NYC Mayoral Candidate Zohran Mamdani’s Rap Past.”
URL: <https://au.rollingstone.com/culture/culture-features/zohran-mamdani-rap-mr-cardamom-78345/>

- Gerbaudo, Paolo. 2026. "TikTok and the algorithmic transformation of social media publics: From social networks to social interest clusters." *New Media & Society* 28(3):1019–1036.
- Graham, Todd, Marcel Broersma, Karin Hazelhoff and Guido van 't Haar. 2013. "Between Broadcasting Political Messages and Interacting with Voters: The use of Twitter during the 2010 UK general election campaign." *Information, Communication & Society* 16(5):692–716.
- Grusell, Marie and Lars Nord. 2023. "Not so Intimate Instagram: Images of Swedish Political Party Leaders in the 2018 National Election Campaign." *Journal of Political Marketing* 22(2):92–107.
- Haley, Brown and Fierick Hannah. 2025. "Andrew Cuomo blasted for 'cringe' NYC campaign ad."
URL: <https://nypost.com/2025/07/16/us-news/clunker-andrew-cuomo-blasted-for-cringe-inauthentic-campaign-ad-where-he-jump-starts-mans-car/>
- Hartmann, Tilo. 2016. Parasocial Interaction, Parasocial Relationships, and Well-Being. In *The Routledge Handbook of Media Use and Well-Being*. Routledge.
- Hartmann, Tilo and Charlotte Goldhoorn. 2011. "Horton and Wohl Revisited: Exploring Viewers' Experience of Parasocial Interaction." *Journal of Communication* 61(6):1104–1121.
- Hayes, Danny. 2005. "Candidate Qualities through a Partisan Lens: A Theory of Trait Ownership." *American Journal of Political Science* 49(4):908–923.
- Horton, Donald and R. Richard Wohl. 1956. "Mass Communication and Para-Social Interaction." *Psychiatry* .
- Jamieson, Kathleen Hall. 1990. *Eloquence in an Electronic Age: The Transformation of Political Speechmaking*. New York: Oxford University Press.

- Jungherr, Andreas. 2023. Chapter 25: Digital campaigning: how digital media change the work of parties and campaign organizations and impact elections. In *Research Handbook on Digital Sociology*. Edward Elgar Publishing.
- Kaczmarek, Bożydar L. J. and Marcin Stencel. 2021. “I believe what I see: Impact of attire and academic degree on evaluation of personality traits of a politician.” *Current Issues in Personality Psychology* 10(4):311–320.
- Kaslovsky, Jaclyn. 2022. “Senators at Home: Local Attentiveness and Policy Representation in Congress.” *American Political Science Review* 116(2):645–661.
- Khamis, Susie, Lawrence Ang and Raymond Welling. 2017. “Self-branding, ‘micro-celebrity’ and the rise of Social Media Influencers.” *Celebrity Studies* 8(2):191–208.
- Kim, Eunji and Shawn Patterson. 2025. “The American Viewer: Political Consequences of Entertainment Media.” *American Political Science Review* 119(2):917–931.
- Klofstad, Casey A. 2017. “Looks and Sounds Like a Winner: Perceptions of Competence in Candidates’ Faces and Voices Influences Vote Choice.” *Journal of Experimental Political Science* 4(3):229–240.
- Kreiss, Daniel, Regina G. Lawrence and Shannon C. McGregor. 2018. “In Their Own Words: Political Practitioner Accounts of Candidates, Audiences, Affordances, Genres, and Timing in Strategic Social Media Use.” *Political Communication* 35(1):8–31.
- Labour. 2019. “Digital Campaigning Guide - Labour.”. “https://labour.org.uk/wp-content/uploads/2019/11/PRINT-READY-12671_19-Digital-Guide-2019-v3-IN-HOUSE.pdf”.
- Laustsen, Lasse. 2021. “Candidate Evaluations Through the Lens of Adaptive Followership Psychology: How and Why Voters Prefer Leaders Based on Character Traits.” *Political Psychology* 42(S1):109–148.

Lauter, David. 1992. "COLUMN ONE : Clinton's Strategy of Triage : 'Don't be greedy little pigs,' remarked an adviser, abandoning some big states as probably hopeless. Those that got immediate, massive attention were key to winning White House." *Los Angeles Times* .

URL: <https://www.latimes.com/archives/la-xpm-1992-11-05-mn-1584-story.html>

Leite, Fernanda Polli and Paulo de Paula Baptista. 2022. "The effects of social media influencers' self-disclosure on behavioral intentions: The role of source credibility, parasocial relationships, and brand trust." *Journal of Marketing Theory and Practice* 30(3):295–311.

Lindholm, Jenny, Tom Carlson and Joachim Högväg. 2021. "See Me, Like Me! Exploring Viewers' Visual Attention to and Trait Perceptions of Party Leaders on Instagram." *The International Journal of Press/Politics* 26(1):167–187.

Loviglio, Jason. 2005. *Radio's Intimate Public: Network Broadcasting and Mass-Mediated Democracy*. University Of Minnesota Press.

Marwick, Alice and danah boyd. 2011. "To See and Be Seen: Celebrity Practice on Twitter." *Convergence* 17(2):139–158.

Mattes, Kyle, Michael Spezio, Hackjin Kim, Alexander Todorov, Ralph Adolphs and R. Michael Alvarez. 2010. "Predicting Election Outcomes from Positive and Negative Trait Assessments of Candidate Images." *Political Psychology* 31(1):41–58.

McGregor, Shannon C. 2020. "'Taking the Temperature of the Room': How Political Campaigns Use Social Media to Understand and Represent Public Opinion." *Public Opinion Quarterly* 84(S1):236–256.

Meta. 2026. "Instagram Feed Recommendations."

URL: <https://transparency.meta.com/features/explaining-ranking/ig-feed-recommendations/>

- Metz, Manon, Sanne Kruikemeier and Sophie Lecheler. 2020. “Personalization of politics on Facebook: examining the content and effects of professional, emotional and private self-personalization.” *Information, Communication & Society* 23(10):1481–1498.
- Miller, Andrew. 2025. “Social media explodes after ‘cringe’ TikTok video of AOC, House Dems goes viral: ‘Couldn’t get any lamer’.”
URL: <https://www.foxnews.com/politics/social-media-explodes-after-cringe-tiktok-video-aoc-house-dems-goes-viral>
- Munger, Kevin and Valerie Li. 2025. “Thirst Traps and Quick Cuts: The Effects of TikTok “Edits” on Evaluations of Politicians.” *Social Media + Society* 11(2):20563051251329990.
- Mutz, Diana C. 2015. *In-Your-Face Politics: The Consequences of Uncivil Media*. Princeton: Princeton University Press.
- Nah, Hye Soo. 2022. “The appeal of “real” in parasocial interaction: The effect of self-disclosure on message acceptance via perceived authenticity and liking.” *Computers in Human Behavior* 134:107330.
- Parnes, Amie. 2026. “Influencers are pushing out media for aspiring Democratic politicians.”
URL: <https://thehill.com/homenews/campaign/5792552-democrats-using-creators-reach-voters/>
- Peng, Yilang. 2021. “What Makes Politicians’ Instagram Posts Popular? Analyzing Social Media Strategies of Candidates and Office Holders with Computer Vision.” *The International Journal of Press/Politics* 26(1):143–166.
- Perrett, Connor. 2020. “Jon Ossoff is betting that TikTok can boost young voters in Georgia’s heated runoff race — and his chances don’t look bad.”

- Petrova, Maria, Ananya Sen and Pinar Yildirim. 2021. “Social Media and Political Contributions: The Impact of New Technology on Political Competition.” *Management Science* 67(5):2997–3021.
- Plouffe, David. 2026. “Opinion | The Old Way of Campaigning Won’t Cut It Anymore.” *The New York Times* .
URL: <https://www.nytimes.com/2026/04/05/opinion/politics-midterms-tiktok-attention-content.html>
- Renteria, Melissa. 2025. “President Ford’s ‘Great Tamale Incident’ happened 49 years ago.” *San Antonio Express News* .
URL: <https://www.expressnews.com/news/article/san-antonio-history-gerald-ford-tamale-husk-20255924.php>
- Ryfe, David Michael. 1999. “Franklin Roosevelt and the Fireside Chats.” *Journal of Communication* 49(4):80–103.
- Shah, Simmone. 2024. “The Top 5 Campaign Moments on TikTok.” Time. Section: Politics.
URL: <https://time.com/7160873/trump-harris-2024-campaign-tiktok-viral/>
- Smith, David. 2020. “AOC’s cooking live streams perfect the recipe for making politics palatable.” *The Guardian* .
- Starita, Giovanni Daniele and Tommaso Trillò. 2022. “‘Happy Monday friends! Coffee?’ Matteo Salvini, good morning selfies, and the influencer politician.” *Contemporary Italian Politics* 14(3):331–351.
- Tamasco, Marc. 2025. “CNN anchor slams Democrats’ awkward TikTok video ‘inviting ridicule’.”
URL: <https://www.foxnews.com/media/cnn-anchor-slams-democrats-awkward-tiktok-video-inviting-ridicule>

TikTok. 2026. “How TikTok recommends videos #ForYou.”.

URL: <https://newsroom.tiktok.com/how-tiktok-recommends-videos-for-you?lang=en>

Todorov, Alexander, Anesu N. Mandisodza, Amir Goren and Crystal C. Hall. 2005. “Inferences of Competence from Faces Predict Election Outcomes.” *Science* 308(5728):1623–1626.

Tomasik, Emily and Katerina Eva Matsa. 2025. “1 in 5 Americans now regularly get news on TikTok, up sharply from 2020.”.

URL: <https://www.pewresearch.org/short-reads/2025/09/25/1-in-5-americans-now-regularly-get-news-on-tiktok-up-sharply-from-2020/>

Towner, Terri L. and Caroline L. Muñoz. 2024. “Tell Me an Instagram Story: Ephemeral Communication and the 2018 Gubernatorial Elections.” *Social Science Computer Review* 42(6):1288–1312.

Tsfati, Yariv, Jonathan Cohen, Shira Dvir-Gvirsman, Keren Tsurriel, Israel Waismelman and R. Lance Holbert. 2022. “Political Para-Social Relationship as a Predictor of Voting Preferences in the Israeli 2019 Elections.” *Communication Research* 49(8):1118–1147.

Tukachinsky, Riva, Nathan Walter and Camille J Saucier. 2020. “Antecedents and Effects of Parasocial Relationships: A Meta-Analysis.” *Journal of Communication* 70(6):868–894.

Vale, Corey. 2026. “Where Is the GOP’s Digital Response?”.

URL: <https://www.yahoo.com/news/articles/where-gops-digital-response-123942129.html>

Vega, Priscella. 2022. “Cool or cringe? Politicians try to connect on TikTok, but risk the dreaded ‘teenage eye roll’.”.

URL: <https://www.yahoo.com/news/politicians-want-reach-young-voters-120017955.html>

- Wells, Dylan. 2025. “Inside the Democrats’ new media strategy to reach voters and take on Trump.” *The Washington Post* .
- URL:** <https://www.washingtonpost.com/politics/2025/03/15/democrats-influencers-congress-media-strategy/>
- Wells, Dylan and Tatum Hunter. 2025. “How Zohran Mamdani used social media to build a movement.” *The Washington Post* .
- URL:** <https://www.washingtonpost.com/politics/2025/06/26/zohran-mamdani-social-media-memes/>
- Xiong, Heyu. 2021. “The Political Premium of Television Celebrity.” *American Economic Journal: Applied Economics* 13(4):1–33.
- Young, Dannagal Goldthwaite. 2020. *Irony and Outrage: The Polarized Landscape of Rage, Fear, and Laughter in the United States*. New York, NY: Oxford University Press.
- Yuan, Shupeí and Chen Lou. 2020. “How Social Media Influencers Foster Relationships with Followers: The Roles of Source Credibility and Fairness in Parasocial Relationship and Product Interest.” *Journal of Interactive Advertising* 20(2):133–147.
- Zulli, Diana and Terri L. Towner. 2021. “The Effects of “Live,” Authentic, and Emotional Instagram Images on Congressional Candidate Evaluations.” *Social Media + Society* 7(4):20563051211062917.

Appendix for “Should Candidates Campaign Like Influencers? Evidence From 16,572 Instagram Reels and Two Generative AI Video Experiments”

Cecilia Y. Sui

A	Data Overview	1
A.1	Candidate Universe and Platform Selection	1
A.2	Construction of the Instagram Reels Archive	1
A.2.1	Distribution of Reels Across Accounts	3
A.2.2	Temporal Coverage	4
A.3	Candidate Characteristics	4
A.4	Engagement-Data Sample	4
A.4.1	Candidate Characteristics and Engagement-Data Availability	6
A.5	Distribution of Engagement Outcomes	6
A.6	Scope and Limitations of the Study 1 Data	8
B	Measuring Influencer-Style Communication	10
B.1	From Concepts to Observable Indicators	11
B.1.1	Aesthetic Proximity	11
B.1.2	Relational Signaling	11
B.2	Full Coding Protocol	12
B.2.1	Dimension 1: Informal Proximity Cues in Production	12
B.2.2	Dimension 2: Informal Proximity Cues in Candidate Presentation	13
B.2.3	Dimension 3: Direct Address to Viewer	14
B.2.4	Dimension 4: Audience acknowledgment	15
B.2.5	Dimension 5: Self-Disclosure	17
B.2.6	Dimension 6: Interaction Invitation	18
B.2.7	Dimension 7: Identity Bridging	19
B.3	Automated Video Annotation	21
B.4	Human Validation	21
B.5	Index Construction	21
B.6	Distribution of the Measures	22
B.7	Alternative Thresholds and Measurement Robustness	23
B.8	Summary	24
C	Generative Stimulus Construction	25
C.1	Experimental Stimulus Design and Identification Strategy	25
C.1.1	Features Manipulated	26
C.1.2	Features Held Constant	27
C.1.3	Relationship to Study 1 Measurement	29
C.2	Script Construction	31
C.2.1	General Principles	31
C.2.2	Study 2 Scripts	31
C.2.3	Study 3 Scripts	32
C.3	Synthetic Voice Construction	33
C.3.1	Voice Design	33
C.3.2	Text-to-Speech Generation	34
C.3.3	Generation and Selection	34

C.4	Candidate Avatar and Still-Image Construction	35
C.4.1	Study 2 Candidate Images	35
C.4.2	Study 3 Candidate Images	37
C.4.3	Image Selection, Reproducibility, and Reverse-Image Search	41
C.5	Image-to-Video Animation	41
C.5.1	Models and Generation Settings	41
C.5.2	Motion Prompts and Treatment Invariance	42
C.5.3	Sequential Segment Generation	43
C.5.4	Generation Screening and Stopping Rule.	44
C.6	Lip Synchronization and Final Video Assembly	45
C.6.1	Lip-Synchronization Model Selection	45
C.6.2	Segment-Level Lip Synchronization	45
C.6.3	Final Video Assembly	46
C.7	Ethics and Safeguards	46
D	Validation and Realism	48
D.1	Validation Strategy	48
D.2	Study 2: Pre-Fielding Stimulus Screening	49
D.2.1	Candidate Videos Entering Screening	49
D.2.2	Pre-Fielding Exclusion of the Male Study 3 Stimuli.	49
D.3	Study 3: Realism Pilot Design.	50
D.3.1	Sample and Fielding	50
D.3.2	Experimental Procedure.	50
D.4	Perceived Realism Measure.	51
D.5	Benchmark Results	51
D.5.1	Pairwise Benchmark Comparisons	52
D.5.2	Stimulus-Level Variation	52
D.6	Realism Across Experimental Conditions.	53
D.6.1	Aesthetic-Proximity Contrast	53
D.6.2	Relational-Signaling Contrast	54
D.6.3	Realism Across the Four Factorial Conditions	54
E	Survey Instruments, Measures, and Preregistration.	56
E.1	Core Outcomes and Scale Construction.	56
E.2	Manipulation Checks.	56
E.2.1	Study 2.	56
E.2.2	Study 3.	58
E.3	Preregistrations and Deviations.	58
E.4	Study 2 Survey Instrument.	58
E.4.1	Pre-Treatment Questions	59
E.4.2	Video Instructions and Presentation.	60
E.4.3	Post-Video Candidate Evaluations	60
E.4.4	Pre-Debriefing Questions	61
E.4.5	Debriefing and Synthetic-Media Information	62
E.5	Study 3 Survey Instrument.	63
E.5.1	Consent Information and Introduction	63
E.5.2	Pre-Treatment Questions	63
E.5.3	Audiovisual Attention Check.	65
E.5.4	Experimental Video Instructions and Assignment	66
E.5.5	Post-Video Candidate Evaluations	66
E.5.6	Parasocial Closeness	67

E.5.7	Manipulation Checks and Additional Perceptions	67
E.5.8	Pre-Debriefing Questions	68
E.5.9	Debriefing and Synthetic-Media Information	69
F	Study 1: Full Results and Robustness	70
F.1	Full Regression Results for Platform Visibility and Engagement	70
F.2	Component-Level Associations	71
G	Study 2: Full Results and Robustness	73
G.1	Full Regression Results	73
G.2	Component-Level Results	74
H	Study 3: Full Results and Robustness	75
H.1	Full Study 3 Regression Results and Component-Level Results	75
H.2	Average Marginal Effects	80
H.3	Estimated Condition Means and Experimental Contrasts	82

A Data Overview

This appendix provides additional information on the data used in Study 1, which examines how state legislative candidates use influencer-style communication in their short-form videos and how those practices are associated with platform reach and engagement. The study distinguishes two dimensions of influencer-style communication. *Aesthetic proximity* captures informal and platform-native visual presentation, while *relational signaling* captures communicative practices that simulate interpersonal exchange with viewers. Appendix B describes the measurement and validation of these dimensions in detail.

Study 1 draws on a near-complete archive of publicly available Instagram Reels from candidates who advanced to the November 2025 general elections for the New Jersey General Assembly and Virginia House of Delegates. The full corpus contains 16,572 Reels from 200 Instagram accounts. A subset of 10,723 Reels from 160 individually attributable candidate accounts could be matched to credible post-level engagement data and is used to examine associations with views, likes, comments, and reposts.

A.1 Candidate Universe and Platform Selection

The candidate universe consists of everyone who advanced to the November 2025 general elections for the New Jersey General Assembly and Virginia House of Delegates, the only two state legislative elections held in the United States that year. The initial roster contains 159 candidates in New Jersey and 196 in Virginia, for a total of 355 candidates. Two Virginia candidates later withdrew from the general election: Republican C. Todd Gilbert and Democrat Sharron A. Burgess. These two candidates are also removed from the all platform-coverage denominators and do not contribute content to the final Reel archive.

An undergraduate research assistant searched for each candidate on both Instagram and TikTok. Candidate accounts were verified against campaign websites, party listings, candidate websites, and other official sources before inclusion. Accounts that could not be confidently attributed to a candidate or campaign were excluded.

Instagram was substantially more widely adopted than TikTok in both states. As Table A1 shows, 111 of the 159 New Jersey candidates (69.8%) had a public Instagram presence, compared with 28 (17.6%) on TikTok. In Virginia, 131 of 194 candidates (67.5%) had a public Instagram presence, compared with 23 (11.9%) on TikTok. Across both states, 242 candidates (68.6%) were identified on Instagram and 51 (14.4%) on TikTok. Because candidates could maintain accounts on both platforms, these categories are not mutually exclusive.

Instagram therefore provides substantially broader coverage of state legislative candidates in these elections while retaining the short-form vertical-video format at the center of the study. For that reason, Study 1 focuses on Instagram Reels.

A.2 Construction of the Instagram Reels Archive

Not every verified Instagram account contained publicly available video content. Of the 111 Instagram presences identified in New Jersey, 99 accounts had publicly available Reels as of February 22, 2026. Of the 131 identified in Virginia, 101 accounts had publicly

Table A1: Social Media Platform Coverage Among 2025 State Legislative Candidates

	New Jersey (<i>N</i> = 159)	Virginia (<i>N</i> = 194)	Total (<i>N</i> = 353)
Public Instagram presence	111 (69.8%)	131 (67.5%)	242 (68.6%)
Public TikTok presence	28 (17.6%)	23 (11.9%)	51 (14.4%)

Note: Entries report the number and percentage of candidates for whom a verified public account was identified on each platform. Instagram and TikTok presence are not mutually exclusive. The candidate universe consists of candidates who advanced to the November 2025 general elections in New Jersey and Virginia. Two Virginia candidates who subsequently withdrew from the general election were excluded from the candidate universe.

available Reels. The remaining accounts posted only image-based content or contained no publicly visible Reels at the time of collection.

For each account with Reel content, I collected all Reels publicly available as of February 22, 2026 using the open-source Python package **Instaloader**. “Complete available history” therefore refers to all Reel content publicly visible when the account was collected. Videos deleted or made private before collection could not be recovered.

The resulting corpus contains 16,572 Reels from 200 Instagram accounts: 99 accounts from New Jersey and 101 from Virginia. Table A2 summarizes the transition from verified Instagram presence to the final video archive.

Table A2: Construction of the Instagram Reels Archive

	New Jersey	Virginia	Total
Verified public Instagram accounts	111	131	242
Accounts with publicly available Reels	99 (89.2%)	101 (77.1%)	200 (82.6%)
Accounts without publicly available Reels	12	30	42
Number of Reels collected	7,088	9,484	16,572

Note: Account counts refer to distinct Instagram accounts rather than necessarily distinct candidates. One New Jersey account is jointly maintained by multiple officeholders, as described in the text. The percentage of accounts with Reels is calculated among verified public Instagram accounts.

One account requires special treatment. The Instagram page *njdistrict38* is the official account for New Jersey’s 38th Legislative District and is maintained jointly by Senator Joe Lagana, Assemblywoman Lisa Swain, and Assemblyman Chris Tully. It therefore carries district-level communication associated with two Assembly candidates in the candidate universe rather than belonging uniquely to one candidate. I retain the account in the video-level corpus because its Reels are official campaign communication involving candidates in the sample. However, I exclude it from descriptive analyses that require assigning a unique candidate-level characteristic to an account. The candidate-level descriptive sample therefore contains 199 individually attributable candidate accounts. The *njdistrict38* was not included in the 10,723-Reel engagement sample since no engagement

metadata was successfully matched to this account.

A.2.1 Distribution of Reels Across Accounts

Posting activity varies substantially across candidate accounts. The 200 accounts contain an average of 82.9 Reels, but the median is only 31, reflecting a heavily right-skewed distribution. One-quarter of accounts contain 12 or fewer Reels, while three-quarters contain 78.5 or fewer. The most active account contains 1,270 Reels. Table A3 reports the full distribution. Figure A1 shows the distribution of Instagram Reels by candidate account.

Table A3: Distribution of Instagram Reels Across Candidate Accounts

Statistic	Reels per Account
Number of accounts	200
Total Reels	16,572
Mean	82.86
Std. Dev.	158.99
Minimum	1
25th percentile	12
Median	31
75th percentile	78.5
IQR	66.5
Maximum	1,270

Note: Statistics are calculated across the 200 Instagram accounts containing at least one publicly available Reel.




Figure A1: Distribution of Instagram Reels Across Candidate Accounts

Note: The figure shows the distribution of the number of Reels collected per Instagram account across the 200 accounts with at least one publicly available Reel. The left panel displays the distribution on a linear scale, while the right panel displays the same distribution with the horizontal axis on a logarithmic scale. Dashed lines indicate the mean (82.9 Reels) and median (31 Reels) per account.

A.2.2 Temporal Coverage

The archive captures how candidates who ran in the 2025 elections used Instagram Reels across their publicly available account histories rather than restricting the sample to content posted during the 2025 campaign. The earliest Reel in the archive was posted on November 22, 2019. This distinction matters for interpretation. Study 1 characterizes the short-form communication of candidates who participated in the 2025 state legislative elections; it does not assume that every Reel in the archive was produced specifically for the 2025 campaign. In ongoing work, I include additional data on posting date, days relative to the general election, and video duration.

A.3 Candidate Characteristics

Table A4 describes the 199 individually attributable candidate accounts in the video archive. The shared *njdistrict38* account is omitted because candidate-level characteristics cannot be assigned uniquely to that account.

The account sample contains 98 individually attributable candidate accounts from New Jersey and 101 from Virginia. Across the two states, 139 candidates (69.8%) are Democrats and 55 (27.6%) are Republicans. The remaining five candidates represent the Independent, Forward, Green, and We the People affiliations. Incumbents make up 55.3% of the sample, and 62.3% of candidates represented in the archive won their 2025 election.

The sample also contains substantial demographic variation. Because official sources did not consistently report candidate gender, gender is coded as *perceived gender* using publicly available names, pronouns, and photographs. This variable should not be interpreted as candidate self-identification.

A.4 Engagement-Data Sample

The descriptive analysis of candidate communication uses the complete 16,572-Reel archive. The analysis of platform reach and engagement requires an additional merge to post-level engagement metadata.

Engagement metrics were collected over five consecutive days, February 23–27, 2026, using the Meta Content Library API and Instaloader. For each successfully matched Reel, I record four platform-reported outcomes: views, likes, comments, and reposts.

Usable engagement metadata were obtained for 10,723 Reels, or 64.7% of the full video archive, representing 160 individually attributable candidate accounts. At the candidate-account level, 39 of the 199 individually attributable accounts could not be matched to sufficiently reliable engagement data and are excluded from the engagement analysis. Potential reasons for failed matches include account changes, deleted or private content, and unavailable historical metrics. The available data do not always permit the precise cause of a failed match to be identified.

Because the engagement measures are cumulative snapshots, the date on which an account was collected determines the endpoint of its exposure window. The five-day collection period creates a maximum difference of four days in accumulation time across accounts.

Table A4: Candidate Characteristics in the Instagram Reels Archive

	New Jersey (<i>N</i> = 98)	Virginia (<i>N</i> = 101)	Total (<i>N</i> = 199)
<i>Party</i>			
Democratic	66 (67.3%)	73 (72.3%)	139 (69.8%)
Republican	30 (30.6%)	25 (24.8%)	55 (27.6%)
Independent	0 (0.0%)	2 (2.0%)	2 (1.0%)
Forward	0 (0.0%)	1 (1.0%)	1 (0.5%)
Green	1 (1.0%)	0 (0.0%)	1 (0.5%)
We the People	1 (1.0%)	0 (0.0%)	1 (0.5%)
<i>Incumbency</i>			
Incumbent	54 (55.1%)	56 (55.4%)	110 (55.3%)
Challenger	44 (44.9%)	45 (44.6%)	89 (44.7%)
<i>Election Outcome</i>			
Won	63 (64.3%)	61 (60.4%)	124 (62.3%)
Lost	35 (35.7%)	40 (39.6%)	75 (37.7%)
<i>Race</i>			
White	68 (69.4%)	54 (53.5%)	122 (61.3%)
Black	12 (12.2%)	34 (33.7%)	46 (23.1%)
Hispanic	13 (13.3%)	5 (5.0%)	18 (9.0%)
Asian	5 (5.1%)	8 (7.9%)	13 (6.5%)
<i>Perceived Gender</i>			
Female	41 (41.8%)	53 (52.5%)	94 (47.2%)
Male	57 (58.2%)	48 (47.5%)	105 (52.8%)

Note: Entries report counts with column percentages in parentheses. The table includes the 199 Instagram accounts attributable to individual candidates and excludes the shared *njdistrict38* account. Perceived gender is researcher-coded from publicly available information and does not necessarily correspond to candidates' self-identified gender.

A.4.1 Candidate Characteristics and Engagement-Data Availability

To assess whether the engagement sample systematically excludes particular types of candidates, I compare accounts with and without usable engagement data on candidate characteristics and communication style. Table A5 shows that there are no statistically detectable differences between the two groups in party, perceived gender, race, incumbency, or election outcome. Missingness does, however, differ substantially by state. Thirty-two of the 39 excluded accounts are from Virginia, whereas only seven are from New Jersey. Thus, 82.1% of excluded accounts are from Virginia compared with 43.1% of accounts retained in the engagement sample ($p < .001$). This imbalance does not affect the full-corpus analysis of how candidates use influencer-style cues, which uses all 16,572 Reels, but it limits the representativeness of the engagement subsample.

Because candidate characteristics alone cannot establish whether retained accounts communicate differently from excluded accounts, I conduct an additional comparison using the influencer-style measures themselves. For each candidate account, I calculate the mean aesthetic-proximity index, mean relational-signaling index, and proportion of Reels containing each component cue. Table A6 compares these account-level averages between the 160 accounts with engagement data and the 40 accounts without them.

None of the account-level differences is statistically significant. Accounts retained in the engagement sample therefore do not show statistically detectable differences in their average use of aesthetic proximity, relational signaling, or the individual influencer-style cues relative to accounts without usable engagement data.

The results provide some reassurance that there are no statistically detectable differences between the two groups in their average communication style, although the relatively small number of excluded accounts limits the precision of these comparisons. The engagement analysis should therefore still be interpreted as applying to the 160 candidate accounts for which credible platform metrics could be recovered rather than to the complete 200-account archive.

The substantial state imbalance remains an important limitation. In ongoing data collection, I examine why engagement metadata were disproportionately unavailable for Virginia accounts and attempt to recover additional observations where possible. I also examine posting date and account history once those variables are finalized, since the availability of historical engagement metadata may vary with the age of an account or its posts.

A.5 Distribution of Engagement Outcomes

Table A7 reports the untransformed distributions of the four engagement outcomes among the 10,723 Reels with usable metadata.

All four measures are highly right-skewed. The median video received 806 views and 31 likes, compared with means of 11,684 views and 427 likes, reflecting a small number of extremely high-engagement posts. The maximum observed video received more than 8.4 million views.

Zero values are also common for the more active forms of engagement. Likes are zero for 2.6% of Reels, comments for 36.2%, and reposts for 84.1%. Views are strictly positive in the matched data. For the regression analyses, I therefore transform each outcome

Table A5: Comparison of Candidate Accounts With and Without Engagement Data

Characteristic	Without Engagement Data (<i>N</i> = 39)	With Engagement Data (<i>N</i> = 160)	<i>p</i>
<i>State</i>			
New Jersey	7 (17.9%)	91 (56.9%)	< .001
Virginia	32 (82.1%)	69 (43.1%)	
<i>Party</i>			
Democratic	28 (71.8%)	111 (69.4%)	0.849
Forward	0 (0.0%)	1 (0.6%)	
Green	0 (0.0%)	1 (0.6%)	
Independent	1 (2.6%)	1 (0.6%)	
Republican	10 (25.6%)	45 (28.1%)	
We the People	0 (0.0%)	1 (0.6%)	
<i>Perceived Gender</i>			
Female	15 (38.5%)	79 (49.4%)	0.296
Male	24 (61.5%)	81 (50.6%)	
<i>Race</i>			
Asian	4 (10.3%)	9 (5.6%)	0.309
Black	12 (30.8%)	34 (21.2%)	
Hispanic	4 (10.3%)	14 (8.8%)	
White	19 (48.7%)	103 (64.4%)	
<i>Incumbency</i>			
Challenger	20 (51.3%)	69 (43.1%)	0.460
Incumbent	19 (48.7%)	91 (56.9%)	
<i>Election Outcome</i>			
Lost	19 (48.7%)	56 (35.0%)	0.161
Won	20 (51.3%)	104 (65.0%)	

Note: Entries report counts with column percentages in parentheses. Reported *p*-values are from chi-square tests of independence comparing the categorical distribution of each characteristic across accounts with and without usable engagement data. The state comparison is calculated from the account counts shown in the table. Several minor-party categories contain very small cell counts.

Table A6: Account-Level Comparison of Influencer-Style Communication With and Without Engagement Data

Measure	With Engagement ($N = 160$)	Without Engagement ($N = 40$)	t	p
Aesthetic proximity index	0.409 (0.270)	0.460 (0.253)	-1.118	0.268
Relational signaling index	0.158 (0.105)	0.154 (0.111)	0.172	0.864
Informal production	0.416 (0.280)	0.455 (0.275)	-0.806	0.423
Informal candidate presentation	0.402 (0.271)	0.464 (0.263)	-1.327	0.189
Direct address	0.344 (0.245)	0.300 (0.230)	1.082	0.284
Mutual awareness	0.063 (0.106)	0.044 (0.057)	1.500	0.136
Self-disclosure	0.204 (0.160)	0.260 (0.225)	-1.480	0.145
Invitation to interact	0.082 (0.129)	0.062 (0.100)	1.050	0.297
Identity bridging	0.096 (0.102)	0.106 (0.166)	-0.367	0.716

Note: Entries report account-level means with standard deviations in parentheses for accounts with and without usable engagement metadata. Video-level measures are aggregated to the account level before comparison, such that component-measure means represent the average share of an account’s Reels containing each feature. The final two columns report two-sample t -tests comparing account-level means across the two groups.

as $(\log(y + 1))$, which reduces the influence of the extreme upper tail while retaining observations with zero engagement.

A.6 Scope and Limitations of the Study 1 Data

Three features of the data are important for interpreting Study 1. First, the archive is an account-history archive rather than a campaign-period sample. It contains all Reels publicly available from the accounts of candidates who ran in the 2025 elections, including content posted before the 2025 campaign. The observational results therefore describe how this population of candidates has used short-form video across their available Instagram histories rather than estimating communication specifically during the 2025 general-election campaign.

Second, the archive captures publicly available content at the time of collection. Reels deleted or made private before February 22, 2026 cannot be observed. The corpus should therefore be understood as a near-complete archive of *available* Reel content rather than a historical record of every video ever posted by these accounts.

Third, the engagement analysis uses a smaller subset of the archive. Engagement metadata were available for 10,723 of 16,572 Reels and 160 of the 199 individually attributable candidate accounts. Candidate characteristics and average communication styles are broadly similar across accounts with and without engagement data, although missingness is disproportionately concentrated in Virginia. The engagement results should therefore be interpreted as within-candidate associations among accounts for which post-level metrics could be recovered rather than as estimates from the entire 200-account archive.

Finally, platform engagement is not equivalent to voter evaluation. Views, likes, comments, and reposts capture observable responses within Instagram, but they cannot

Table A7: Descriptive Statistics for Engagement Outcomes

	Views	Likes	Comments	Reposts
Observations	10,723	10,723	10,723	10,723
Mean	11,684.13	427.46	47.10	3.18
Std. Dev.	131,235.60	3,214.89	344.57	68.65
Minimum	1	0	0	0
25th percentile	291	11	0	0
Median	806	31	1	0
75th percentile	2,627.5	99	8	0
Maximum	8,422,513	124,849	13,210	4,432
Zero count	0 (0.0%)	279 (2.6%)	3,887 (36.2%)	9,019 (84.1%)

Note: Descriptive statistics are calculated across the 10,723 candidate-video observations with usable engagement metadata. The final row reports the number and percentage of observations with zero engagement for each outcome. The main analyses transform each engagement outcome as $\log(y + 1)$.

establish whether a video made viewers more favorable toward the candidate or more willing to support them electorally. Study 1 consequently addresses candidate adoption and platform feedback descriptively. The experimental studies separately examine the causal effects of aesthetic proximity and relational signaling on voter evaluations.

B Measuring Influencer-Style Communication

This appendix describes the measurement of influencer-style communication used in Study 1. The study analyzes 16,572 publicly available Instagram Reels from 200 accounts associated with candidates who advanced to the 2025 general elections for the New Jersey General Assembly and Virginia House of Delegates. Appendix A describes the candidate universe, platform selection, data collection, and sample in detail.

Influencer-style communication is not treated as a single visual or rhetorical feature. Instead, I distinguish two analytically separate dimensions: *aesthetic proximity* and *relational signaling*. Aesthetic proximity captures visual and embodied features that reduce the apparent distance between a candidate and an ordinary social media user. Relational signaling captures communicative practices that frame the audience as participants in an interpersonal or quasi-interpersonal exchange rather than simply as recipients of a broadcast message.

I operationalize these dimensions using seven binary indicators coded independently at the video level. Two indicators measure aesthetic proximity: informal production and informal candidate presentation. Five measure relational signaling: direct address, audience acknowledgment, self-disclosure, interaction invitation, and identity bridging. A video can contain any combination of these cues; the categories are not mutually exclusive. Table A8 provides a high-level overview of the seven indicators.

Table A8: Overview of Influencer-Style Communication Indicators

Dimension	Indicator	Core Coding Question
<i>Aesthetic proximity</i>	Informal production	Does the filming style or setting resemble casual, native short-form content rather than formal campaign or broadcast production?
	Informal candidate presentation	Does the candidate’s attire, posture, movement, or delivery present them in a casual or off-stage manner?
<i>Relational signaling</i>	Direct address	Does the candidate treat the viewer as the person being spoken to?
	Audience acknowledgment	Does the candidate acknowledge or simulate awareness of viewers’ questions, reactions, or presence?
	Self-disclosure	Does the viewer learn something about the candidate’s personal life, background, everyday activities, or internal emotional experience?
	Interaction invitation	Does the candidate explicitly invite the viewer to respond, communicate, or participate?
	Identity bridging	Does the candidate explicitly place themselves and the viewer within the same social group or community?

Note: All indicators are coded independently at the video level as present (1) or absent (0). A video may contain multiple indicators.

The remainder of this appendix explains how the indicators were defined and coded, how automated annotations were validated against human judgments, how the two sum-

mary indices were constructed, and whether the descriptive conclusions depend on alternative ways of combining or thresholding the indicators.

B.1 From Concepts to Observable Indicators

The measurement strategy follows the conceptual distinction between how a candidate *looks* and how a candidate *communicates a relationship* with the audience.

B.1.1 Aesthetic Proximity

Aesthetic proximity refers to visual and embodied cues that make a candidate’s presentation resemble ordinary platform-native content rather than conventional campaign or broadcast communication. Research on politicians’ visual self-presentation distinguishes professional and institutional presentation from private, personal, and informal forms of communication (Farkas and Bene, 2021; Metz, Kruikemeier and Lecheler, 2020; Peng, 2021). I separate this dimension into two indicators because the informality of the production and the informality of the candidate’s own presentation need not coincide.

Informal production captures filming style and setting, including features such as selfie or handheld framing, everyday environments, and low institutional polish. *Informal candidate presentation* captures the candidate’s embodied appearance and performance, including casual attire, relaxed posture or movement, and informal delivery. A candidate can therefore score positively on one indicator without scoring positively on the other. For example, a formally dressed candidate may deliver a speech into a handheld phone, while a casually dressed candidate may appear in footage recorded with professional camera work.

These two indicators jointly capture aesthetic proximity without requiring a particular location, item of clothing, or production technique to define the construct on its own.

B.1.2 Relational Signaling

Relational signaling refers to communicative cues that simulate features of interpersonal exchange between the candidate and the viewer. The indicators draw on parasocial research emphasizing direct acknowledgment, self-disclosure, perceived interaction, and other practices through which media figures can create the experience of being personally addressed (Hartmann and Goldhoorn, 2011; Horton and Wohl, 1956; Leite and Baptista, 2022).

I distinguish five observable forms of relational signaling. *Direct address* treats the viewer as the conversational addressee. *Audience acknowledgment* signals awareness of viewers’ questions, reactions, or presence. *Self-disclosure* reveals personal experiences, identity, everyday activities, or emotions beyond the candidate’s official political role. *Interaction invitation* explicitly asks viewers to respond, communicate, or participate. Finally, *identity bridging* places the candidate and viewer within a shared social identity or community.

These indicators capture related but distinct practices. For example, a candidate can directly address viewers without inviting them to respond; acknowledge previous audience questions without disclosing anything personal; or describe their own background

without explicitly placing themselves in the same group as the audience. Each indicator is therefore coded separately.

B.2 Full Coding Protocol

The unit of annotation is one complete video. Coders evaluate the audiovisual content of the Reel, including spoken language, candidate-intended on-screen text, and visible visual cues. Platform-generated captions or surrounding interface elements are not treated as part of the candidate’s communication unless they reproduce content intentionally embedded in the video.

Each of the seven indicators is coded independently as present (1) or absent (0). A video may therefore receive multiple positive labels. The coding protocol is deliberately conservative: when the available evidence is ambiguous, the indicator is coded as absent. This rule limits the extent to which weak or ambiguous cues are treated as evidence of influencer-style communication.

The complete codebook follows. For every indicator, it provides a conceptual definition, positive coding criteria, negative examples, and edge-case rules intended to distinguish adjacent constructs. Below is the exact codebook used by human annotators.

General Coding Instructions:

1. Unit: one video (one TikTok/Instagram Reel/YouTube Shorts)
2. Code each dimension as 1 = present, 0 = absent.
3. Code based on the video content (spoken words + on-screen text + visible visual cues), not captions added by the platform unless they clearly reflect candidate-intended text (e.g., edited subtitles in-video)
4. Each dimension captures a different communication cue and should be coded independently.
5. A video may receive multiple 1’s if multiple cues appear.
6. When uncertain, default to 0 and leave a short note in the “Notes” column.

B.2.1 Dimension 1: Informal Proximity Cues in Production

Variable name: informal_production (0/1)

Definition: The video uses filming style, camera work, or setting that reduces institutional distance and visually resembles native short-form social media rather than formal campaign or broadcast production.

Code 1 when the overall filming style or setting resembles casual, native short-form social media content rather than formal campaign production.

The following cues often signal this style:

1. Selfie or arm’s length framing
2. Handheld / phone-native framing or candid camera angle

3. Private/domestic/personal setting: e.g., kitchen, living room, car, backyard, hallway at home
4. Everyday, non-staged environment is visibly part of the shot
5. Overall production style clearly signals low institutional polish and resembles casual user-generated social media content rather than professional campaign or broadcast production

Code 0 when the overall production style resembles formal campaign communication or broadcast-style production.

Examples include:

1. Podium/ institutional/ legislative/ press-conference style
2. Studio or campaign-ad style with stable professional framing and staged backdrop
3. Professionally edited broadcast-style communication, even if filmed in a home or casual setting
4. Scenic or B-roll phone footage without candidate presence or candidate speech

Edge-case rules:

1. Minor shakiness alone does not qualify for 1.
2. If informal and formal cues conflict, code based on the dominant overall visual impression.
3. If uncertain, default to 0.

B.2.2 Dimension 2: Informal Proximity Cues in Candidate Presentation

Variable name: informal_candidate (0/1)

Definition: The candidate's own embodied presentation signals casualness or low formality through attire, posture, movement, or delivery style.

Code 1 if any of the following is clearly present:

1. Casual or everyday attire: e.g., t-shirt, hoodie, sweatshirt, polo, casual sweater, baseball cap, athleisure
2. Relaxed or non-official posture: e.g., leaning on counter, sitting casually, walking while speaking, perched informally
3. Casual delivery style: low-script, or spontaneous-seeming delivery rather than formal speech performance

Code 0 if the candidate is primarily presented as:

1. Formally dressed for office-seeking communication: e.g., suit, blazer, formal wear
2. Standing or seated in a conventionally official speech posture

3. Delivering a clearly staged campaign speech, even if recorded on a phone or in a casual location

Edge-case rules:

1. Code 1 when the candidate appears visibly relaxed or off-stage, even if the clothing itself is technically business attire: e.g., top button undone, sleeves rolled, shirt untucked, loosened tie
2. Conversational refers to the delivery style, not relational content. Direct address, self-disclosure, or identity bridging need not to be present for conversational delivery style.
3. If uncertain, default to 0.

B.2.3 Dimension 3: Direct Address to Viewer

Variable name: `direct_address` (0/1)

Definition: The candidate directly addresses the viewer as the conversational addressee, using second-person language or equivalent forms that clearly treat the audience as the person being spoken to. This cue signals a speaker-to-viewer interaction frame, rather than a broadcast-style statement.

Code 1 if any of the following is clearly present:

1. Clear second-person pronouns directed toward the viewer:
e.g., “You deserve better leadership.”
“Your family should feel safe.”
“I’m asking you to vote.”
“Thank you for watching.”
2. Second-person plural or audience collective used as addressee:
e.g., “Folks, we have a problem.”
“Friends, I need your help.”
“Neighbors, we can fix this.”
“Everyone, listen to this.”

These terms, must function as the addressee, not just as a noun. For example, “everyone wants lower taxes” would be coded as 0.

3. Inclusive “we/us” that clearly includes the viewer.
e.g., “We Jersey families deserve better.”
“Together we can fix this.”
“We all know this system is broken.”

Code 0 if any of the following apply:

1. No addressee framing. Purely informational or broadcast speech.
e.g., “This bill will lower taxes.”
“Atlantic County needs better leadership.”
2. “We” refers to an institutional actor.
e.g., “We in the legislature passed this bill.”
“Our campaign is working hard.”
“We Republicans/Democrats believe...”
3. Audience groups mentioned but not addressed.
e.g., “Families in New Jersey are struggling.”
“Working people deserve better.”

Edge-case rules:

1. If it is unclear whether the viewer is included, code as 0.
2. Calls to action without second-person are coded as 0.
e.g., “Vote this November.”
“Watch this.”
3. Generic greeting, code as 0. e.g., “Hi everyone!”

B.2.4 Dimension 4: Audience acknowledgment

Variable name: `audience_acknowledgment` (0/1)

Definition: The candidate signals awareness of the audience’s presence, reactions, or questions, framing the message as part of an ongoing or imagined interaction with viewers. This cue simulates two-way awareness, even though the interaction is not actually occurring.

Code 1 if any of the following occur:

1. Candidate references audience feedback or reactions.
e.g., “I saw your comments about this issue.”
“A lot of you have been asking me about property taxes.”
“People keep messaging me about this.”
“Many of you asked what this bill actually does.”
2. Candidate anticipates viewer questions.
e.g., “You might be wondering why this matters.”
“You’re probably asking what this means for you.”
“You may be thinking this sounds expensive.”

3. Candidate frames the video as responding to audience curiosity, like a back-and-forth explanatory interaction.

e.g., “Let me explain what’s going on here.”

“Let me show you what this bill actually does.”

“Let me walk you through this.”

4. Candidate directs viewer attention in a way that assumes shared perception.

e.g., “Look at this.”

“Watch what happens here.”

“Take a look at this chart.”

Code 0 if any of the following apply:

1. Purely informational statements:

e.g., “This bill lowers taxes.”

“Here’s what the legislation does.”

2. Generic rhetorical traditions:

e.g., “Let me be clear.”

“Let me say this.”

These are speech markers, not audience awareness.

3. References to opponents or other political actors without audience awareness.

e.g., “My opponent keeps saying this policy works.”

However, if the script signals audience awareness, it is coded as 1. e.g., “Republicans want me to ignore you.”

4. Generic commands without interaction framing:

e.g., “Vote this November.”

Calls to actions without audience awareness, framed as part of an ongoing or imagined interaction with viewers are coded as 0.

Edge-case rules:

1. Does the candidate acknowledge the audience’s reactions or simulate responding to them? If yes, code as 1. If the candidate is simply delivering a statement, code as 0.

B.2.5 Dimension 5: Self-Disclosure

Variable name: self_disclosure (0/1)

Definition: The candidate reveals personal life experiences, identity information, everyday activities or personal emotional experiences that portray them as a private individual rather than solely as a political actor. The disclosed information is not strictly necessary to explain the policy argument and instead serves to humanize the candidate.

Code 1 if any of the following occur:

1. Personal life anecdotes outside official political roles.
e.g., “When I was in school, my family struggled to pay rent.”
“Growing up, I watched my parents work two jobs.”
“When I first became a parent, I realized how expensive childcare is.”
2. Family references.
e.g., “As a mother of five, I care deeply about education.”
“My daughter asked me yesterday why this issue matters.”
“I’m spending the evening with my kids after this meeting.”
3. Everyday personal activities that show ordinary daily-life behavior like commuting, cooking, grocery shopping, working out, walking the dog, morning routines.
e.g., “I’m on my way to pick up my kids from school.”
“Just finished my morning run.”
4. Personal background not tied to official office.
e.g., “I used to work in construction.”
“I grew up in a working-class neighborhood.”
“My parents are immigrants.”
“As a small business owner, I ran and operated business for 10 years.”
5. Personal emotional disclosure, where the candidate reveals their own emotional experiences or feelings about events or personal situations, especially when tied to personal life experiences.
e.g., “I was really scared when my son got sick.”
“It breaks my heart to see families struggling like this.”
“I’m proud of the community we’ve built.”

These statements reveal the candidate’s internal emotional state, rather than simply using emotions rhetoric, such as “Happy to announce that we passed this bill.”

Code 0 if any of the following apply:

1. Political credentials.
e.g., “I served on the city council for 10 years.”
“As your state senator, I passed this bill.”

2. Policy justification without personal revelation.
e.g., “I support education funding because it’s important to families.”
“This policy would help working parents.”
3. Emotional language without personal disclosure.
e.g., “This is outrageous.”
“The policy is devastating.”

Edge-case rules:

1. Does the viewer learn something about the candidate’s personal life? If yes, code as 1. If the viewer only learns about their political role or policy stance, code as 0.

B.2.6 Dimension 6: Interaction Invitation

Variable name: `interaction_invite` (0/1)

Definition: The candidate explicitly invites the viewer to respond, interact, or participate, signaling that audience participation is expected or welcome. These invitations imply two-way communication or participatory social media behavior, even if no response actually occurs.

Code 1 if any of the following occur:

1. Requests for viewer feedback.
e.g., “Comment down below.”
“Tell me what you think.”
“Let me know your thoughts.”
“What do you think about this?”
2. Direct communication invitations with the candidate or campaign.
e.g., “DM me.”
“Message me if you have questions.”
“Reach out to me.”
“Call our office.”
3. Platform-native participation behaviors.
e.g., “Follow for more updates.”
“Share this with your friends.”
“Stitch this/Duet this.”
4. Invitations to participate in a shared event or conversation (online or offline).
e.g., “Join us tonight.”
“Come talk with us at the town hall.”
“See you there!”

Code 0 if any of the following apply:

1. Pure electoral mobilization.
e.g., “Vote this November.”
“Vote him out.”
“Support this bill.”
2. Informational redirection.
e.g., “Learn more at my websites.”
“Visit my website for details.”
3. Collective political statements.
e.g., “Together we can win.”

Edge-case rules:

1. Does the candidate ask the viewer to respond, interact, or participate? If yes, code as 1. If the candidate only asks the viewer to support a political outcome, such as a bill, code as 0.
2. Audience acknowledgment signals “I know you are there and reacting.” and reciprocity signals “You should respond and engage.” They often co-occur, but neither logically contains the other.
 - Example of audience acknowledgment without reciprocity: “A lot of you have been asking me about property taxes, so let me explain.”
 - Example of reciprocity without audience acknowledgment: “Comment down below and tell me what you think.”
 - Example of both: “I’ve been getting a lot of questions about this - comment below if you have more.”

B.2.7 Dimension 7: Identity Bridging

Variable name: community (0/1)

Definition: The candidate constructs a shared social identity or community that explicitly includes both the candidate and the audience, framing the relationship as membership in the same group. This cue signals collective belonging between the candidate and the viewer, rather than simply describing or representing a group.

Code 1 if any of the following occur:

1. The candidate explicitly identifies themselves as part of the same social group as the audience.
e.g., “As moms like us, we know how hard childcare is.”
“Like many of you, I run a small business.”

2. The candidate frames a place as a shared community identity.
e.g., “We Jersey families deserve better schools.”
“Here in St. Louis, we take care of each other.”
“Our community here in King County.”
3. The candidate constructs a collective group identity with supporters, where the audience is framed as members of the same collective effort.
e.g., “Together we’re building something bigger.”

Code 0 if any of the following apply:

1. Generic “we” without identity construction.
e.g., “We need better leadership.”
“We must fix this system.”
2. Representation rhetoric where the candidates speaks for a group rather than with them.
e.g., “I represent the people of Virginia.”
“I fight for working families.”
3. Third-person group references where groups are described from outside, not shared membership.
e.g., “American workers deserve better.”
“Members of the community are struggling.”
“People in King County are facing challenges.”

Edge-case rules:

1. Identity statements without bridging.
e.g., “As a small business owner, I understand the challenges.”
This is self-disclosure, but not identity bridging.
2. Identity bridging with comparison.
e.g., “As a small business owner like many of you ...”
Now the candidate explicitly links themselves with the viewers.
3. Does the candidate explicitly place themselves and the viewer inside the same social group? If yes, code as 1. If the candidate only talks about a group or represents them from outside, code as 0.

B.3 Automated Video Annotation

Manually coding all 16,572 videos would require substantial annotation time and would make it difficult to apply a detailed seven-indicator codebook consistently across the full audiovisual corpus. I therefore use a multimodal large language model to scale the annotation after developing and validating the codebook manually.

I first applied the complete coding protocol to a random sample of 100 Reels by hand. I then supplied the same coding instructions to **Gemini-3.1-Pro**, which evaluated each video’s full audiovisual content and returned a binary classification for each of the seven indicators. The model received the full uploaded video in .mp4 format and the complete codebook in the prompt rather than a transcript or still frame alone, allowing the coding procedure to incorporate visual presentation, spoken delivery, and verbal content.

The model was instructed to apply the indicators independently and to follow the same conservative decision rule used in human coding: ambiguous cases should be coded as absent.

B.4 Human Validation

To assess whether the automated classifications recover the judgments generated by the codebook, I compare model labels against human labels on the 100-video random validation sample. The human labels were generated before the model classifications so that model outputs did not influence human coding decisions.

Across the seven indicators and 100 videos, the human and automated labels agree on 95.1% of classifications overall. Agreement is also high for each individual indicator: 88% for informal production, 98% for informal candidate presentation, 90% for direct address, 96% for audience acknowledgment, 90% for self-disclosure, 100% for interaction invitation, and 98% for identity bridging.

B.5 Index Construction

The main analyses summarize the seven indicators into two theoretically defined indices rather than combining all cues into a single influencer-style score. The *aesthetic-proximity index* is the normalized sum of informal production and informal candidate presentation:

$$A_i = \frac{\text{Informal Production}_i + \text{Informal Candidate Presentation}_i}{2}.$$

Because both components are binary, the index takes the values (0), (0.5), and (1). A value of 0 indicates that neither aesthetic cue is present, 0.5 that one is present, and 1 that both are present.

The *relational-signaling index* is the normalized sum of direct address, audience acknowledgment, self-disclosure, interaction invitation, and identity bridging:

$$R_i = \frac{1}{5}(\text{Direct Address}_i + \text{Audience acknowledgment}_i + \text{Self-Disclosure}_i \\ + \text{Interaction Invitation}_i + \text{Identity Bridging}_i).$$

It therefore ranges from 0 to 1 in increments of 0.2. A value of 0 indicates that none

of the five relational cues is present, while a value of 1 indicates that all five appear in the same video.

The normalization gives equal weight to each component within a dimension while placing both summary indices on the same 0–1 scale. I do not weight indicators by their empirical prevalence or estimated predictive performance because the indices are intended to represent theoretically defined communication bundles rather than optimized statistical scores.

I also do not combine the two dimensions in the main analyses. Doing so would obscure the distinction at the center of the paper between looking like an influencer and communicating relationally with viewers.

B.6 Distribution of the Measures

Table A9 reports the prevalence of each component indicator across the full 16,572-video archive and the distributions of the two continuous indices. Figure A2 shows the distributions of the two indices.

Table A9: Descriptive Statistics and Prevalence of Influencer-Style Communication Indicators

Indicator	<i>N</i>	Mean	SD	Min.	25th	Median	75th	Max.	Prevalence
Informal production	16,572	0.534	0.499	0	0	1	1	1	53.4%
Informal candidate presentation	16,572	0.478	0.500	0	0	0	1	1	47.8%
Direct address	16,572	0.362	0.481	0	0	0	1	1	36.2%
Audience acknowledgment	16,572	0.114	0.318	0	0	0	0	1	11.4%
Self-disclosure	16,572	0.227	0.419	0	0	0	0	1	22.7%
Interaction invitation	16,572	0.103	0.304	0	0	0	0	1	10.3%
Identity bridging	16,572	0.109	0.312	0	0	0	0	1	10.9%

Note: All indicators are binary measures coded at the Reel level as present (1) or absent (0). Prevalence reports the percentage of Reels in which each indicator is present. Because the indicators are binary, the mean is equivalent to prevalence expressed as a proportion. The sample contains 16,572 Reels.




Figure A2: Distribution of Aesthetic Proximity and Relational Signaling Indices

B.7 Alternative Thresholds and Measurement Robustness

The primary regression analyses retain both indices in their continuous 0–1 form and therefore do not depend on a high-versus-low threshold. A threshold is used only for the descriptive joint-distribution analysis that summarizes the combinations of aesthetic proximity and relational signaling appearing in candidate videos.

In the main descriptive analysis, I classify an index as “high” when it is at least 0.5. For aesthetic proximity, this means that at least one of the two aesthetic cues is present. For relational signaling, it means that at least three of the five relational cues are present.

Because the two indices contain different numbers of component indicators, the substantive interpretation of a common 0.5 threshold is not identical across dimensions. I therefore examine whether the conclusion that candidates disproportionately adopt aesthetic rather than relational features depends on this particular cutoff. Table A10 reports the joint distribution under alternative definitions of “high” communication cues.

Table A10: Joint Distribution of Aesthetic Proximity and Relational Signaling Under Alternative Thresholds

Aesthetic Threshold	Relational Threshold	Low/Low	High Aesthetic Only	High Relational Only	High/High
$\geq .50$	$\geq .50$	40.0%	50.2%	1.4%	8.4%
$\geq .50$	$\geq .20$	27.5%	20.1%	13.8%	38.5%
$\geq .50$	$\geq .40$	35.6%	38.6%	5.8%	20.0%
$\geq .50$	$\geq .60$	40.0%	50.2%	1.4%	8.4%
$= 1.00$	$\geq .50$	55.1%	35.0%	2.3%	7.6%
> 0	> 0	27.5%	20.1%	13.8%	38.5%

Note: Entries report percentages of the complete archive of 16,572 Reels. “Low/Low” denotes videos below both stated thresholds. “High aesthetic only” denotes videos at or above the stated aesthetic threshold but below the relational threshold, while “high relational only” denotes videos below the aesthetic threshold but at or above the relational threshold. “High/High” denotes videos meeting both thresholds. The regression analyses use the original continuous indices and therefore do not depend on these cutoffs.

Across these alternative definitions, the size of the aesthetic–relational gap depends on how much relational signaling is required to classify a video as “high.” Under the most permissive definition, where the presence of any relational cue qualifies ($R \leq .20$), 52.3% of videos contain high relational signaling, compared with 58.6% containing at least one aesthetic-proximity cue. The difference becomes substantially larger as the relational threshold increases. Requiring at least two of the five relational cues ($R \leq .40$) reduces the share classified as high relational to 25.8%, while requiring at least three ($R \leq .60$) reduces it to 9.8%. By comparison, 58.6% of videos contain at least one of the two aesthetic cues, and even the stricter requirement that both aesthetic cues be present classifies 42.6% of videos as high aesthetic proximity.

These comparisons qualify, rather than overturn, the main descriptive finding. Candidates frequently use individual relational cues, especially when a single cue is sufficient for classification, but they much less often combine multiple relational practices in the same

video. Aesthetic proximity is more commonly adopted as a bundle: under every specification, videos classified as high aesthetic but low relational outnumber those classified as high relational but low aesthetic. The main conclusion therefore does not depend on a single cutoff, but it is most accurately stated as an asymmetry in the *depth* of adoption: candidates commonly adopt the visual side of the influencer repertoire, whereas sustained or multi-cue relational signaling remains considerably less common.

B.8 Summary

The measurement strategy is designed to preserve the distinction between two components that often co-occur in influencer content but need not move together in political communication. Aesthetic proximity measures the visual and embodied presentation of the candidate, while relational signaling measures communicative practices that simulate interpersonal exchange with viewers. Seven independently coded indicators operationalize these dimensions across the complete Reel archive.

The automated annotation procedure closely reproduces human coding on the random validation sample, with 95.1% overall agreement and indicator-level agreement between 88% and 100%. The primary analyses use continuous normalized indices, while alternative threshold and component-level analyses assess whether the descriptive conclusions depend on a particular method of aggregation.

C Generative Stimulus Construction

Studies 2 and 3 use generative AI to construct controlled audiovisual treatments that vary how a political candidate communicates while holding other features of the candidate and message as constant as possible. This appendix documents the full stimulus-construction process. The goal of the pipeline is not simply to generate realistic political videos, but to create experimental stimuli in which theoretically meaningful communication cues can be manipulated independently. In ordinary campaign communication, visual presentation, candidate identity, voice, policy content, production quality, and interpersonal style often change together. The modular pipeline separates these components so that aesthetic proximity and relational signaling can be varied while preserving the features that should remain constant for causal identification.

Three principles guided the design of the pipeline. First, *identification*: matched treatment stimuli should differ primarily on the intended communication cues rather than on candidate identity, voice, policy position, or other characteristics that could independently affect voter evaluations. Second, *ecological plausibility*: the resulting stimuli should fall within the range of audiovisual styles that viewers might reasonably encounter in contemporary short-form political communication. The relevant standard is therefore not that every generated clip be indistinguishable from human-recorded video, but that the stimuli are sufficiently credible for respondents to evaluate them as candidate communication and that unintended differences in realism across treatment conditions are minimized and assessed directly. Appendix D reports the separate stimulus-validation procedures used to evaluate these properties.

Third, the pipeline is designed around *archival reproducibility*. Contemporary generative models are stochastic, and many of the tools used here operate through interfaces that do not guarantee that identical prompts and settings will reproduce identical outputs. Reproducibility therefore depends on preserving the exact stimuli shown to respondents, together with the scripts, prompts, model and provider information, generation settings where available, intermediate outputs, and the researcher decisions involved in selecting the final stimuli. The contribution of the pipeline is accordingly architectural rather than vendor-specific: it provides a general workflow for separating manipulated from controlled components, generating each component in modular stages, screening the resulting stimuli, and archiving the final experimental objects. The sections below document each stage of that process in detail.

C.1 Experimental Stimulus Design and Identification Strategy

The experimental stimuli were designed to isolate the two dimensions of influencer-style communication developed in Section 3.1 in the main text and in Appendix B: aesthetic proximity and relational signaling. The central challenge is that these dimensions are difficult to manipulate cleanly with naturally occurring political videos. A candidate who records a casual vertical video from a kitchen rather than a professionally produced campaign advertisement may also differ in identity, appearance, voice, policy position, delivery, resources, and reputation. Similarly, a candidate who speaks directly to viewers and shares personal experiences may differ from a candidate using a conventional broadcast register on dimensions other than relational signaling. Comparing such videos would

therefore make it difficult to attribute differences in voter response to communication style itself.

The stimulus-construction strategy addresses this problem by decomposing a political video into separately controllable components. Candidate identity and appearance, voice, policy content, visual presentation, relational language, animation, and final video assembly are treated as distinct elements of the stimulus. The experimental treatments deliberately vary the components corresponding to aesthetic proximity and relational signaling while holding other components as constant as possible within the relevant treatment comparison. Study 2 manipulates aesthetic proximity while holding the spoken message fixed within each matched comparison. Study 3 independently manipulates aesthetic proximity and relational signaling in a 2×2 factorial design, allowing the two dimensions and their interaction to be evaluated separately.

Importantly, the treatments are designed as *bundles* of theoretically related cues rather than manipulations of individual stylistic features. This choice mirrors how the constructs are defined in Study 1 and how these communication styles typically appear in practice. The resulting estimates therefore identify the effect of assignment to a higher- rather than lower-aesthetic-proximity presentation, or to higher rather than lower relational signaling, but they do not identify which particular component within a treatment bundle produces an observed effect. Cue-level decomposition would require additional experimental conditions and is beyond the scope of the present studies.

C.1.1 Features Manipulated

Aesthetic proximity. The aesthetic-proximity manipulation changes the visual and embodied presentation of the candidate while leaving the substantive political message unchanged. It was designed to move both components of aesthetic proximity measured in Study 1: *informal production* and *informal candidate presentation*. Informal production concerns whether the filming style and setting resemble casual, platform-native short-form content rather than conventional campaign or broadcast production. Informal candidate presentation concerns whether the candidate’s attire, posture, movement, and delivery present them in a casual or off-stage manner. In the low-aesthetic condition, the candidate appears in a more conventional political presentation characterized by horizontal framing, formal attire, an institutional or professional setting, and polished lighting and production. In the high-aesthetic condition, the same candidate appears in a more platform-native presentation characterized by vertical or selfie-style framing, casual attire, and an everyday domestic setting. These changes were implemented jointly rather than one at a time, producing a treatment contrast between two recognizable modes of candidate self-presentation rather than between isolated visual attributes. Study 2 manipulates only this aesthetic dimension. Within each issue-position comparison, respondents encounter versions of the candidate that deliver the same substantive message but differ in aesthetic presentation. Study 3 uses the same conceptual manipulation of aesthetic proximity as one factor in the 2×2 factorial design.

Relational signaling. The relational-signaling manipulation changes the candidate’s communicative relationship with the viewer through the spoken script. It corresponds directly to the five indicators measured in Study 1: direct address, audience acknowledg-

ment, self-disclosure, interaction invitation, and identity bridging. These cues are manipulated through verbal content rather than through the candidate’s visual appearance. The low-relational condition uses a conventional broadcast-style policy register. The candidate communicates the policy problem and proposed response without presenting the viewer as a conversational partner: the script contains no personal anecdote, no simulated acknowledgment of viewers’ prior reactions or questions, no explicit construction of shared identity, and no invitation for viewers to respond. The high-relational condition preserves the same issue domain, policy position, and substantive proposals while adding the five relational cues. The candidate speaks directly to the viewer, acknowledges what viewers have been saying or experiencing, discloses a personal experience related to the policy problem, explicitly places herself and the audience within a shared social group, and closes by inviting further interaction. In Study 3, for example, the high-relational cost-of-living message references a personal grocery-shopping experience, while retaining the same substantive policy content as the low-relational version. Section C.2 reports the complete scripts and maps each treatment component onto the corresponding codebook indicator. Study 2 does not manipulate relational signaling: the relational content of the script is held fixed within the aesthetic treatment comparisons. Study 3 interacts the low- and high-relational scripts with the low- and high-aesthetic presentations, yielding four conditions: low aesthetic/low relational, high aesthetic/low relational, low aesthetic/high relational, and high aesthetic/high relational.

C.1.2 Features Held Constant

The modular design is intended to prevent the treatment manipulations from inadvertently changing features of the candidate or message that could independently affect voter evaluations. Table A11 summarizes the resulting identification strategy. The meaning of “held constant” differs slightly across the two studies because Study 2 deliberately exposes respondents to multiple issue and ideological configurations, whereas Study 3 uses a single candidate and policy message.

Candidate identity and appearance. Within each relevant aesthetic comparison, the stimuli depict the same fictional candidate. Rather than generating unrelated “professional” and “influencer” candidates, the low- and high-aesthetic images were constructed using a shared facial reference. The low-aesthetic or professional image served as the identity anchor, and the corresponding high-aesthetic or influencer-style image was generated using the same facial reference so that facial structure and other identity markers remained as stable as possible while attire, framing, and setting changed. Section C.4 describes the image-generation and identity-matching procedure in detail. Study 2 contains four fictional candidate identities. The gun-policy stimuli use two male candidate avatars and the reproductive-policy stimuli use two female candidate avatars. The avatars are randomized across policy positions within each policy domain. Study 3 uses the same female candidate identity employed in the reproductive-policy portion of Study 2. Because issue domain and implied candidate sex are confounded in Study 2 by design, the experiment does not use their comparison to estimate either an issue-domain or candidate-sex effect.

Table A11: Manipulated and Controlled Features of the Experimental Stimuli

Feature	Study 2	Study 3	Implementation
Facial identity	Held constant within matched aesthetic comparisons	Held constant across all four conditions	Reference-image anchoring preserves the same apparent candidate identity across visual conditions
Voice	Held constant within matched comparisons	Held constant across all four conditions	The same synthetic voice identity is used across the relevant treatment versions
Issue domain	Held constant within each aesthetic comparison; varies across experimental blocks	Held constant across all four conditions	Study 2 uses gun-policy and reproductive-policy domains; Study 3 uses a cost-of-living message
Policy position	Held constant within each aesthetic comparison; varies across ideological versions	Held constant across all four conditions	The same substantive position and policy proposals are maintained within treatment contrasts
Spoken policy content	Held constant within each aesthetic comparison	Held substantively constant across relational conditions	Study 2 uses an identical script across aesthetic versions; Study 3 uses matched low- and high-relational scripts
Approximate video duration	Held as constant as possible within matched comparisons	Held as constant as possible across conditions	Audio and script lengths are matched as closely as possible (roughly 25s for Study 2 and 48s for Study 3).
Aesthetic presentation	Manipulated	Manipulated	Framing and orientation, attire, setting, candidate presentation, and production cues are varied across conditions
Relational signaling	Held constant within aesthetic comparisons	Manipulated	Script features vary direct address, audience acknowledgment, self-disclosure, identity bridging, and interaction invitation

Note: “Held constant” refers to the treatment contrast relevant for identifying the effect of interest. Study 2 intentionally varies issue domain and ideological position across stimuli, but these features do not vary between the high- and low-aesthetic versions of the same matched message.

Voice and vocal identity. Synthetic voices were generated separately from the visual candidate images. The same voice identity is used across the high- and low-aesthetic versions of a matched message so that aesthetic treatment assignment does not also change the candidate’s vocal identity. Study 3 similarly preserves the same female synthetic voice across all four factorial conditions. The exact audio file is used across low/high aesthetic conditions when the script is identical. When the relational script changes in Study 3, new audio necessarily accompanies the alternative text, but the same underlying synthetic voice specification is used. Section C.3 provides the complete voice-generation procedure and reports.

Policy substance and ideological position. Policy content is held constant within every treatment contrast used to estimate an aesthetic effect. In Study 2, the high- and low-aesthetic versions of a given stimulus use the same script, so the candidate’s issue position, arguments, and wording do not vary with aesthetic treatment assignment. The study separately includes liberal and conservative positions within the gun-policy and reproductive-policy domains, but ideological position is not itself part of the aesthetic manipulation. In Study 3, manipulating relational signaling requires some wording to change. The two scripts were therefore constructed to preserve the same policy problem, substantive position, and proposed policy response while changing the relational register through which that content is delivered. The low- and high-relational versions are matched as closely as possible in length and spoken duration. The final versions last approximately 48 seconds. Section C.2 presents the full scripts side by side and reports additional comparisons of their observable textual and acoustic properties.

Animation and final assembly. The treatment design also requires that downstream production choices not systematically introduce new treatment differences after the still images and audio have been generated. The same general animation, lip-synchronization, rendering, and screening procedure was therefore applied across experimental conditions. Identical generation settings and motion prompts were used across conditions. Final videos were screened for facial instability, unnatural movement, lip-sync errors, and other generation artifacts before entering the realism validation test and the experimental stimulus set. Sections C.5–C.6.3 document these procedures and the points at which researcher judgment entered the selection process. Holding these features constant is necessarily a design objective rather than a claim that synthetic videos are perceptually identical on every unintended dimension. Generative models may introduce small differences that are not explicitly requested by the researcher, and the visual manipulation itself may alter perceptions beyond the intended theoretical construct. For this reason, stimulus construction was followed by separate validation procedures examining treatment fidelity and perceived realism. Appendix D reports those analyses, including the observed difference in perceived realism across aesthetic conditions.

C.1.3 Relationship to Study 1 Measurement

The experimental treatments were constructed from the same conceptual and operational framework used to measure naturally occurring candidate communication in Study 1.

This correspondence is important because the observational and experimental components are intended to address different questions about the same communication dimensions. Study 1 asks whether candidates use aesthetic-proximity and relational-signaling cues and whether those cues are associated with platform feedback. Studies 2 and 3 then manipulate those dimensions to estimate their causal effects on voter evaluations.

Table A12 makes this correspondence explicit. The two aesthetic indicators used in the Reel codebook are translated into visual treatment features. Informal production is represented through differences in framing, camera perspective, setting, and production polish, while informal candidate presentation is represented through differences in attire and posture. The five relational indicators are implemented directly in the Study 3 high-relational script.

Table A12: Mapping Study 1 Measures onto the Experimental Treatments

Dimension	Study 1 Indicator	Experimental Operationalization
<i>Aesthetic proximity</i>	Informal production	High-aesthetic stimuli use platform-native visual cues, including vertical or selfie-style framing, an everyday setting, and lower institutional production polish; low-aesthetic stimuli use horizontal framing, institutional or professional settings, and a more polished campaign presentation.
	Informal candidate presentation	High-aesthetic stimuli present the candidate in casual attire and relaxed posture; low-aesthetic stimuli use formal attire and formal posture.
<i>Relational signaling</i>	Direct address	The high-relational script treats the individual viewer as the conversational addressee, whereas the low-relational script uses a more broadcast-oriented register.
	Audience acknowledgment	The high-relational script acknowledges or simulates awareness of what viewers have said, asked, or experienced; corresponding acknowledgment is absent from the low-relational script.
	Self-disclosure	The high-relational script reveals a personal experience related to the policy problem; corresponding personal disclosure is absent from the low-relational script.
	Interaction invitation	The high-relational script explicitly invites viewers to respond or communicate; the low-relational script does not.
	Identity bridging	The high-relational script explicitly places the candidate and viewer within a shared group or community; the low-relational script discusses the policy problem without constructing this shared identity.

Note: The experimental treatments operationalize the two broader dimensions as bundles. The table identifies the intended correspondence between the observational indicators and treatment features; it does not imply that the causal effect of any individual cue is separately identified.

C.2 Script Construction

The first stage of the pipeline constructs the spoken messages independently of the visual treatments. The objective is to preserve the candidate’s substantive political position while varying relational signaling only where required by the experimental design. Study 2 holds the script exactly constant within each aesthetic treatment comparison. Study 3 uses two versions of the same policy message that differ in relational signaling while preserving the same issue, policy position, and substantive proposals.

C.2.1 General Principles

I drafted the substantive content of the scripts manually. The scripts were written to sound like plausible candidate messages while keeping the treatment language confined to the intended experimental dimension. In Study 2, this required identical spoken content across the high- and low-aesthetic versions of each message. In Study 3, it required changing the relational register without changing the underlying policy argument. I revised the scripts iteratively for naturalness, treatment fidelity, and spoken length. I used ChatGPT (with GPT-5.1) as a writing aid during this process, primarily to identify wording changes that could reduce differences in length across matched scripts. The model was instructed to rephrase sentences without changing the policy position, substantive proposals, or intended relational cues. It did not determine the experimental treatments or independently generate the final scripts. I manually reviewed and edited all suggested revisions before producing the final versions. No LLM-generated wording was accepted without manual review. The final scripts were then converted to synthetic speech as described in Section C.3. Because spoken duration depends on both text length and vocal delivery, I evaluated balance using the rendered audio rather than word count alone.

C.2.2 Study 2 Scripts

Study 2 isolates aesthetic proximity. The spoken message therefore remains identical between the low- and high-aesthetic versions of each stimulus. Only the candidate’s visual presentation changes. Respondents were assigned to either a gun-policy or reproductive-policy issue domain. Within each domain, the experiment included two ideological positions, producing four substantive scripts:

- Reproductive Rights (Liberal): *When it comes to pregnancy, no two stories are the same. Some are filled with joy, others with heartbreak. But one thing is always true: those decisions belong to women – not politicians! I will fight to protect reproductive freedom, so every person can make choices about their body, their family, and their future, with dignity and without interference.*
- Reproductive Rights (Conservative): *Every life is precious, and every child is a gift. I believe we can build a culture that values both women and unborn children — one where families get the support they need to choose life. We should make adoption easier, expand care for new mothers, and ensure no woman ever feels alone in that decision. Protecting life and offering compassion can go hand in hand.*

- Gun Policy (Liberal): *How many times do we have to watch the same story unfold — another shooting, another community in pain? I refuse to accept that this is normal. It’s time for stronger background checks, safe storage laws, and keeping weapons of war off our streets. We can protect our kids, our communities, and our future — if we finally choose courage over silence.*
- Gun Policy (Conservative): *Every parent wants their kids safe at school and in their community. We can protect families without taking away the rights of responsible gun owners. That means enforcing existing laws, improving mental health care, and keeping guns out of the hands of criminals — not law-abiding citizens. Safety and freedom can coexist if we have the courage to act responsibly.*

C.2.3 Study 3 Scripts

Study 3 manipulates relational signaling through two versions of the same cost-of-living message. Both versions describe the same policy problem and advocate the same substantive response. They differ in whether the candidate communicates in a conventional broadcast register or uses the five relational cues defined in Appendix B. The low-relational script uses a conventional policy-oriented register. It does not include a personal anecdote, explicit acknowledgment of viewers’ prior reactions, shared-identity language, or an invitation to interact. The high-relational script preserves the same substantive policy content while adding all five relational cues. The candidate directly addresses the viewer, acknowledges what viewers have been telling her, shares a personal grocery-shopping experience involving her children, invokes a shared identity through language such as “families like ours,” and ends by inviting viewers to comment or message her. These additions were designed to change the candidate’s relationship to the audience rather than her policy position. The final scripts were matched in spoken duration at approximately 48 seconds. The complete scripts are reproduced below. In the high-relational version, **bold text** marks language introduced to implement relational signaling.

Below are the scripts used in Experiment 2:

1. Low Relational Signaling: (106 words with 778 characters, about 48 seconds) *“The rising cost of living has increased significantly in recent years, placing sustained pressure on household budgets across the country. Recent data show that groceries, rent, childcare, and prescription drugs continue to rise faster than overall wages for many households. Addressing these challenges requires targeted policy solutions that strengthen supply chains, promote market competition, and reduce structural cost pressures across key sectors of the economy. Government should focus on lowering everyday costs, preventing price gouging, and expanding access to affordable housing and childcare so that households can achieve greater financial stability over time. These steps are essential to ensuring long-term economic security and overall financial stability.”*
2. High Relational Signaling: (127 words 741 characters, about 48 seconds) *“**So a lot of you have been telling me** how expensive everyday life has gotten lately. **Honestly, I’m feeling it too. The other day I was at the grocery store**”*

*with my kids, looking at the total, and I just thought: how did it get this high? Groceries, rent, childcare, and prescription drugs cost way more than they should, and **families like ours** shouldn't have to carry that kind of pressure alone. **We should** focus on policies to lower everyday costs, stop price gouging, and expand affordable housing and childcare so we can build real financial stability together. **Tell me** what rising costs are hitting your family the most? **Comment down below or send me a message!** I really want to hear what you're going through."*

C.3 Synthetic Voice Construction

The second stage of the pipeline converts the experimental scripts into synthetic speech. I used synthetic rather than recorded human voices for two reasons. First, a reusable synthetic voice allows vocal identity to remain constant across matched treatment conditions. Second, it avoids cloning or appropriating the voice of an identifiable person. The male and female voices used in the experiments were created specifically for this project and were not based on recordings of real speakers.

C.3.1 Voice Design

I created the underlying voice identities using ElevenLabs Voice Design, which generates a new synthetic voice from a natural-language description rather than from a recording of an existing speaker. Voices were designed in October, 2025. Voice Design generated candidate voices from the specifications below, after which I selected one male and one female voice for use in the experiments. The voice descriptions targeted candidates in their mid-30s speaking native American English with a General American accent. I aimed for voices that sounded credible for a political candidate without sounding like professional announcers. The desired delivery was clear, confident, warm, and conversational, with moderate pacing and natural intonation.

Male voice-design prompt.

Native American English with a neutral General American accent. Male, in the mid-30s. Excellent audio quality. Persona: credible political candidate. Emotion: confident, warm, conversational. A natural mid-range male voice with clear articulation, moderate resonance, and an approachable conversational quality. Speaks at a steady, natural pace with varied intonation. Professional without sounding like a broadcaster or commercial voice-over artist; polished enough for public speaking but natural enough to sound like an ordinary candidate speaking directly to voters.

Female voice-design prompt.

Native American English with a neutral General American accent. Female, in the mid-30s. Excellent audio quality. Persona: credible political candidate. Emotion: confident, warm, conversational. A natural mid-range female

voice with clear articulation, moderate resonance, and an approachable conversational quality. Speaks at a steady, natural pace with varied intonation. Professional without sounding like a broadcaster or commercial voice-over artist; polished enough for public speaking but natural enough to sound like an ordinary candidate speaking directly to voters.

I generated 10 rounds of candidate voices with 5 male and 5 female voices in total under the Creator plan on ElevenLabs. Each voice generated an audio under the default settings of the following script:

Campaigns are contests of competing visions. The country chooses one or the other. And we accept the choice the country made. I've said many times you can't love your country only when you win. You can't love your neighbor only when you agree. Something I hope we can do no matter who you voted for is see each other not as adversaries, but as fellow Americans, bring down the temperature.

Then, I manually selected the final voices based on naturalness, intelligibility, age plausibility, and suitability for a candidate speaking in both conventional and social-media-style settings.

C.3.2 Text-to-Speech Generation

After selecting the two voice identities, I generated the experimental audio using ElevenLabs' Eleven Multilingual v2 text-to-speech model (`eleven_multilingual_v2`). Across scripts, I fixed stability at 0.60, similarity at 0.65, and style exaggeration at 0.50. These settings were chosen based on results from a grid-search of each parameter from 0 to 1 at 0.1 intervals. These settings were held constant so that changes in the experimental audio did not arise from different synthesis parameters.

Speed was the only synthesis parameter adjusted across the Study 3 relational scripts. Because adding relational language changed the natural duration of the high-relational message, I set its speed to 1.10 while retaining the default speed of 1.00 for the low-relational message. This produced final audio clips of approximately 48 seconds for both scripts. The adjustment was made to match exposure duration rather than to create a difference in vocal style. Within each aesthetic treatment pair, the script and synthetic voice were identical. The same final audio file was used for both the high- and low-aesthetic videos, eliminating audio differences within the aesthetic treatment contrast.

C.3.3 Generation and Selection

Text-to-speech generation is stochastic, so identical text and settings can yield somewhat different prosody across renders. I therefore generated 100 candidate audio clips for each script and screened the resulting clips for pronunciation, fluency, pacing, unnatural pauses, prosodic artifacts, and overall naturalness. The final clip for each script was selected from this candidate set using the same criteria. Clips were rejected when they contained clear pronunciation errors, unnatural pauses or emphasis, abrupt changes in pitch or tempo, audible synthesis artifacts, or delivery that sounded implausible for a candidate message. Selection was based on audio quality and naturalness. The same underlying

male and female voice identities were retained across studies. The male voice was used for Study 2’s gun-policy stimuli, while the female voice was used for Study 2’s reproductive-policy stimuli and Study 3. Thus, treatment comparisons change the intended visual or relational features without simultaneously substituting a different speaker identity. The replication archive preserves the final audio files shown in the experiments. Because ElevenLabs speech synthesis is nondeterministic, these archived audio files constitute the exact experimental objects; regenerating the same script with the same voice and settings is not expected to produce a bit-for-bit identical recording.

C.4 Candidate Avatar and Still-Image Construction

The third stage of the pipeline constructs the fictional candidate avatars used as the visual basis for the experimental videos. The goal is to vary aesthetic presentation while holding the candidate’s facial identity as constant as possible. I generated the low-aesthetic still images first and then used Kling AI’s facial-consistency tool to convert them into corresponding high-aesthetic versions. I used the KlingAI web interface rather than the API, so the generation process did not expose a user-specified seed. The platform’s default stochastic generation settings were therefore used.

The low-aesthetic image was generated first and served as the identity anchor. The corresponding high-aesthetic image was then created from that reference. This sequence was intended to change the candidate’s presentation without independently regenerating the candidate’s face.

C.4.1 Study 2 Candidate Images

Study 2 used two female and two male fictional candidates.¹³ Both were designed to appear white, approximately in their mid-30s, and plausible as relatively young political candidates. The female candidates were used for the reproductive-policy stimuli and the male candidate for the gun-policy stimuli. For each candidate, I first generated a pool of conventional campaign-style images in horizontal 16:9 format. The low-aesthetic condition used formal attire, upright posture, a professional office setting, an American flag and bookshelf, steady waist-up framing, and polished lighting. Candidate expression and eye contact were held relatively neutral and consistent so that the treatment centered on presentation rather than emotion. The consolidated low-aesthetic prompts were:

Female low-aesthetic image.

Create a photorealistic image of a fictional white female political candidate approximately in her mid-30s. She has shoulder-length blonde hair, neatly but naturally styled. She is wearing a tailored navy blazer over a white blouse and sitting upright in a composed, professional posture. Show her waist-up in

¹³I originally generated four female and four male candidates. After pairwise comparison of the final videos conducted by three undergraduate research assistants, I picked the two most realist pairs for female and male candidates to use as the final experimental stimuli.

a clean political campaign office, with an American flag and bookshelf visible in the background. She looks directly into the camera with a serious, attentive, neutral-to-warm expression. Use stable horizontal 16:9 framing and balanced professional lighting. Preserve realistic skin texture, natural facial proportions, and ordinary human imperfections. The image should resemble a professionally recorded political campaign video rather than a glamour portrait or stock photograph. No text, logos, watermark, exaggerated expression, or stylized features.

Male low-aesthetic image.

Create a photorealistic image of a fictional white male political candidate in his mid-30s. He has short dark-brown hair and light stubble. He is wearing a navy suit, white dress shirt, and red tie and sitting upright in a composed, professional posture. Show him waist-up in a clean political campaign office, with an American flag and bookshelf visible in the background. He looks directly into the camera with a serious, attentive, neutral-to-warm expression. Use stable horizontal 16:9 framing and balanced professional lighting. Preserve realistic skin texture, natural facial proportions, and ordinary human imperfections. The image should resemble a professionally recorded political campaign video rather than a glamour portrait or stock photograph. No text, logos, watermark, exaggerated expression, or stylized features.

I generated 50 images for each candidate avatar using Kling Image 1.5, the most updated version of Kling AI Image models at the time of generation, from which the final image was selected to be used for video generation.¹⁴ Candidate images were screened for realism, facial plausibility, and suitability for conversion into a matched high-aesthetic version. The high-aesthetic images were generated by supplying the corresponding low-aesthetic image as a facial reference and instructing Kling AI Image 1.5 to preserve the candidate’s identity while changing the visual presentation. For both male and female conversions, facial-reference strength was set to 90.¹⁵ The intended changes were attire, posture, setting, framing, and production style rather than facial identity.

A consolidated Study 2 conversion prompt is:

Study 2 high-aesthetic conversion.

Use the person in the reference image as the identity anchor. Preserve the same facial identity, facial proportions, skin tone, eye color, hair color, hairstyle, apparent age, and body build. Do not make the person younger, older, more attractive, or more expressive. Change only the presentation to a casual, platform-native social-media style. Dress the candidate in a plain fitted gray athletic T-shirt and casual joggers with no logos or visible branding.

¹⁴All images for Study 2 were generated in October, 2025.

¹⁵The reference strength setting is no longer available with the new models such as Kling AI Image

2.1 and 3.0. Users can use textual prompt to directly customize the reference strength.

Place the candidate sitting cross-legged beside a beige couch in a realistic, lived-in living room. Include ordinary background details such as a sofa and small coffee table with books and a porcelain coffee mug. Use a slightly elevated selfie-style camera angle and vertical 9:16 composition, framed approximately waist-up to three-quarter view. The candidate should look directly into the camera with the same serious, attentive, neutral-to-warm expression as the reference image. Use soft natural daylight, realistic skin texture, and ordinary household lighting rather than studio lighting. The image should resemble authentic user-generated short-form social-media content while remaining photorealistic. No phone frame, interface elements, text, logo, watermark, beauty filter, dramatic expression, or changes to facial identity.

For each candidate avatar (generated in the low-aesthetic condition), I then generated 50 high-aesthetic conversions and manually screened the resulting pairs. The final output from this step includes one low-aesthetic image and one high-aesthetic image for each candidate avatar. I originally created four male and four female candidate avatars, resulting in 16 generated images in total. Figure A3 and A4 show the example candidate avatars generated from the models.




Figure A3: Low-Aesthetic Proximity Condition

Note: Examples of real candidate stimuli used in the low-aesthetic proximity condition. The top row shows the four female candidate identities and the bottom row shows the four male candidate identities.

C.4.2 Study 3 Candidate Images

Study 3 reused the candidate identities developed and screened in Study 2 rather than generating a new set of candidates. The male and female reference images selected for the final experiment are shown in Figure A5. Both had performed well during the Study 2 stimulus-selection process. Before constructing the Study 3 stimuli, however, I updated these reference images using the newer Kling Image 3.0 model, which became available soon after Study 2. Specifically, I prompted the model to increase the photorealism of each image while preserving the candidate’s facial identity, appearance, and overall composition. The exact prompt was:

Female Candidates



Male Candidates



Figure A4: High-Aesthetic Proximity Condition

Note: Examples of real candidate stimuli used in the high-aesthetic proximity condition. The top row shows the four female candidate avatars and the bottom row shows the four male candidate avatars.

Enhance Image Realism.

Using the provided reference image, create a more photorealistic version of the same image. Preserve the subject's identity as closely as possible, including the same facial structure, facial features, hairstyle, hair color, skin tone, apparent age, gender presentation, body proportions, expression, and overall appearance. Preserve the same clothing, pose, camera angle, framing, background, lighting direction, and composition.

The resulting reference images, shown in Figure A6, served as the identity-consistent starting points for all subsequent Study 3 stimulus generation.

The Study 3 aesthetic manipulation was also refined to create a less extreme contrast between professional and platform-native presentation. Rather than defining high aesthetic proximity through workout clothing and sitting on the floor, the revised condition



(a) Female candidate



(b) Male candidate

Figure A5: Study 3 Candidate Stimuli Adopted from Study 2



(a) Female candidate



(b) Male candidate

Figure A6: Study 3 Candidate Stimuli Generated from Images Adopted from Study 2

depicts the candidate seated at a kitchen table in ordinary casual clothing. The purpose was to preserve the informal, self-produced features of short-form communication without making the candidate appear unusually careless or unserious. The low-aesthetic condition retained the conventional campaign presentation of the reference image: professional attire, stable centered framing, and a polished office setting. The high-aesthetic condition was generated from the reference images shown in Figure A6 using Kling Image 3.0. Figure A7 shows the female and male candidate images after the conversion. The consolidated high-aesthetic conversion prompt was:

Study 3 high-aesthetic image.

Use the reference image as the identity anchor and preserve the candidate's exact facial identity, facial proportions, skin tone, eye color, hair color, hairstyle, and apparent age. Do not alter facial structure or make the candidate younger, older, more attractive, or more expressive. Show the same candidate seated at a light maplewood kitchen table in a realistic modern home kitchen with white shaker cabinets during daytime. Dress the candidate in a plain light heather-gray crew-neck knit sweater with a relaxed fit and no logos or visible branding. Hair should remain natural and casually styled. Keep makeup minimal and natural. Frame the candidate chest-up, slightly above eye level and slightly off-center, as if a smartphone were being held at arm's length with the right arm. The candidate leans slightly forward with one forearm resting casually on the table and looks directly into the camera with a calm,

serious, neutral-to-warm expression. Show the scene directly through the camera; do not include a visible phone, phone frame, user-interface overlay, captions, or social-media graphics. Use natural daylight entering from one side, producing soft directional shadows rather than studio lighting. Preserve realistic skin texture, including natural pores and lines. Include subtle household details such as a small plant, wooden cutting board, and fruit bowl in a slightly out-of-focus background with realistic smartphone depth of field. The overall appearance should be casual, self-produced, and platform-native while remaining credible for a political candidate. It should not look like a professional portrait, advertisement, glamour photograph, or highly stylized image.




Figure A7: Study 3 Candidate Stimuli for the High-Aesthetic Condition

For each candidate reference image, I generated 50 candidate images, yielding 100 images in total, and selected the strongest output based on naturalness, identity consistency, and photorealism. The female conversion was generated first. For the male conversion, I additionally tested the selected high-aesthetic female image as a style reference while retaining the male low-aesthetic image as the identity reference.¹⁶ The model was instructed to reproduce the visual style of the female reference while preserving the male candidate’s facial identity, skin tone, eye color, and hair color. Stimuli for both

¹⁶Kling Image 3.0 allows up to ten reference images per generation. This version was selected as the final image to use in Study 3.

candidates were initially constructed for Study 3. During pre-fielding screening, however, the male stimuli were excluded because recurrent artifacts around the mouth and frown lines reduced their perceived naturalness, and their audio–visual lip synchronization was less consistent than that of the female stimuli. The final Study 3 experiment therefore used only the female candidate. Appendix D reports the validation and stimulus-selection procedure.

C.4.3 Image Selection, Reproducibility, and Reverse-Image Search

Image generation involved substantial researcher screening because facial consistency and realism cannot be guaranteed by the prompt alone. I evaluated candidate images on three criteria: (1) treatment fidelity, (2) facial consistency across aesthetic conditions, and (3) visual realism. Images were rejected for facial drift, distorted anatomy, inconsistent hair or facial features, implausible clothing or backgrounds, exaggerated expressions, or obvious generation artifacts. Selection occurred before the video generation process. The validation procedure used to select among completed video stimuli is reported separately in Appendix D. Because generation occurred through Kling AI’s consumer web interface, I cannot provide a user-set random seed or guarantee that the prompts will regenerate identical images. Reproducibility therefore rests on preserving the exact selected images and their provenance. The replication archive includes all final still images, the original generation prompts, reference images, recorded settings where available, and discarded or intermediate generations are also archived. Before deployment, I also screened the final synthetic candidate faces using Google’s reverse-image search to reduce the risk that a generated identity closely resembled a real identifiable person with an online presence. None of the final images produced a match to an existing real-person image online.

C.5 Image-to-Video Animation

The next stage of the pipeline animates the selected candidate still images into continuous talking-head videos. I used Kling AI for all image-to-video generation in both studies. Study 2 used Kling Video 2.5 Turbo, while Study 3 used Kling Video 3.0, which became available after the Study 2 stimuli had been produced. Video generation for Study 2 took place from October through December 2025, and Study 3 videos were generated from January through April 2026. The goal of this stage was to introduce natural head and upper-body movement without changing candidate identity, expression, setting, or other features established in the still images. Lip synchronization was performed separately in the next stage.

C.5.1 Models and Generation Settings

Table A13 summarizes the image-to-video configuration. Because generation was conducted through Kling AI’s web interface, I did not specify a random seed. The default platform model seed was used.

Table A13: Kling Image-to-Video Generation Settings

Setting	Study 2	Study 3
Provider	Kling AI	Kling AI
Model	Kling Video 2.5 Turbo	Kling Video 3.0
Mode	Image-to-video	Image-to-video
Generation period	Oct.–Dec. 2025	Jan.–Apr. 2026
Maximum generated clip length	10 seconds	15 seconds
Number of source segments per final video	3	5
Final stimulus length	Approx. 25 seconds	Approx. 48 seconds
Aspect ratio	16:9 for low-aesthetic condition; 9:16 for high-aesthetic condition	16:9 for low-aesthetic condition; 9:16 for high-aesthetic condition
Resolution	720p	720p
Seed	Model Default	Model Default

Note: Generated source segments were longer than some corresponding audio segments and were trimmed before final assembly. Final stimulus durations therefore reflect the matched audio rather than the maximum duration of the generated source clips.

C.5.2 Motion Prompts and Treatment Invariance

Animation prompts were designed to preserve the still-image treatment while limiting additional variation introduced by video generation. I used the same prompt for repeated generations from the same reference image. Thus, rejected and replacement clips within a candidate-by-aesthetic condition differed because of stochastic generation rather than changes in researcher instructions. The low- and high-aesthetic conditions required different camera behavior because filming style is itself part of the aesthetic-proximity treatment. The low-aesthetic condition therefore used a fixed, centered camera, whereas the high-aesthetic condition allowed subtle handheld movement. In both conditions, candidate expression, gaze, and bodily movement were kept restrained to avoid introducing an unintended difference in emotional intensity or expressiveness.

A consolidated low-aesthetic animation prompt is:

Low-aesthetic animation prompt.

Animate the person in the reference image while preserving the exact facial identity, facial proportions, hairstyle, skin tone, clothing, background, lighting, and overall composition. The candidate remains seated and looks directly into the camera with a serious, attentive, neutral-to-warm expression. Use a fixed, steady, centered camera consistent with a professionally recorded political campaign video. The candidate makes only small, natural head movements and subtle upper-body gestures while speaking. Keep movement restrained and realistic. Maintain stable posture, consistent eye contact, and natural blinking. Do not change the candidate’s face, age, hairstyle, clothing,

body shape, background objects, lighting, or camera position. Avoid exaggerated gestures, large head movements, camera movement, zooming, sudden expression changes, background motion, object deformation, facial distortion, or other visible AI artifacts. The result should resemble a continuous, professionally recorded talking-head campaign video.

A consolidated high-aesthetic prompt is:

High-aesthetic animation prompt.

Animate the person in the reference image while preserving the exact facial identity, facial proportions, hairstyle, skin tone, clothing, background, lighting, and overall composition. The candidate remains seated and looks directly into the camera with a serious, attentive, neutral-to-warm expression. Create subtle natural movement consistent with a self-recorded smartphone video from social media influencers. Allow only slight handheld camera micro-movement and small natural variations in framing. The candidate makes restrained head movements and subtle upper-body gestures while speaking, with steady eye contact and natural blinking. The movement should feel informal and user-generated without becoming shaky, dramatic, or distracting. Do not change the candidate’s face, age, hairstyle, clothing, body shape, background objects, or lighting. Avoid exaggerated gestures, large head movements, abrupt camera movement, zooming, sudden expression changes, background distortion, object deformation, facial drift, or other visible AI artifacts. The result should resemble a continuous, natural short-form social-media video rather than a cinematic or highly stylized animation.

In Study 3, animation instructions depended only on the aesthetic condition, not the relational condition. The high- and low-relational treatments were introduced through the audio and subsequent lip-synchronization stage. Thus, relational treatment assignment did not change the image-to-video prompt or intended movement style.

C.5.3 Sequential Segment Generation

Kling AI’s maximum clip length required each experimental video to be constructed from multiple generated segments. Study 2 used Kling Video 2.5 Turbo, which generated clips of up to 10 seconds. I therefore generated three 10-second source clips for each approximately 25-second final video. Study 3 used Kling Video 3.0, which allowed clips of up to 15 seconds. I generated five 15-second source clips to construct each approximately 48-second final video. Before animation, each final audio track was divided into shorter segments that did not exceed the corresponding model’s maximum clip length while stopping at natural pauses. Some audio segments were substantially shorter than the maximum. The generated video segment was therefore trimmed to the duration of its matched audio segment rather than stretching or altering the audio to fit the video. Generation proceeded sequentially. The selected still image served as the starting frame for the first segment. Once an acceptable first segment had been selected, I extracted its

final frame and used that frame as the starting image for the next generation. The same procedure was repeated for every subsequent segment:

Reference image $\rightarrow$ Segment 1 $\rightarrow$ final frame $\rightarrow$ Segment 2 $\rightarrow \dots$

Using the ending frame of the preceding segment as the starting frame of the next was intended to maintain facial identity, body position, background placement, lighting, and camera composition across segment boundaries. This procedure also reduced visible discontinuities when the segments were later concatenated.

C.5.4 Generation Screening and Stopping Rule

Video generation was stochastic and frequently produced unusable outputs. I therefore screened every generated segment before using its final frame to initiate the next segment. Clips were rejected when they contained visible artifacts or movement inconsistent with the intended treatment. Because video generation was substantially more costly than still-image generation, I did not generate a fixed large number of alternatives for every segment. Instead, I regenerated each segment until I obtained an acceptable output, then used that segment as the basis for the next stage of the sequence. The target was three acceptable sequential segments for each Study 2 video and five for each Study 3 video. The number of failed generations therefore varied across segments and conditions. Table A14 summarizes the principal rejection criteria.

Table A14: Image-to-Video Screening Criteria

Failure Mode	Screening Criterion	Response
Facial drift	Candidate facial structure or apparent identity changes visibly during the clip	Reject and regenerate
Unnatural movement	Implausible head, hand, torso, or facial movement, including repetitive or exaggerated gestures	Reject and regenerate
Background distortion	Objects appear, disappear, morph, or move implausibly	Reject and regenerate
Sudden camera motion	Abrupt camera movement inconsistent with the intended stable or subtle-handheld presentation	Reject and regenerate
Visible AI artifact	Warping, flickering, anatomical distortion, texture failure, or another obvious generation artifact	Reject and regenerate
Continuity failure	A segment ending provides a poor starting point for the next segment or creates a visible discontinuity across clips	Reject and regenerate
Treatment contamination	Generated motion or camera behavior introduces features inconsistent with the assigned aesthetic condition	Reject and regenerate

Note: Screening was conducted during stimulus generation, before validating the realism and fielding the survey experiments.

This procedure necessarily includes researcher judgment. There was no fixed ex ante number of generation attempts per segment; generation continued until a clip satisfied

the predefined substantive criteria of identity preservation, treatment fidelity, continuity, and visual naturalness. Once accepted, the clip was fixed and its final frame anchored the next generation. The exact final video files are therefore the reproducible experimental objects, rather than deterministic outputs that can be regenerated from the prompt alone. The completed animation sequences were then passed to the lip-synchronization stage described in Section C.6.

C.6 Lip Synchronization and Final Video Assembly

The following stage of the pipeline synchronizes each animated video segment to its corresponding synthetic-audio segment. Lip synchronization was performed at the segment level, before final assembly. Thus, each generated animation clip was independently synchronized to the matching portion of the speech, and the completed lip-synced clips were then concatenated to form the final stimulus.

C.6.1 Lip-Synchronization Model Selection

For Study 2, I used Kling AI’s built-in lip-synchronization tool. For Study 3, before selecting the final system, I tested several available lip-synchronization tools on the same male high-aesthetic sample clip. These included Wav2Lip, Kling’s built-in lip-synchronization tool, PixVerse, Dzine AI, and Sync Labs. The comparison was qualitative and application-specific rather than a general benchmark of the models. I evaluated whether each system could synchronize the supplied audio while preserving the candidate’s facial identity and producing a visually coherent mouth region across the full clip. The alternative systems performed poorly on this sample. Common failures included inaccurate phoneme-to-mouth alignment, unnatural or missing teeth, a generated mouth region whose skin tone or appearance did not match the rest of the face, degradation of facial identity, temporal instability across frames, and visible synthetic artifacts. Some outputs appeared to overlay a separately generated mouth onto the original face rather than preserving a coherent facial appearance. Sync Labs produced substantially cleaner results on the same clip, with better audio–mouth alignment, facial consistency, tooth rendering, and temporal stability. I therefore used Sync Labs for all final experimental stimuli. The final model used was Sync-3, accessed through the web interface during April, 2026. This screening procedure was intended only to select a tool capable of producing usable stimuli for this application. It should not be interpreted as evidence that Sync Labs performs better than the alternative systems generally.

C.6.2 Segment-Level Lip Synchronization

Each animation segment described in Section C.5 was paired with the corresponding audio segment generated in Section C.3. Lip synchronization was then performed separately for every segment rather than on the fully assembled video. For each segment, I generated two Sync Labs renders using the same source video, audio, and default synchronization settings. I visually inspected both outputs and retained the better render. Selection focused on:

- alignment between speech and mouth movement;

- natural teeth and mouth appearance;
- preservation of facial identity and skin tone;
- temporal stability of the mouth and surrounding facial region; and
- absence of visible lip-sync or generative artifacts.

If both renders contained clear failures, additional renders were generated. Otherwise, the better of the two renders was retained. The same procedure was applied across treatment conditions. Lip-sync settings did not vary by aesthetic or relational treatment. The synchronization stage therefore served only to align the previously generated speech with the candidate’s mouth movements rather than to introduce additional treatment variation. Where necessary, video segments were trimmed to match the duration of their corresponding audio segments before going through the lip-synchronization stage. No speech was stretched or compressed during the lip-synchronization stage.

C.6.3 Final Video Assembly

After lip synchronization, the selected segments were concatenated in chronological order to construct the final experimental video. Study 2 stimuli consisted of three synchronized segments and lasted approximately 25 seconds. Study 3 stimuli consisted of five synchronized segments and lasted approximately 48 seconds. Final assembly was performed in iMovie on a MacBook Pro with an Apple M2 processor. The already lip-synced clips were placed directly adjacent to one another in the timeline and exported as a single video. No transition effects, motion smoothing, color correction, filters, or other visual editing were applied during assembly. Because the final frame of each animation segment had been used to initialize the next segment during the generation stage, the adjacent clips were designed to join without requiring an artificial transition.

The exported files then resolution-reduced to 360p, and the lower-resolution videos were the exact stimuli presented to respondents. The replication archive preserves these final videos together with audio segments, lip-synced segments, and intermediate renders.

C.7 Ethics and Safeguards

The use of generative AI in this study raises issues of provenance, impersonation, participant disclosure, and reproducibility. I therefore adopted several safeguards across both the measurement and experimental components of the project. In Study 1, Gemini-3.1-Pro is used as a measurement tool to apply a predefined codebook to publicly available candidate videos. Model classifications are treated as generated measurements rather than ground truth and are validated against independent human coding, as described in Appendix B. The model is not used to generate, modify, or attribute statements to the candidates in the observational corpus.

For Studies 2 and 3, all experimental candidates are fictional. No real politician’s face or voice was used or cloned in constructing the stimuli. Candidate voices were generated from natural-language descriptions rather than recordings of identifiable speakers. I also screened the final synthetic candidate images using reverse-image search to reduce the risk that a generated identity closely resembled a real identifiable person with an

online presence. The resulting videos were created specifically as experimental stimuli and should not be interpreted as recordings or representations of actual candidates. Additional human screening and respondent-based validation of the resulting stimuli are reported in Appendix D.

Respondents were not informed that the videos were AI-generated before reporting the primary outcomes, because doing so could itself alter evaluations of the candidates and compromise the experimental comparison. At the conclusion of each experiment, respondents were explicitly informed that the videos were synthetically generated and did not depict real political candidates. The debriefing also provided information about identifying AI-generated and manipulated audiovisual content and discussed potential risks from synthetic media. The complete survey instruments, debriefing language, and IRB information are reported in Appendix E.

Finally, I treat provenance of the exact experimental objects as part of research reproducibility. Generative systems are stochastic, and several tools used here do not permit identical outputs to be regenerated from a prompt alone. The replication materials therefore preserve the final stimuli shown to respondents together with the treatment scripts, prompts, model information, generation settings where available, and intermediate materials needed to document how the stimuli were constructed and selected. This archival approach makes the realized experimental treatments, rather than the prompts alone, the primary reproducible research objects.

D Validation and Realism

This appendix evaluates the validity and perceived realism of the audiovisual stimuli used in Studies 2 and 3. Appendix C documents how the stimuli were constructed and the analyses below instead assess whether the completed videos preserved the intended experimental contrasts and were sufficiently credible for respondents to evaluate as candidate communication. I first describe the validation strategy, then report the pre-fielding screening conducted for Study 2 and the independent realism pilot conducted before Study 3.

D.1 Validation Strategy

Using fully synthetic audiovisual stimuli creates a tradeoff between experimental control and perceived realism. The generative pipeline makes it possible to vary aesthetic proximity and relational signaling cues while holding features such as candidate identity, voice, policy position, and substantive message comparatively constant. At the same time, generative models can introduce unintended differences in facial appearance, movement, lip synchronization, or other perceptual features that are not part of the intended treatment. Validation is therefore necessary to assess not only whether the treatments manipulate the intended communication styles, but also whether unintended properties of the generated videos could complicate their interpretation.

I evaluate the stimuli along two related dimensions. The first is treatment fidelity: whether the completed videos preserve the intended differences in aesthetic proximity and relational signaling cue bundles while maintaining candidate identity and avoiding salient generation defects. A treatment contrast is useful only if the intended stylistic differences survive the generation process without being overwhelmed by facial instability, unnatural motion, synchronization errors, or other artifacts. The second is perceptual credibility and comparability: whether viewers can plausibly evaluate the videos as instances of candidate communication and, especially, whether perceived realism differs systematically across experimental conditions. This distinction matters because a synthetic stimulus need not be indistinguishable from authentic footage to provide an informative experimental contrast. However, if one treatment condition appears substantially more artificial than another, the estimated treatment effect may partly reflect that unintended difference rather than the intended communication cues alone.

The validation process therefore proceeded in two stages. Before Study 2, three trained research assistants independently screened 32 generated candidate videos using structured evaluations of naturalness, facial consistency, movement, lip synchronization, and overall plausibility. These evaluations were used to select the final matched stimuli before fielding. Before Study 3, I conducted a separate realism pilot with 108 respondents recruited through Prolific. Respondents evaluated the experimental stimuli alongside authentic candidate videos from the Study 1 corpus and AI-generated political videos already circulating online, allowing the experimental stimuli to be located relative to both human-recorded and synthetic benchmarks. The pilot also permits a direct test of whether perceived realism covaries with assignment to the aesthetic-proximity or relational-signaling treatments.

The purpose of these analyses is therefore not to establish that the synthetic videos

are perceptually indistinguishable from human-recorded campaign footage. Rather, the relevant questions are whether the stimuli are credible enough to support candidate evaluations, how they compare with realistic external benchmarks, and whether any unintended realism differences align with the experimental treatments. The sections that follow document each validation stage and its implications for interpreting the treatment effects.

D.2 Study 2: Pre-Fielding Stimulus Screening

Study 2 underwent a structured pre-fielding screening procedure to select the final experimental stimuli from a larger set of 32 completed synthetic candidate videos. The purpose of this screening was not to determine which treatment condition was more appealing, but to identify matched high- and low-aesthetic videos that preserved candidate identity, implemented the intended visual contrast, and minimized visible generation artifacts. Screening occurred after the candidate images had been animated, synchronized to the synthetic audio, and assembled into completed videos. Appendix C documents the generation process itself, and this section describes how the completed candidate videos were evaluated and selected for fielding.

D.2.1 Candidate Videos Entering Screening

The initial candidate pool contained four female and four male fictional candidate identities. Female candidates were developed for the reproductive-policy stimuli and male candidates for the gun-policy stimuli. For each identity, matched low- and high-aesthetic versions were constructed so that the candidate’s facial identity and spoken message remained constant while visual presentation varied. The low-aesthetic versions used conventional campaign presentation, including formal attire, professional settings, horizontal framing, and polished production, whereas the corresponding high-aesthetic versions used casual attire, everyday settings, and platform-native framing. Appendix C provides the complete image-generation and video-construction procedure.

A total of 32 completed videos entered the pre-fielding screening. These videos represented four female and four male avatars, where the female avatars were assigned the reproductive-policy issue domain and the male avatars were assigned the gun-policy issue domain. Each avatar generated four videos: liberal-high-aesthetic, liberal-low-aesthetic, conservative-high-aesthetic, and conservative-low-aesthetic. Before entering the RA screening, the videos had already passed generation-stage quality control described in Appendix C generated segments containing salient facial drift, implausible movement, discontinuities, or other visible generation failures were rejected and regenerated during construction. The RA screening therefore constituted a second stage of validation applied to completed, otherwise usable stimuli rather than an initial screen of raw model outputs.

D.2.2 Pre-Fielding Exclusion of the Male Study 3 Stimuli

Both female and male stimulus sets were initially developed when extending the design to Study 3. During pre-fielding review, however, the male videos repeatedly exhibited unnatural artifacts that became especially visible during facial movement and

lip-synchronization. These artifacts reduced the perceived naturalness of the videos and could not be eliminated consistently through generations. I therefore excluded the male stimulus set before Study 3 was fielded and conducted the factorial experiment using the female candidate only. This decision was made on stimulus-quality grounds before fielding the survey experiment. The resulting Study 3 design consequently identifies the effects of aesthetic proximity and relational signaling for the fictional female candidate used in the experiment. It does not provide an experimental comparison of whether those effects vary by candidate sex. Candidate-level moderation, including potential differences by sex or gender, remains a question for future work.

D.3 Study 3: Realism Pilot Design

Before fielding Study 3, I conducted a separate respondent-based pilot to evaluate the perceived realism of the completed experimental stimuli. Unlike the Study 2 screening, which used trained research assistants to select among candidate videos, this pilot measured how ordinary viewers perceived the final videos. The design served two purposes. First, it located the experimental stimuli relative to authentic candidate videos and AI-generated political content already circulating online. Second, because respondents were randomly assigned to one of the four Study 3 treatment videos, it allowed perceived realism to be compared directly across the aesthetic-proximity and relational-signaling conditions.

D.3.1 Sample and Fielding

The pilot was fielded through Prolific on May 11, 2026. Participants were eligible if they were at least 18 years old, resided in the United States, and used English as their primary language. Recruitment used demographic quotas to balance respondents by sex. A total of 123 respondents were recruited, where respondents whose total completion time was less than one-third of the sample median completion time were excluded, resulting in a final analytic sample of 108.

D.3.2 Experimental Procedure

Each respondent viewed nine short political videos. One was randomly selected from the four completed Study 3 experimental stimuli: traditional broadcast (low aesthetic, low relational), aesthetic only (high aesthetic, low relational), relational only (low aesthetic, high relational), or the full influencer package (high aesthetic, high relational). The remaining eight videos were common benchmark stimuli consisting of four authentic candidate videos drawn from the Study 1 Instagram corpus and four existing AI-generated political videos collected from Instagram and Facebook.

Assignment to the experimental stimulus was randomized at the respondent level. The resulting cell sizes were: traditional style ($n = 27$), aesthetic only ($n = 28$), relational only ($n = 27$), and full influencer style bundle ($n = 26$). Because each participant viewed only one experimental stimulus, comparisons of perceived realism across the four treatment conditions are between-subjects. By contrast, every respondent evaluated all four authentic and all four benchmark-AI videos, resulting in 108 ratings for each.

All nine videos were presented in randomized order. Thus, the assigned experimental stimulus could appear at any position in the sequence rather than systematically preceding or following the benchmark videos. Respondents could replay videos, and video completion was required before the realism question became available. Audio was required during respondent recruitment, and an audiovisual manipulation check was required before proceeding to the videos.

The authentic clips provide a reference for the realism of human-recorded candidate communication, whereas the AI-generated and AI-manipulated clips provide a comparison with synthetic political content that viewers could plausibly encounter online. The benchmark AI stimuli were deliberately selected to represent the heterogeneous and imperfect quality of political synthetic media circulating on social media rather than highly polished model demonstrations.

D.4 Perceived Realism Measure

Immediately after each video, respondents evaluated whether the clip appeared to be AI-generated or real. The exact question was:

Do you think this video is real or AI-generated?

Responses were recorded on a four-category scale:

- Definitely real
- Probably real
- Probably AI-generated
- Definitely AI-generated

I code the measure from 1(=Definitely AI-generated) to 4(=Definitely real) so that higher values indicate greater perceived realism. Each respondent therefore contributed nine realism evaluations: one for an experimental stimulus, four for authentic candidate videos, and four for benchmark AI videos.

D.5 Benchmark Results

The realism pilot produced a clear ordering across the three video categories. Authentic candidate videos received the highest realism ratings (mean=3.04, SD=0.78), followed by the experimental stimuli (mean=1.94, SD=0.85), and then the existing AI-generated or AI-manipulated benchmark videos (mean=1.50, SD=0.77).

The experimental stimuli did not fully reproduce the perceived realism of authentic campaign footage. On average, experimental videos were rated 1.09 points lower than authentic candidate videos on the four-point scale ((mean_experimental=1.94) versus (mean_authentic=3.04); difference (=−1.09), 95% CI [−1.271, −0.914]). The magnitude of this difference indicates that respondents were generally able to distinguish the generated experimental videos from human-recorded campaign footage.

At the same time, the experimental stimuli were perceived as more realistic than the existing synthetic political videos used as the external AI benchmark. Experimental

stimuli averaged 1.94 compared with 1.50 for benchmark AI videos, a difference of 0.45 points with 95% CI [0.269, 0.625]. Thus, although the experimental stimuli fell short of the realism of authentic candidate footage, they exceeded the perceived realism of the AI-generated and AI-manipulated political content included as a benchmark for synthetic media that viewers could encounter online.

As expected, the largest difference occurred between the authentic and benchmark-AI categories. Authentic candidate videos were rated 1.54 points more realistic on average than benchmark AI videos mean=3.04 versus M=1.50 and 95% CI [1.435, 1.643]. These results establish that the realism measure meaningfully distinguishes known human-recorded footage from known synthetic or manipulated political content.

D.5.1 Pairwise Benchmark Comparisons

Table A15 reports the complete pairwise comparisons. For transparency, I report both differences on the original four-point realism scale and standardized effect sizes. The raw-scale differences are particularly useful because the outcome has a directly interpretable range from definitely AI-generated to definitely real.

Table A15: Pairwise Comparisons of Perceived Video Realism

Comparison	M_1 (SD)	M_2 (SD)	Difference	95% CI	t (df)	p	Cohen's d
My Videos ($n = 108$) vs. Real Videos ($n = 432$)	1.944 (0.852)	3.037 (0.782)	-1.093	[-1.271, -0.914]	-12.113 (155.1)	< .001	-1.372
My Videos ($n = 108$) vs. Benchmark AI ($n = 432$)	1.944 (0.852)	1.498 (0.774)	0.447	[0.269, 0.625]	4.962 (154.0)	< .001	0.566
Real Videos ($n = 432$) vs. Benchmark AI ($n = 432$)	3.037 (0.782)	1.498 (0.774)	1.539	[1.435, 1.643]	29.089 (861.9)	< .001	1.979

Note: Entries report pairwise Welch two-sample t -tests comparing perceived video realism across stimulus types. M_1 and M_2 correspond to the first and second groups listed in each comparison, with standard deviations in parentheses. Differences are calculated as the first group minus the second group. Confidence intervals are 95% confidence intervals for the mean difference. d denotes Cohen's d .

The standardized differences reinforce the same pattern. The experimental-authentic gap was large ($d = -1.37$), whereas experimental stimuli were moderately more realistic than the benchmark AI videos ($d = 0.57$). The authentic-versus-benchmark-AI difference was especially large ($d = 1.98$). These standardized quantities are secondary to the original-scale comparisons, but they provide an additional indication of the magnitude of the differences across benchmark categories.

D.5.2 Stimulus-Level Variation

The pooled comparisons could obscure heterogeneity among individual videos. Figure A8 therefore reports mean perceived realism and 95% confidence intervals separately for each of the four experimental stimuli, four authentic candidate videos, and four benchmark AI videos.




Figure A8: Perceived Realism Across Video Stimuli in Study 3

Note: Perceived realism ranges from 1 (“definitely AI-generated”) to 4 (“definitely real”). Points report mean ratings, and error bars report 95% confidence intervals. Experimental stimuli were randomly assigned such that each respondent evaluated one of the four treatment videos. All respondents evaluated the eight benchmark videos.

Taken together, the benchmark analyses place the experimental stimuli between authentic and existing synthetic political content. They were clearly less realistic than human-recorded campaign videos, but also more realistic than the selected benchmark AI videos. The next section examines the more consequential issue for causal interpretation: whether perceived realism varies systematically across the experimental treatment conditions.

D.6 Realism Across Experimental Conditions

The benchmark analyses establish that the experimental stimuli fall between authentic candidate footage and existing synthetic political content in perceived realism. For interpreting the Study 3 treatment effects, however, the more consequential question is whether perceived realism itself varies across the experimental conditions. If one treatment systematically appears more artificial than another, the corresponding treatment effect may reflect both the intended communication cues and an unintended change in perceived realism.

Because each pilot respondent was randomly assigned to one of the four Study 3 experimental stimuli, comparisons across experimental conditions are between subjects. I examine realism separately across the aesthetic-proximity and relational-signaling dimensions and then report the four factorial cells individually.

D.6.1 Aesthetic-Proximity Contrast

Perceived realism differed slightly across the aesthetic-proximity manipulation. Respondents assigned to a high-aesthetic stimulus rated the video an average of 1.74 on the

four-point realism scale ($SD=0.76$, $n = 54$), compared with 2.15 among respondents assigned to a low-aesthetic stimulus ($SD=0.90$, $n = 54$). The estimated difference was (-0.41) points (95% CI [-0.73,-0.09], $p=.012$), corresponding to a standardized difference of ($d = -0.49$).

Thus, high-aesthetic videos were perceived as less realistic than low-aesthetic videos, producing an unintended realism difference of moderate magnitude. This comparison does not establish that perceived artificiality causes the substantive treatment effects reported in Study 3. It shows that assignment to the aesthetic treatment changed perceived realism in addition to the intended differences in attire, setting, framing, and production style. The aesthetic-proximity treatment should therefore be interpreted as the effect of the complete visual treatment bundle rather than as a perfectly isolated manipulation of informal presentation.

D.6.2 Relational-Signaling Contrast

The corresponding difference across the relational-signaling manipulation was considerably smaller. Respondents assigned to a high-relational stimulus reported an average realism rating of 1.98 ($SD=0.82$, $n = 53$), compared with 1.91 among respondents assigned to a low-relational stimulus ($SD=0.89$, $n = 55$). The estimated difference was 0.07 points (95% CI [-0.25,0.40], $p = .662$), with a standardized difference of ($d=0.08$).

The point estimate therefore provides no evidence that adding relational signaling systematically reduced perceived realism. The estimated difference is close to zero and is slightly positive rather than negative. At the same time, the confidence interval is sufficiently wide that the pilot does not establish tight equivalence between the relational conditions. The pilot reveals no systematic directional realism penalty associated with relational signaling. Table A16 shows the perceived realism scores across experimental treatment dimensions.

Table A16: Perceived Realism Across Experimental Treatment Conditions

Treatment Contrast	Low Condition	High Condition	Difference (High – Low)	95% CI	Cohen’s d	p
Aesthetic proximity	$M = 2.148$, $SD = 0.899$, $n = 54$	$M = 1.741$, $SD = 0.757$, $n = 54$	-0.407	[-0.725, -0.090]	-0.490	.012
Relational signaling	$M = 1.909$, $SD = 0.888$, $n = 55$	$M = 1.981$, $SD = 0.820$, $n = 53$	0.072	[-0.254, 0.398]	0.084	.662

Note: Perceived realism ranges from 1 (“definitely AI-generated”) to 4 (“definitely real”). Positive differences indicate that the high-treatment condition was perceived as more realistic than the corresponding low-treatment condition. Each respondent evaluated one randomly assigned experimental stimulus.

D.6.3 Realism Across the Four Factorial Conditions

Because the marginal contrasts above average across the other treatment dimension, I also examine perceived realism separately for each of the four factorial conditions. Table A17 reports the mean realism rating for the traditional-broadcast condition (low aesthetic, low relational), aesthetic-only condition (high aesthetic, low relational), relational-only

condition (low aesthetic, high relational), and full influencer condition (high aesthetic, high relational).

Table A17: Pairwise Comparisons of Perceived Realism Across Study 3 Experimental Conditions

Comparison	M_1 (SD)	M_2 (SD)	Difference	95% CI	t (df)	p	Cohen's d
Creator-High ($n = 26$) vs. Creator-Low ($n = 28$)	1.808 (0.749)	1.679 (0.772)	0.129	[-0.287, 0.545]	0.623 (51.9)	.536	0.170
Creator-High ($n = 26$) vs. Traditional-High ($n = 27$)	1.808 (0.749)	2.148 (0.864)	-0.340	[-0.786, 0.105]	-1.534 (50.5)	.131	-0.420
Creator-High ($n = 26$) vs. Traditional-Low ($n = 27$)	1.808 (0.749)	2.148 (0.949)	-0.340	[-0.811, 0.131]	-1.452 (49.1)	.153	-0.397
Creator-Low ($n = 28$) vs. Traditional-High ($n = 27$)	1.679 (0.772)	2.148 (0.864)	-0.470	[-0.914, -0.026]	-2.122 (51.9)	.039	-0.574
Creator-Low ($n = 28$) vs. Traditional-Low ($n = 27$)	1.679 (0.772)	2.148 (0.949)	-0.470	[-0.939, -0.000]	-2.009 (50.1)	.050	-0.544
Traditional-High ($n = 27$) vs. Traditional-Low ($n = 27$)	2.148 (0.864)	2.148 (0.949)	0.000	[-0.496, 0.496]	0.000 (51.6)	1.000	0.000

Note: Entries report pairwise Welch two-sample t -tests comparing perceived realism across the four Study 3 experimental conditions. Creator-High denotes the high-aesthetic, high-relational condition; Creator-Low denotes the high-aesthetic, low-relational condition; Traditional-High denotes the low-aesthetic, high-relational condition; and Traditional-Low denotes the low-aesthetic, low-relational condition. M_1 and M_2 correspond to the first and second conditions listed in each comparison, with standard deviations in parentheses. Differences are calculated as the first condition minus the second. Confidence intervals are 95% confidence intervals for the mean difference.

E Survey Instruments, Measures, and Preregistration

E.1 Core Outcomes and Scale Construction

Studies 2 and 3 use the same four core outcomes: parasocial closeness, perceived competence, candidate favorability, and candidate support. Parasocial closeness and perceived competence are measured using four-item indices, while favorability and support are measured with single items. The parasocial-closeness items adapt measures from Cohen and Holbert (2021) to capture respondents’ immediate sense of social and interpersonal connection to the candidate following a single exposure. Because the experiments present fictional candidates whom respondents encounter only within the survey, the index measures immediate parasocial closeness rather than an enduring parasocial relationship. The competence items capture judgments about the candidate’s experience, seriousness, professionalism, and electoral competitiveness.

Table A18 reproduces the exact wording of the core outcome measures. All measures use five-point response scales. Favorability ranges from “very unfavorable” to “very favorable”. The remaining measures use agreement scales ranging from “strongly disagree” to “strongly agree.” Higher values are coded to represent greater parasocial closeness, perceived competence, favorability, or candidate support.

For the two multi-item outcomes, I first construct the composite measure from the raw survey responses and then standardize the completed index. The parasocial-closeness index is the sum of the four items measuring group membership, interpersonal comfort, naturalness, and authenticity. The perceived-competence index is similarly the sum of the four items measuring political experience, seriousness, professionalism, and electoral competitiveness. Because each constituent item ranges from 1 to 5, each raw four-item index ranges from 4 to 20, with higher scores indicating greater parasocial closeness or perceived competence.

In the primary analyses, I standardize each completed composite index within study to have mean zero and standard deviation one. Favorability and candidate support are standardized directly from their original five-point measures using the same procedure. Treatment effects can therefore be interpreted in standard-deviation units and compared across outcomes despite differences in their original scales.

E.2 Manipulation Checks

Both experiments included manipulation checks to assess whether respondents perceived the intended differences in communication style. These measures are used as tests of treatment fidelity rather than as substantive outcomes. Study 2 includes a check for the aesthetic-proximity manipulation, while Study 3 separately measures perceived aesthetic proximity and relational signaling.

E.2.1 Study 2

Study 2 tests whether the visual treatment altered respondents’ perceptions of the candidate’s presentation while holding the spoken message constant. After viewing each

Table A18: Core Outcome Measures in Studies 2 and 3

Construct	Survey Item	Response Scale
Parasocial Closeness	When I watch the video, I feel that I am part of the candidate's group.	1 = Strongly disagree; 2 = Somewhat disagree; 3 = Neither agree nor disagree; 4 = Somewhat agree; 5 = Strongly agree
	The candidate makes me feel comfortable, as if I am with a friend.	Same five-point agreement scale
	I see in the candidate a natural and real person.	Same five-point agreement scale
	The candidate seemed authentic rather than scripted.	Same five-point agreement scale
Perceived Competence	The candidate seemed politically experienced.	Same five-point agreement scale
	The candidate came across as serious.	Same five-point agreement scale
	The candidate behaved professionally in this video.	Same five-point agreement scale
	I believe this candidate would be competitive in an election.	Same five-point agreement scale
Favorability	How favorable is your overall impression of this candidate?	1 = Very unfavorable; 2 = Unfavorable; 3 = Neither unfavorable nor favorable; 4 = Favorable; 5 = Very favorable
Candidate Support	I would support this candidate through voting, donating, or volunteering.	Same five-point agreement scale

Note: Exact capitalization and response labels reproduce the survey instrument as administered.

candidate video, respondents were asked: *This video looked like something an influencer would post.* The same 5-point Likert scale labels are used from “Strongly disagree” to “Strongly agree,” with higher values indicating that the video was perceived as more influencer-like.

For analysis, the manipulation-check measure was standardized within the Study 2 analytic sample in the same manner as the other single-item outcomes. I estimate the effect of the high-aesthetic condition using the same specification as the primary Study 2 analyses. Because each respondent evaluated multiple videos, standard errors are clustered by respondent.

The manipulation produced the intended perceptual difference. Respondents rated videos in the high-aesthetic-proximity condition as substantially more influencer-like than the corresponding low-aesthetic videos ($\beta = 0.230$, $SE = 0.020$, $p < .001$). Thus, the visual changes in attire, setting, framing, and production style were perceptible to respondents and shifted evaluations in the intended direction.

E.2.2 Study 3

Study 3 separately checks both dimensions of the factorial manipulation. The same influencer-style manipulation check was adopted from Study 2 to serve as the aesthetic-proximity manipulation check in Study 3. The relational-signaling manipulation checks assess whether respondents perceived the candidate as communicating directly and relationally with the audience. The exact items were:

- *The candidate invited interaction.*
- *The candidate shared something personal.*

Responses were recorded in the same 5-point Likert scale ranging from “Strongly disagree” to “Strongly agree.” For each manipulation check, I estimate the same (2×2) factorial specification used for the substantive Study 3 outcomes.

E.3 Preregistrations and Deviations

Both survey experiments were preregistered here: https://osf.io/vjrw4/overview?view_only=0be885f4fd1047f1a684f12fc21bc4f3. More details on the deviations will be included in future versions of the draft.

E.4 Study 2 Survey Instrument

Study 2 was administered as a module within the Graduate Weidenbaum Center Survey Wave 1 in 2026 at Washington University in St. Louis. This section reproduces the experimental module in the order in which respondents encountered it. The module consisted of pre-treatment questions, instructions for the video task, four repeated video-evaluation blocks, pre-debriefing questions about the stimuli, and a debriefing that disclosed the use of synthetic media and provided information about identifying AI-generated video. Informed-consent form and general demographic questions were administered by the omnibus survey outside this module.

E.4.1 Pre-Treatment Questions

Respondents first reported their recent use of four social media platforms.

Q1. In the past month, how often have you used the following platforms?

- TikTok
- Instagram
- YouTube
- Facebook

For each platform, respondents selected one of five categories:

1. Never
2. Once or twice
3. About once a week
4. Several times a week
5. Everyday

Respondents were then randomly assigned to one of two issue domains and answered a corresponding policy-attitude question. The instrument therefore measured respondents' prior issue position before exposure to the candidate videos.

Respondents assigned to the reproductive-rights domain received:

Q2a. There has been some discussion about abortion during recent years. Which one of the opinions on this page best agrees with your view?

- By law, abortion should never be permitted.
- The law should permit abortion only in case of rape, incest, or when the woman's life is in danger.
- The law should permit abortion other than for rape/incest/danger to woman but only after the need for the abortion has been clearly established.
- By law, a woman should always be able to obtain an abortion as a matter of personal choice.
- Other

Respondents assigned to the gun-policy domain received:

Q2b. Do you think the federal government should make it more difficult for people to buy a gun than it is now, make it easier for people to buy a gun, or keep these rules about the same as they are now?

- More difficult
- Easier
- Keep these rules about the same
- Don't know

E.4.2 Video Instructions and Presentation

Respondents next received the following instructions:

You will now watch four short political videos, each about 20–30 seconds long.

Please watch each video carefully. After each video, you will be asked to answer questions about the video content and your impressions of the candidates.

There are no right or wrong answers. We are interested in your honest opinions.

The four videos were then presented one at a time in randomized order. After each video, respondents completed the same evaluation battery before proceeding to the next stimulus.

E.4.3 Post-Video Candidate Evaluations

After each video, respondents first evaluated their overall impression of the candidate.

Q3. How favorable is your overall impression of this candidate?

- Very unfavorable
- Unfavorable
- Neither unfavorable nor favorable
- Favorable
- Very favorable

Respondents then completed the perceived-competence battery.

Q4. Based on your impression of the candidate shown in the video, how much do you agree or disagree with the following statements?

- The candidate behaved professionally in this video.
- The candidate seemed politically experienced.
- The candidate came across as serious.
- The candidate would be competitive in an election.

Each statement used the following five-point response scale:

1. Strongly disagree
2. Somewhat disagree
3. Neither agree nor disagree
4. Somewhat agree
5. Strongly agree

The next two question blocks measured parasocial closeness, candidate support, and the aesthetic manipulation check.

Q5. Thinking about the video you just watched, please indicate how much you agree or disagree with each of the following statements.

- When I watch the video, I feel that I am part of the candidate’s group.
- The candidate makes me feel comfortable, as if I am with a friend.
- I see in the candidate a natural and real person.

Q6. Thinking about the video you just watched, please indicate how much you agree or disagree with each of the following statements.

- The candidate seemed authentic rather than scripted.
- I would support this candidate through voting, donating, or volunteering.
- This video looked like something an influencer would post.

All Q5 and Q6 items used the same five-point agreement scale from “strongly disagree” to “strongly agree.”

The first four items across Q5 and Q6 form the parasocial-closeness index described in Section 3. The support item is analyzed as a single downstream candidate-evaluation outcome. The final item “This video looked like something an influencer would post” serves as the Study 2 manipulation check for the aesthetic-proximity treatment rather than as a substantive outcome.

E.4.4 Pre-Debriefing Questions

After completing all four video evaluations, respondents answered three questions about the stimuli before being informed that the candidate videos were synthetically generated.

First, respondents answered a content-recall question:

Q7. What did the candidates talk about?

- Gun Control
- Reproductive Rights
- Immigration
- Tax Reform

Respondents then rated the overall realism of the videos:

Q8. Overall, how realistic did the videos seem to you?

- Very realistic
- Somewhat realistic
- Neither realistic nor unrealistic
- Somewhat unrealistic
- Unrealistic

Finally, respondents reported how they believed the videos had been produced:

Q9. Which of the following best describes how you think the videos were produced?

- Recorded by real people without AI assistance.
- Recorded by real people but digitally edited using AI tools.
- Fully generated by AI (e.g., synthetic person or voice).

- A mix of these.
- I am not sure.

These measures were administered only after respondents had completed evaluations of all four candidates, so respondents were not informed of the synthetic nature of the stimuli before the substantive outcome measures were recorded.

E.4.5 Debriefing and Synthetic-Media Information

At the conclusion of the experiment, respondents were explicitly informed that the videos had been created using artificial intelligence for research purposes and did not depict real political candidates:

The videos you watched were created using artificial intelligence (AI) and were designed specifically for research purposes. They do not depict real political candidates.

The debriefing then provided an eight-part tutorial on identifying AI-generated or manipulated video. Respondents were advised to consider:

1. **When a video was created.** Most realistic AI-generated videos became widely available relatively recently, so older videos are less likely to be fully AI-generated.
2. **The duration and continuity of camera shots.** AI-generated videos may rely on relatively short clips before cutting to a new angle, whereas long continuous shots may be more indicative of authentic footage.
3. **Background text.** Text on clothing, signs, posters, or subtitles may be blurry, misspelled, or malformed in AI-generated video.
4. **Faces, mouths, and teeth.** Synthetic video may contain blurry teeth, inconsistent mouth shapes, or unnatural movements around the lips and jaw.
5. **Visual continuity.** AI-generated videos may contain sudden changes in a person's appearance, disappearing or reappearing objects, or unexplained changes in the surrounding scene.
6. **Logical inconsistencies in the scene.** Synthetic videos may contain physically implausible actions or objects that appear, disappear, or move in ways inconsistent with normal cause and effect.
7. **Urgency and verification in potential scams.** Respondents were advised to independently verify urgent claims involving requests for money or apparent emergencies rather than relying exclusively on audiovisual content.
8. **Common AI-video tools.** The tutorial identified tools such as Sora, Kling, Runway, Luma, Pika, and Synthesia as examples of platforms capable of generating synthetic video.

The debriefing concluded by emphasizing that generative AI can have legitimate creative uses while also creating risks of misinformation and fraud, and encouraged respondents to evaluate online audiovisual content critically.

E.5 Study 3 Survey Instrument

Study 3 was fielded as a standalone survey on Prolific. This section reproduces the survey instrument in the order in which respondents encountered it. The survey consisted of informed consent, pre-treatment demographic and political questions, an audiovisual attention check, exposure to one randomly assigned experimental video, post-treatment candidate evaluations and manipulation checks, pre-debriefing questions about the video, and a debriefing that disclosed the use of synthetic media and provided information about identifying AI-generated video.

E.5.1 Consent Information and Introduction

Respondents first received the following consent information:

You are invited to participate in a research study conducted by researchers at Washington University in St. Louis.

The purpose of this study is to understand how people perceive political videos commonly seen on social media. If you choose to participate, you will watch several short videos and answer questions about your impressions of the candidates and the videos.

The candidates shown in the videos are hypothetical individuals running for election. Some of the content you view may contain AI-generated content for research purposes.

Participation is voluntary. The survey will take approximately 3–4 minutes to complete. You may skip any question you prefer not to answer and may stop participating at any time by closing your browser.

Videos contain political content. There are no known risks beyond those encountered in everyday online activities. The results may help researchers better understand online political communication.

Your responses will be kept confidential and analyzed only in aggregate form.

The consent page also provided contact information for the research team and the Washington University Human Research Protection Office. Respondents were then informed:

By clicking “Next,” you indicate that you are at least 18 years old and consent to participate in this study.

E.5.2 Pre-Treatment Questions

Respondents next answered a series of demographic, media-use, and political-orientation questions.

Q1. What is your age?

- 18–20
- 21–24

- 25–29
- 30–34
- 35–39
- 40–44
- 45–49
- 50–54
- 55–60
- 61–64
- 65+

Q2. What is your gender?

- Woman
- Man
- Non-binary
- Other
- Prefer not to say

Q3. Which of the following best describes your race or ethnicity?

- White, non-Hispanic
- Black, non-Hispanic
- Hispanic
- Asian or Native Hawaiian/other Pacific Islander, non-Hispanic
- Native American/Alaska Native, non-Hispanic
- Multiple races, non-Hispanic
- Other

Q4. What is the highest level of education you have completed?

- High school or less
- Some college
- College grad
- Postgrad

Q5. In the past month, how often have you used the following platforms?

- TikTok
- Instagram
- YouTube
- Facebook

For each platform, respondents selected one of five categories:

1. Never

2. Once or twice
3. About once a week
4. Several times a week
5. Everyday

Q6. Here is a seven-point scale on which the political views that people might hold are arranged from extremely liberal to extremely conservative. Where would you place yourself on this scale?

1. Extremely Liberal
2. Liberal
3. Slightly Liberal
4. Moderate
5. Slightly Conservative
6. Conservative
7. Extremely Conservative

Q7. Generally speaking, do you think of yourself as a Republican, a Democrat, an Independent, or what?

- Democrat
- Independent
- Republican
- Other

Respondents identifying as Democrats or Republicans received a follow-up party-strength question:

Q8. Would you call yourself a strong Republican/Democrat or a not very strong Republican/Democrat?

- Strong Republican/Democrat
- Weak Republican/Democrat
- Independent Republican/Democrat

E.5.3 Audiovisual Attention Check

Before viewing the experimental stimulus, respondents completed an audiovisual attention check designed to verify that they could play the embedded video and hear its audio.

Q9a. A short video will appear below. Please click play, turn on sound, and watch the video carefully before continuing. The number you hear corresponds to the shape you see. Please choose the correct combination in the next question.

Q9b. Please choose the correct combination based on the video you just watched.

- One – Square

- Two – Triangle
- Three – Circle
- Four – Triangle

The correct response was Three – Circle. Respondents who failed this attention check were automatically directed to the end page without going through the rest of the survey, and were excluded from the analytical sample.

E.5.4 Experimental Video Instructions and Assignment

Respondents next received the following instructions:

You will now watch one short political video about 50 seconds long.

Please watch each video carefully. After the video, you will be asked to answer questions about the video content and your impressions of the candidates.

There are no right or wrong answers. We are interested in your honest opinions.

Each respondent was then randomly assigned to one of the four Study 3 experimental videos:

1. Low aesthetic proximity / low relational signaling;
2. High aesthetic proximity / low relational signaling;
3. Low aesthetic proximity / high relational signaling; or
4. High aesthetic proximity / high relational signaling.

Only one experimental video was shown to each respondent.

E.5.5 Post-Video Candidate Evaluations

Immediately after viewing the experimental video, respondents evaluated the candidate on a series of outcomes.

Q10. How favorable is your overall impression of this candidate?

- Very unfavorable
- Unfavorable
- Neither unfavorable nor favorable
- Favorable
- Very favorable

Respondents next evaluated several candidate characteristics.

Q11. Based on your impression of the candidate shown in the video, how much do you agree or disagree with the following statements?

- The candidate seemed trustworthy.
- The candidate seemed politically experienced.
- The candidate came across as serious.

- The candidate behaved professionally in this video.

Each item used the following five-point response scale:

1. Strongly disagree
2. Somewhat disagree
3. Neither agree nor disagree
4. Somewhat agree
5. Strongly agree

Respondents then answered three additional candidate-evaluation items.

Q12. Based on the video you just watched, to what extent do you agree or disagree with the following statements?

- I would support this candidate through voting, donating, or volunteering.
- I would like to learn more about this candidate.
- I believe this candidate would be competitive in an election.

These items used the same five-point agreement scale.

The four items used to construct the primary perceived-competence index are political experience, seriousness, professionalism, and electoral competitiveness. The trustworthiness and willingness-to-learn-more items are additional Study 3 outcomes and are not included in the primary competence index.

E.5.6 Parasocial Closeness

Respondents next completed the four-item parasocial-closeness battery.

Q13. Thinking about the video you just watched, please indicate how much you agree or disagree with each of the following statements.

- When I watch the video, I feel that I am part of the candidate's group.
- The candidate makes me feel comfortable, as if I am with a friend.
- I see in the candidate a natural and real person.
- The candidate seemed authentic rather than scripted.

All four items used the same five-point agreement scale from “strongly disagree” to “strongly agree.” These items form the parasocial-closeness index described in Section B.

E.5.7 Manipulation Checks and Additional Perceptions

The final post-treatment battery assessed respondents' perceptions of the candidate's communication style and campaign presentation.

Q14. Thinking about the video you just watched, please indicate how much you agree or disagree with each of the following statements.

- This video looked like something an influencer would post.
- This video looked professional.
- The candidate seems to have a well-funded campaign.

- The candidate seemed responsive to what people have been saying.
- The candidate shared something personal.
- The candidate invited interaction.

All items used the same five-point agreement scale.

The first two items capture perceptions of the candidate's visual presentation and provide checks on the aesthetic-proximity manipulation. The final three items correspond to central components of the relational-signaling treatment: audience acknowledgment or responsiveness, personal self-disclosure, and invitation to interact. The perceived campaign-funding item measures an additional downstream inference about the candidate's campaign resources.

E.5.8 Pre-Debriefing Questions

After completing the candidate-evaluation questions, respondents answered several questions before the survey disclosed that the experimental video had been generated using AI.

Q15. What did the candidate talk about?

- Gun Control
- Reproductive Rights
- Immigration
- Cost of Living

Respondents then reported the salience of the policy issue used in Study 3:

Q16. How concerned are you about rising costs?

- Not at all concerned
- Slightly concerned
- Moderately concerned
- Very concerned
- Extremely concerned

Respondents next evaluated the overall realism of the experimental video:

Q17. Overall, how realistic did the videos seem to you?

- Very realistic
- Somewhat realistic
- Neither unrealistic nor realistic
- Somewhat unrealistic
- Unrealistic

Finally, respondents reported how they believed the video had been produced:

Q18. Which of the following best describes how you think the videos were produced?

- Recorded by real people without AI assistance.

- Recorded by real people but digitally edited using AI tools.
- Fully generated by AI (e.g., synthetic person or voice).
- A mix of these.
- I am not sure.

These questions were administered after all substantive candidate evaluations and manipulation checks, so respondents did not receive an explicit disclosure that the stimulus was AI-generated before the primary outcome measures were recorded.

E.5.9 Debriefing and Synthetic-Media Information

At the conclusion of the study, respondents were explicitly informed that the experimental video had been created using artificial intelligence:

Thank you for participating in this study!

The videos you watched were created using artificial intelligence (AI) and were designed specifically for research purposes. They do not depict real political candidates.

The debriefing then provided an eight-part tutorial on identifying AI-generated or manipulated videos. Respondents were advised to consider:

1. **When the video was created.** Realistic generative-video systems became widely available relatively recently, making substantially older videos less likely to have been fully AI-generated.
2. **The duration and continuity of camera shots.** AI-generated videos may rely on relatively short clips before cutting to another shot.
3. **Background text.** Text on clothing, signs, posters, and subtitles may appear blurry, misspelled, or malformed in generated footage.
4. **Faces, mouths, and teeth.** Synthetic videos may contain blurry teeth, inconsistent mouth shapes, or unnatural movement around the lips and jaw.
5. **Visual continuity.** People, objects, or backgrounds may unexpectedly change, disappear, or reappear across generated sequences.
6. **Logical inconsistencies.** Generated footage may contain physically implausible objects, movements, or sequences of events.
7. **Urgency and independent verification in potential scams.** Respondents were encouraged to verify urgent audiovisual claims through trusted people or institutions rather than relying on the apparent authenticity of a video or voice.
8. **Common AI-video tools.** The tutorial identified tools such as Sora, Kling, Runway, Luma, Pika, and Synthesia as examples of platforms capable of producing synthetic video.

The debriefing concluded by emphasizing that generative AI can be used creatively and responsibly but can also facilitate misinformation and fraud, and encouraged respondents to critically evaluate audiovisual content encountered online.

F Study 1: Full Results and Robustness

This appendix reports the complete results and supporting analyses for Study 1. The main text emphasizes the substantive patterns in candidate adoption and platform performance, while the sections below provide the regression estimates underlying the main figures and additional analyses assessing the robustness of those patterns. Because candidates choose both what to post and how to present themselves, all Study 1 engagement estimates are interpreted as descriptive associations rather than causal effects.

F.1 Full Regression Results for Platform Visibility and Engagement

Table A19 reports the full regression results underlying Figure 2 Panel (b) in the main text. The analysis uses the 10,723 Reels from 160 candidate accounts for which post-level engagement metadata could be recovered. The four outcomes are views, likes, comments, and reposts. Because the engagement measures are highly right-skewed and, with the exception of views, may contain zeros, each is transformed as $\log(1 + y)$.

For each outcome, I jointly regress engagement on the aesthetic-proximity and relational-signaling indices. Both indices range from 0 to 1, so each coefficient represents the association corresponding to a shift from none to all of the component cues in that dimension. All models include candidate-account fixed effects and standard errors clustered by account. The fixed effects absorb stable differences across accounts, including time-invariant differences in candidate prominence, party, district, and baseline audience size. The resulting estimates therefore compare differently styled videos posted by the same candidate.

The full estimates reinforce the distinction between the two dimensions reported in the main text. Aesthetic proximity is positively associated with three of the four platform outcomes. A full-range increase in the aesthetic-proximity index is associated with a 0.164 increase in logged views ($SE = 0.070$, $p = .018$), a 0.233 increase in logged likes ($SE = 0.054$, $p < .001$), and a 0.194 increase in logged comments ($SE = 0.041$, $p < .001$). Using $(100[\exp(\beta) - 1])$ to translate the log coefficients, these estimates correspond to approximately 17.8% more views, 26.3% more likes, and 21.4% more comments. The association with reposts is essentially zero ($\beta = -0.002$, $SE = 0.027$, $p = .953$).

Relational signaling produces a different pattern. The coefficient for views is negative and comparatively large ($\beta = -0.400$, $SE = 0.220$), corresponding to approximately 33.0% fewer views across the full range of the index, but the estimate falls just outside conventional levels of statistical significance ($p = .070$). The associations with likes ($\beta = -0.055$, $SE = 0.155$, $p = .723$) and comments ($\beta = 0.018$, $SE = 0.175$, $p = .920$) are small relative to their uncertainty. Reposts provide the only statistically significant positive association: a full-range increase in relational signaling is associated with a 0.086 increase in logged reposts ($SE = 0.043$, $p = .046$), equivalent to approximately 9.0% more reposts.

The contrast across the two dimensions is therefore not that every aesthetic cue receives a platform reward while every relational cue is penalized. Rather, aesthetic proximity shows a comparatively consistent positive association with the most visible measures of reach and engagement—views, likes, and comments, whereas the associations for relational signaling are weaker and less consistent across outcomes. This pattern complements

Table A19: Associations Between Influencer-Style Communication and Platform Visibility and Engagement

	Dependent Variable: $\log(1 + y)$			
	Views	Likes	Comments	Reposts
Aesthetic proximity	0.164*	0.233***	0.194***	-0.002
	(0.070)	(0.054)	(0.041)	(0.027)
Relational signaling	-0.400	-0.055	0.018	0.086*
	(0.220)	(0.155)	(0.175)	(0.043)
Account fixed effects	Yes	Yes	Yes	Yes
SE clustered by account	Yes	Yes	Yes	Yes
Observations	10,723	10,723	10,723	10,723
Candidate accounts	160	160	160	160
R^2	0.650	0.653	0.493	0.162

Note: Entries are OLS coefficients from models with candidate-account fixed effects. Standard errors clustered by candidate account are reported in parentheses. All dependent variables are transformed as $\log(1 + y)$. The aesthetic-proximity and relational-signaling indices range from 0 to 1, so each coefficient corresponds to a shift across the full range of the relevant index. The analytic sample contains $N = 10,723$ videos from 160 candidate accounts. These estimates describe within-candidate associations and should not be interpreted causally.

* $p < .05$, ** $p < .01$, *** $p < .001$.

the adoption results in the main text: the dimension candidates use most extensively is also the one most consistently associated with observable platform performance. Because these comparisons remain observational even with account fixed effects, they establish the structure of platform feedback candidates encounter rather than the causal effect of adopting a particular communication style.

F.2 Component-Level Associations

Decomposing the two indices into their seven component cues, shown in Table A20, shows that the aggregate patterns mask meaningful variation across specific practices. The positive association for aesthetic proximity is driven primarily by informal production, which is associated with more likes ($\beta = 0.160$, $p = .006$) and comments ($\beta = 0.162$, $p = .018$), while informal candidate presentation is not significantly associated with any engagement outcome. Relational cues are more heterogeneous: direct address is associated with more reposts ($\beta = 0.078$, $p < .001$) and mutual awareness with more comments ($\beta = 0.156$, $p = .016$), whereas invitations to interact are associated with fewer likes and comments and the identity-bridging cue with fewer views and likes. Thus, the weak aggregate returns to relational signaling do not reflect uniformly null effects across its components, but rather a mixture of positive, negative, and near-zero associations.

Table A20: Associations Between Individual Communication Cues and Platform Visibility and Engagement

	Dependent Variable: $\log(1 + y)$			
	Views	Likes	Comments	Reposts
<i>Aesthetic Proximity</i>				
Informal production	0.114 (0.089)	0.160** (0.058)	0.162* (0.069)	-0.031 (0.022)
Informal candidate presentation	0.036 (0.058)	0.058 (0.052)	0.016 (0.054)	0.026 (0.027)
<i>Relational Signaling</i>				
Direct address	-0.049 (0.082)	0.016 (0.063)	-0.020 (0.069)	0.078*** (0.022)
Audience acknowledgment	-0.031 (0.092)	0.059 (0.060)	0.156* (0.065)	-0.018 (0.019)
Self-disclosure	-0.030 (0.041)	0.051 (0.037)	0.106 (0.058)	-0.014 (0.034)
Interaction invitation	-0.146 (0.077)	-0.123* (0.055)	-0.188* (0.076)	-0.033 (0.025)
Identity bridging	-0.172* (0.078)	-0.088* (0.043)	-0.037 (0.044)	0.020 (0.024)
Account fixed effects	Yes	Yes	Yes	Yes
SE clustered by account	Yes	Yes	Yes	Yes
Observations	10,723	10,723	10,723	10,723
Candidate accounts	160	160	160	160
R^2	0.651	0.654	0.496	0.164

Note: Entries are OLS coefficients from models with candidate-account fixed effects. Standard errors clustered by candidate account are reported in parentheses. Each model includes the seven individual communication cues shown in the table. All dependent variables are transformed as $\log(1 + y)$. Cue indicators are coded as present (1) or absent (0), so coefficients represent the within-account association between the presence of a cue and the corresponding engagement outcome, conditional on the other cues included in the model. The analytic sample contains $N = 10,723$ videos from 160 candidate accounts. These estimates describe within-candidate associations and should not be interpreted causally.

* $p < .05$, ** $p < .01$, *** $p < .001$.

G Study 2: Full Results and Robustness

This appendix reports the complete treatment-effect results and supporting analyses for Study 2. The main text focuses on the composite outcomes and their substantive interpretation. The sections below provide the full regression estimates and constituent-item results.

G.1 Full Regression Results

Table A21 reports the full regression results underlying Figure 3 in the main text. Each model estimates the effect of assignment to the high-aesthetic-proximity cue bundle relative to the low-aesthetic condition. The respondent-video evaluation is the unit of analysis, yielding 7,808 observations from 1,952 respondents. Because each respondent evaluated four candidate videos, standard errors are clustered by respondent. The models include no additional covariates. All outcomes are standardized within Study 2, so coefficients are expressed in standard-deviation units.

Table A21: Effects of Aesthetic Proximity on Candidate Evaluations

	Parasocial Closeness	Perceived Competence	Favorability	Support	Influencer Style
Aesthetic Proximity Cue Bundle	-0.050*** (0.015) [$p < .001$]	-0.366*** (0.018) [$p < .001$]	-0.092*** (0.016) [$p < .001$]	-0.087*** (0.015) [$p < .001$]	0.230*** (0.020) [$p < .001$]
Observations	7,808	7,808	7,808	7,808	7,808
Respondents	1,952	1,952	1,952	1,952	1,952
R^2	0.001	0.033	0.002	0.002	0.013
SE clustered by respondent	Yes	Yes	Yes	Yes	Yes

Note: Entries are OLS coefficients comparing the high-aesthetic-proximity cue bundle with the low-aesthetic-proximity condition. Standard errors clustered by respondent are reported in parentheses. All outcomes are standardized within Study 2, so coefficients are expressed in standard-deviation units. The influencer-style measure serves as a manipulation check rather than a substantive outcome.

⁺ $p < .10$, * $p < .05$, ** $p < .01$, *** $p < .001$.

The manipulation check confirms that the treatment altered perceived communication style as intended. Assignment to the high-aesthetic condition increased perceived influencer style by 0.230 standard deviations ($SE = 0.020$, $p < .001$). This shift did not translate into greater parasocial closeness. Instead, the treatment produced a small decline in the closeness index ($\beta = -0.050$, $SE = 0.015$, $p < .001$).

The largest effect appears in perceived competence. High aesthetic proximity reduced the competence index by 0.366 standard deviations ($SE = 0.018$, $p < .001$), substantially larger than its effect on the other outcomes. The treatment also reduced favorability by 0.092 standard deviations ($SE = 0.016$, $p < .001$) and candidate support by 0.087 standard deviations ($SE = 0.015$, $p < .001$). The full regression results therefore reinforce the main-text pattern: making the candidate appear more influencer-like produces no relational benefit, while generating a pronounced competence penalty and smaller downstream reductions in candidate evaluation and support.

G.2 Component-Level Results

Figure A9 decomposes the two composite indices into their constituent items. The small negative effect on parasocial closeness is not uniform across its four components. High aesthetic proximity reduced respondents' sense that they were part of the candidate's group ($\beta = -0.092$, $SE = 0.015$, $p < .001$), as well as perceptions that the candidate appeared natural and real ($\beta = -0.055$, $SE = 0.017$, $p = .001$) and authentic rather than scripted ($\beta = -0.043$, $SE = 0.017$, $p = .011$). The treatment had essentially no effect on feeling comfortable with the candidate ($\beta = 0.016$, $SE = 0.016$, $p = .299$).




Figure A9: Effects of High- versus Low-Aesthetic-Proximity Cue Bundles on Candidate Evaluations

Note: Points report standardized OLS estimates of the average treatment effect of the high- versus low-aesthetic-proximity cue bundle. Error bars report 95% confidence intervals. Standard errors are clustered by respondent. The influencer-style measure serves as the manipulation check.

The competence result is considerably more uniform. High aesthetic proximity lowered all four constituent measures: perceived political experience ($\beta = -0.396$, $SE = 0.019$, $p < .001$), professionalism ($\beta = -0.343$, $SE = 0.019$, $p < .001$), electoral competitiveness ($\beta = -0.311$, $SE = 0.019$, $p < .001$), and seriousness ($\beta = -0.228$, $SE = 0.018$, $p < .001$). Thus, the decline in the composite competence index is not driven by a single item, and the aesthetic treatment consistently weakened multiple judgments associated with fitness for political office. For comparison, Figure A9 also displays the smaller reductions in favorability and support and the positive manipulation-check effect.

H Study 3: Full Results and Robustness

This appendix reports the complete results and supporting analyses for Study 3. The main text emphasizes the substantive patterns across the four experimental conditions, while Figure 4 displays the corresponding estimated condition means. The analyses below use the full analytic sample of 1,575 respondents and standardized outcomes, estimated with OLS and heteroskedasticity-robust standard errors. Treatments are coded 0/1, so the lower-order aesthetic-proximity and relational-signaling coefficients represent simple effects when the other treatment is held at its low level. The sections that follow report the full regression estimates, component-level results, and condition means and contrasts.

H.1 Full Study 3 Regression Results and Component-Level Results

Table A22 reports the full regression estimates for the Study 3 outcomes using the 0/1 treatment coding described in the main text. Figures A10 and A11 decompose the composite outcomes into their constituent items and report the corresponding estimates for the additional outcomes and manipulation checks. Table A23 then shows the corresponding regression estimates. Because these are the lower-order coefficients from the factorial models, Figure A10 reports the effect of aesthetic proximity when relational signaling is low, while Figure A11 reports the effect of relational signaling when aesthetic proximity is low.

The component-level results reinforce the competence pattern in the main analysis. As shown in Figure A10, aesthetic proximity lowers evaluations of political experience and professionalism in particular, while the estimates for seriousness and electoral competitiveness are smaller and less precise. The parasocial-closeness components generally move in a positive direction, with the clearest increase appearing for respondents' reported comfort with the candidate, but these shifts do not produce a statistically distinguishable increase in the overall closeness index.

Relational signaling produces an even broader competence penalty. Figure A11 shows negative effects across all four constituent competence items, including political experience, seriousness, professionalism, and electoral competitiveness. At the same time, the relational manipulation produces little consistent movement across the components of parasocial closeness. The manipulation-check outcomes confirm that the treatment strongly increased perceived self-disclosure and interaction invitation, indicating that respondents recognized the intended relational features even though those cues did not translate into greater overall closeness.

The additional outcomes point in the same general direction. Both treatments reduce perceptions that the video was professionally produced and that the campaign was well funded, with substantially larger reductions under the aesthetic-proximity treatment. Relational signaling also produces negative estimates for trustworthiness and willingness to learn more about the candidate. These outcomes are secondary to the preregistered composite measures and are therefore interpreted descriptively here.

Table A22: Conditional Effects of Aesthetic Proximity and Relational Signaling on Candidate Evaluations from Study 3

	Parasocial Closeness	Perceived Competence	Favorability	Support
Aesthetic Proximity (<i>When Relational Signaling Is Low</i>)	0.122	−0.251***	0.018	−0.022
	(0.068)	(0.064)	(0.066)	(0.068)
	$p = .074$	$p < .001$	$p = .780$	$p = .750$
Relational Signaling (<i>When Aesthetic Proximity Is Low</i>)	0.034	−0.363***	−0.129	−0.123
	(0.071)	(0.066)	(0.070)	(0.071)
	$p = .627$	$p < .001$	$p = .067$	$p = .082$
Aesthetic Proximity × Relational Signaling	−0.055	−0.028	−0.074	−0.039
	(0.101)	(0.098)	(0.100)	(0.101)
	$p = .587$	$p = .778$	$p = .461$	$p = .700$
Num. Obs.	1,575	1,575	1,575	1,575
R^2	0.002	0.053	0.007	0.006
Std. Errors	Robust	Robust	Robust	Robust

Note: Entries are OLS coefficients from the 2×2 factorial experiment. All outcomes are standardized, so coefficients are expressed in standard-deviation units. Heteroskedasticity-robust standard errors are shown in parentheses. Treatment indicators use 0/1 coding. Accordingly, the lower-order aesthetic-proximity coefficient represents the effect of high versus low aesthetic proximity when relational signaling is low, while the lower-order relational-signaling coefficient represents the effect of high versus low relational signaling when aesthetic proximity is low. The interaction coefficient indicates how the effect of one treatment changes across levels of the other. * $p < .05$, ** $p < .01$, and *** $p < .001$.




Figure A10: Conditional Effects of Aesthetic Proximity on Study 3 Outcomes. Points report OLS estimates comparing the high- and low-aesthetic-proximity conditions when relational signaling is held at its low level. Bars show 95% confidence intervals based on heteroskedasticity-robust standard errors. All outcomes are standardized within Study 3.




Figure A11: Conditional Effects of Relational Signaling on Study 3 Outcomes. Points report OLS estimates comparing the high- and low-relational-signaling conditions when aesthetic proximity is held at its low level. Bars show 95% confidence intervals based on heteroskedasticity-robust standard errors. All outcomes are standardized within Study 3.

Table A23: Conditional Effects of Aesthetic Proximity and Relational Signaling Across Study 3 Outcomes at Component-Level

	Parasocial Closeness	Perceived Competence	Part of Group	Comfortable	Natural Authentic	Politically Experienced	Serious	Professional	Electronically Competitive	Favorability	Support	Influencer Style	Responsive	Personal Self-Disclosure	Interaction Invitation	Trustworthy	Learn More	Professionally Produced Video	Well-Funded Campaign
Aesthetic Proximity	0.122 (0.068) $p = .074$	-0.251*** (0.064) $p < .001$	0.016 (0.068) $p = .816$	0.208** (0.068) $p = .002$	0.115 (0.069) $p = .098$	0.092 (0.070) $p = .188$	-0.306*** (0.068) $p < .001$	-0.113 (0.062) $p = .070$	-0.342*** (0.066) $p < .001$	-0.100 (0.067) $p = .136$	0.018 (0.066) $p = .780$	-0.022 (0.068) $p = .750$	0.539*** (0.070) $p < .001$	-0.011 (0.069) $p = .874$	0.123* (0.057) $p = .032$	0.077 (0.057) $p = .178$	0.019 (0.067) $p = .772$	-0.032 (0.068) $p = .643$	-0.752*** (0.064) $p < .001$
Relational Signaling	0.034 (0.071) $p = .627$	-0.363*** (0.066) $p < .001$	0.039 (0.071) $p = .583$	0.104 (0.070) $p = .137$	-0.053 (0.070) $p = .455$	0.041 (0.071) $p = .568$	-0.430*** (0.067) $p < .001$	-0.370*** (0.067) $p < .001$	-0.241*** (0.062) $p < .001$	-0.187** (0.069) $p = .007$	-0.129 (0.070) $p = .067$	-0.123 (0.071) $p = .082$	0.234** (0.072) $p = .001$	0.077 (0.068) $p = .263$	1.287*** (0.056) $p < .001$	1.431*** (0.051) $p < .001$	-0.114 (0.071) $p = .109$	-0.162* (0.069) $p = .020$	-0.272*** (0.062) $p < .001$
Aesthetic Proximity × Relational Signaling	-0.055 (0.101) $p = .587$	-0.028 (0.098) $p = .778$	-0.042 (0.101) $p = .674$	-0.101 (0.100) $p = .313$	-0.033 (0.101) $p = .742$	-0.020 (0.101) $p = .845$	-0.041 (0.097) $p = .675$	0.038 (0.099) $p = .704$	-0.020 (0.098) $p = .839$	-0.064 (0.100) $p = .522$	-0.074 (0.100) $p = .461$	-0.039 (0.101) $p = .700$	-0.052 (0.097) $p = .593$	-0.124 (0.101) $p = .217$	-0.304*** (0.083) $p < .001$	-0.075 (0.072) $p = .298$	-0.080 (0.101) $p = .426$	0.003 (0.101) $p = .973$	0.161 (0.094) $p = .088$
Num. Obs.	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575
R^2	0.002	0.053	0.000	0.008	0.004	0.002	0.078	0.033	0.017	0.007	0.006	0.075	0.002	0.328	0.487	0.006	0.007	0.124	0.168
Std. Errors	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust

Note: Entries are OLS coefficients from the Study 3 2×2 factorial experiment. All outcomes are standardized, so coefficients are expressed in standard-deviation units. Heteroskedasticity-robust standard errors are reported in parentheses. Treatment indicators use 0/1 coding. Accordingly, the lower-order aesthetic-proximity coefficient represents the effect of high versus low aesthetic proximity when relational signaling is low, while the lower-order relational-signaling coefficient represents the effect of high versus low relational signaling when aesthetic proximity is low. The interaction coefficient indicates how the effect of one treatment changes across levels of the other. The first two columns report the composite parasocial-closeness and perceived-competence indices; subsequent columns report component items and additional evaluative and manipulation-check outcomes. * $p < .05$, ** $p < .01$, and *** $p < .001$.

H.2 Average Marginal Effects

Tables A24 and A25 report the average effects of aesthetic proximity and relational signaling across both levels of the other treatment. For these models, the treatment indicators are centered at -0.5 and 0.5 , so the lower-order coefficients represent average marginal effects rather than the conditional effects reported in the main text. This reparameterization does not change the fitted values or interaction estimates, and it changes only the interpretation of the lower-order treatment coefficients. The first table reports the four primary outcomes, while the second extends the same specification to the component-level, additional evaluative, and manipulation-check outcomes.

Table A24: Average Effects of Aesthetic Proximity and Relational Signaling on Candidate Evaluations from Study 3

	Parasocial Closeness	Perceived Competence	Favorability	Support
Aesthetic Proximity (Averaged Across Relational Conditions)	0.095 (0.050) $p = .060$	-0.264^{***} (0.049) $p < .001$	-0.019 (0.050) $p = .711$	-0.041 (0.050) $p = .415$
Relational Signaling (Averaged Across Aesthetic Conditions)	0.007 (0.050) $p = .891$	-0.377^{***} (0.049) $p < .001$	-0.166^{**} (0.050) $p = .001$	-0.143^{**} (0.050) $p = .005$
Aesthetic Proximity $\times$ Relational Signaling	-0.055 (0.101) $p = .587$	-0.028 (0.098) $p = .778$	-0.074 (0.100) $p = .461$	-0.039 (0.101) $p = .700$
Num. Obs.	1,575	1,575	1,575	1,575
R^2	0.002	0.053	0.007	0.006
Std. Errors	Robust	Robust	Robust	Robust

Note: Entries are OLS coefficients from the Study 3 2×2 factorial experiment. All outcomes are standardized, so coefficients are expressed in standard-deviation units. Heteroskedasticity-robust standard errors are shown in parentheses. Treatment indicators are centered at -0.5 and 0.5 , so each lower-order coefficient represents the average effect of that treatment across the two levels of the other. The interaction coefficient indicates how the effect of one treatment changes across levels of the other. $*p < .05$, $**p < .01$, and $***p < .001$.

Table A25: Average Effects of Aesthetic Proximity and Relational Signaling Across Study 3 Outcomes at Component-Level

	Parasocial Closeness	Perceived Competence	Part of Group	Comfortable	Natural and Real	Authentic	Politically Experienced	Serious	Professional	Electorally Competitive	Favorability	Support	Influencer Style	Responsive	Personal Self-Disclosure	Interaction Invitation	Trustworthy	Learn More	Professionally Produced Video	Well-Funded Campaign
Aesthetic Proximity (Averaged Across Relational Conditions)	0.095 (0.050) $p = .060$	-0.264*** (0.049) $p < .001$	-0.005 (0.050) $p = .914$	0.158** (0.050) $p = .002$	0.098 (0.050) $p = .051$	0.082 (0.050) $p = .104$	-0.328*** (0.048) $p < .001$	-0.094 (0.050) $p = .059$	-0.352*** (0.049) $p < .001$	-0.132** (0.050) $p = .008$	-0.019 (0.050) $p = .711$	-0.041 (0.050) $p = .415$	0.504*** (0.049) $p < .001$	-0.073 (0.050) $p = .147$	-0.029 (0.041) $p = .481$	0.040 (0.036) $p = .271$	-0.021 (0.050) $p = .680$	-0.030 (0.050) $p = .553$	-0.671*** (0.047) $p < .001$	-0.789*** (0.046) $p < .001$
Relational Signaling (Averaged Across Aesthetic Conditions)	0.007 (0.050) $p = .891$	-0.377*** (0.049) $p < .001$	0.018 (0.050) $p = .726$	0.053 (0.050) $p = .292$	-0.069 (0.050) $p = .169$	0.031 (0.050) $p = .542$	-0.450*** (0.048) $p < .001$	-0.351*** (0.050) $p < .001$	-0.251*** (0.049) $p < .001$	-0.219*** (0.050) $p < .001$	-0.166** (0.050) $p = .001$	-0.143** (0.050) $p = .005$	0.208*** (0.049) $p < .001$	0.015 (0.050) $p = .773$	1.135*** (0.041) $p < .001$	1.394*** (0.036) $p < .001$	-0.154** (0.050) $p = .002$	-0.160** (0.050) $p = .001$	-0.192*** (0.047) $p < .001$	-0.217*** (0.046) $p < .001$
Aesthetic Proximity × Relational Signaling	-0.055 (0.101) $p = .587$	-0.028 (0.098) $p = .778$	-0.042 (0.101) $p = .674$	-0.101 (0.100) $p = .313$	-0.033 (0.101) $p = .742$	-0.020 (0.101) $p = .845$	-0.041 (0.097) $p = .675$	0.038 (0.099) $p = .704$	-0.020 (0.098) $p = .839$	-0.064 (0.100) $p = .522$	-0.074 (0.100) $p = .461$	-0.039 (0.101) $p = .700$	-0.052 (0.097) $p = .593$	-0.124 (0.101) $p = .217$	-0.304*** (0.083) $p < .001$	-0.075 (0.072) $p = .298$	-0.080 (0.101) $p = .426$	0.003 (0.101) $p = .973$	0.161 (0.094) $p = .088$	-0.027 (0.092) $p = .765$
Num. Obs.	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575	1,575
R ²	0.002	0.053	0.000	0.008	0.004	0.002	0.078	0.033	0.047	0.017	0.007	0.006	0.075	0.002	0.328	0.487	0.006	0.007	0.124	0.168
Std. Errors	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust

Note: Entries are OLS coefficients from the Study 3 2 × 2 factorial experiment. All outcomes are standardized, so coefficients are expressed in standard-deviation units. Heteroskedasticity-robust standard errors are reported in parentheses. Treatment indicators are centered at −0.5 and 0.5, so each lower-order coefficient represents the average effect of that treatment across the two levels of the other treatment. The interaction coefficient indicates how the effect of one treatment changes across levels of the other. The first two columns report the composite parasocial-closeness and perceived-competence indices; subsequent columns report component items and additional evaluative and manipulation-check outcomes. * $p < .05$, ** $p < .01$, and *** $p < .001$.

H.3 Estimated Condition Means and Experimental Contrasts

Table A26 reports estimated means for each of the four experimental conditions and selected contrasts among them. Panel A presents the estimated mean of each standardized outcome for the traditional-broadcast (low-aesthetic-low-relational), aesthetic-only, relational-only, and full-influencer conditions. Panel B reports contrasts that correspond to the simple treatment effects in the 0/1-coded regression model, the full-influencer versus traditional-broadcast comparison, and the effect of each treatment when the other treatment is high.

The condition means show little separation in parasocial closeness. Relative to the traditional-broadcast condition, neither the aesthetic-only condition ($\beta = 0.122$, $SE = 0.068$, $p = .074$), the relational-only condition ($\beta = 0.034$, $SE = 0.071$, $p = .627$), nor the full-influencer condition ($\beta = 0.102$, $SE = 0.070$, $p = .149$) produces a statistically distinguishable increase in closeness.

The pattern is substantially clearer for perceived competence. Relative to traditional broadcast, perceived competence is 0.251 standard deviations lower in the aesthetic-only condition ($SE = 0.064$, $p < .001$), 0.363 standard deviations lower in the relational-only condition ($SE = 0.066$, $p < .001$), and 0.642 standard deviations lower in the full-influencer condition ($SE = 0.068$, $p < .001$). The competence penalties also remain large when either treatment is introduced while the other is already high.

For the downstream outcomes, the full influencer package produces clearer costs than either lower-order contrast alone. Compared with traditional broadcast, the full-influencer condition reduces favorability by 0.184 standard deviations ($SE = 0.068$, $p = .007$) and support by 0.184 standard deviations ($SE = 0.070$, $p = .008$). Relational signaling also reduces favorability when aesthetic proximity is already high ($\beta = -0.203$, $SE = 0.072$, $p = .005$) and reduces support under the same condition ($\beta = -0.162$, $SE = 0.071$, $p = .023$). Taken together, these contrasts reinforce the main result: combining aesthetic proximity and relational signaling does not generate greater parasocial closeness, while the full influencer package produces substantial competence penalties and lower downstream candidate evaluations.

Table A26: Estimated Condition Means and Experimental Contrasts in Study 3

	Parasocial Closeness	Perceived Competence	Favorability	Support
<i>Panel A. Estimated Condition Means</i>				
Traditional Broadcast	−0.065 (0.047)	0.315 (0.042)	0.074 (0.044)	0.082 (0.047)
Aesthetic Only	0.057 (0.049)	0.064 (0.048)	0.092 (0.049)	0.061 (0.049)
Relational Only	−0.031 (0.052)	−0.048 (0.051)	−0.055 (0.055)	−0.041 (0.053)
Full Influencer	0.037 (0.052)	−0.327 (0.054)	−0.110 (0.052)	−0.101 (0.052)
<i>Panel B. Experimental Contrasts</i>				
Aesthetic Only – Traditional	0.122 (0.068) $p = .074$	−0.251*** (0.064) $p < .001$	0.018 (0.066) $p = .780$	−0.022 (0.068) $p = .750$
Relational Only – Traditional	0.034 (0.071) $p = .627$	−0.363*** (0.066) $p < .001$	−0.129 (0.070) $p = .067$	−0.123 (0.071) $p = .082$
Full Influencer – Traditional	0.102 (0.070) $p = .149$	−0.642*** (0.068) $p < .001$	−0.184** (0.068) $p = .007$	−0.184** (0.070) $p = .008$
Aesthetic Effect When Relational = High	0.068 (0.074) $p = .361$	−0.278*** (0.074) $p < .001$	−0.056 (0.076) $p = .462$	−0.060 (0.074) $p = .416$
Relational Effect When Aesthetic = High	−0.020 (0.072) $p = .776$	−0.391*** (0.072) $p < .001$	−0.203** (0.072) $p = .005$	−0.162* (0.071) $p = .023$

Note: Panel A reports estimated means for the four cells of the Study 3 2×2 factorial experiment. Standard errors are reported in parentheses. Panel B reports selected linear contrasts among experimental conditions. All outcomes are standardized within Study 3. Estimates are based on OLS models with heteroskedasticity-robust standard errors. Traditional Broadcast denotes the low-aesthetic, low-relational condition; Aesthetic Only denotes high aesthetic proximity and low relational signaling; Relational Only denotes low aesthetic proximity and high relational signaling; and Full Influencer denotes high levels of both dimensions. * $p < .05$, ** $p < .01$, and *** $p < .001$.