Exposure to Negative Advertising Contributes to Affective Polarization*

Shun Yamaya,† Brian Wu,‡ Shanto Iyengar,§
David Brady,¶ and Douglas Rivers‖ [1]

[1]Department of Political Science, Stanford University, Stanford, CA

July 2026

Abstract

Negative advertising in American election campaigns and the level of partisan hostility toward political opponents have both increased substantially over the past several decades. We argue that these parallel trends are causally linked. We use a blend of experiments and survey time series to estimate the effects of exposure to negative campaigns. The experiments, which manipulate advertising tone while holding other features of the treatment ads constant, were administered during the 1992, 1996, and 2012 campaigns. Our survey analyses rest on a within-election-cycle difference-in-differences design along media market borders, implemented for the 2000, 2004, 2008, 2020, and 2024 campaigns. The survey and experimental results both show that exposure to negative—but not positive—ads significantly increased affective polarization in the 1990s and 2000s. The magnitude of this effect is substantial: for example, a typical week of negative advertising during the height of the 2000 election accounts for roughly one-tenth of the affective polarization observed that year. Exposure to negative ads drove affective polarization by primarily heightening animus toward the out-party. A dynamic specification suggests that campaign effects persist beyond the election season. We further show that the polarizing effect of advertising has diminished in more recent election cycles, consistent with increased partisan sorting or the possibility that measures of out-party animus have approached a ceiling. These multi-method findings showcase a contextual mechanism of polarization and suggest that efforts to reduce polarization should focus on incentives shaping elite behavior, rather than the predispositions of partisans. (242 words)

*We thank Justin Grimmer and Yiqing Xu for methodological guidance. Financial support for this research was provided by the Institute for Humane Studies.

†syamaya@stanford.edu

‡wubrian@stanford.edu

§siyengar@stanford.edu

¶dbrady@stanford.edu

‖rivers@stanford.edu

1 Introduction

There is now ample evidence documenting the phenomenon of affective polarization. Democrats and Republicans habitually treat their opponents as a stigmatized out group. Partisan animus manifests itself in the form of unfavorable evaluations of the opposition party and its leaders, aversion to social contact with partisans from the other side, and in indicators of implicit or sub-conscious prejudice. Most tellingly, partisans subject their opponents to discrimination in a variety of behavioral settings (Iyengar and Westwood 2015; Gift and Gift 2015; Sinclair, Nilsson, and Agerström 2023).

Intensified hostility across the political divide is a relatively recent development. The American National Election Studies feeling thermometer time series shows that between the mid-1970s and 1990, the average rating of the out-party hovered around 50 (the midpoint), suggesting that partisans felt either neutral or ambivalent toward opponents. The thermometer scores began a slow but steady decline in the 1990s, with the average rating of the out-party falling to 40 in 2000. The declining trend then accelerated significantly, with the average out-party rating dropping from 40 to 20 by 2024. Since 2012, the modal thermometer score for the out-party has plummeted to zero.

The onset of increased partisan animus closely tracks another well-documented trend in American politics corresponding to the advent of large-scale televised negative campaigning. At the dawn of the broadcast era, as candidates and their consultants discovered the power of television advertising, campaign ads tended to mimic product ads, with messages extolling the virtues and strengths of the sponsoring candidate. Positive appeals made up around two-thirds of presidential ads that aired between 1960 and 1984 (Geer 2012).[1] The turning point in the trajectory of advertising tone occurred in 1988 when the infamous "Willie Horton" and "Boston Harbor" attacks directed at Democrat Michael Dukakis enabled George Bush to neutralize Dukakis' early lead in the polls, leading consultants to conclude that negative

[1] This tabulation is based on a classification of presidential ads collected by the advertising archive at the University of Oklahoma.

appeals were especially persuasive (for discussions of the 1988 advertising campaign, see Jamieson 1993; Mendelberg 1997).

The proportion of negative ads in presidential campaigns nearly doubled between 1992 and 2008 (Geer 2012) exceeding 60 percent. For the post-2008 period, for which there is more reliable and precise data on advertising buys and airings, the move toward negativity has continued – attack and contrast ads took up more than 70 percent of the advertising budget in congressional and gubernatorial elections held in 2012 and 2014 (Wesleyan Media Project 2014). In 2024, based on a database of all televised ads that aired before Election Day, the level of negativity reached 82 percent.[2] In the case of the Trump campaign, virtually every ad attacked Kamala Harris (see Wesleyan Media Project 2024). Clearly, "going negative" is a dominant strategy in modern campaigns.

Despite the rapid proliferation of social media, television remains the primary news source for Americans. Studies based on behavioral data show that television news consumption dominates online sources by a factor of five (Allen et al. 2020; Muise et al. 2022). Although the Internet has fundamentally changed how the world communicates in the 21st century, television remains the dominant platform for campaign communication. In short, the audience for televised campaign ads is vast.

The correspondence between the onset of increased out-party animus and the widespread adoption of negative campaigning can hardly be coincidental. In this paper, we argue that there are strong theoretical reasons to expect that amplified negativity in campaign discourse is causal to increased partisan animus. Decades of research into the origins of public opinion demonstrates that partisan voters take on the views of their leaders (Zaller 1992). When Kamala Harris exclaims that Donald Trump "is unstable and unfit for office" and Trump counters with "she is a radical left lunatic," partisans naturally conclude that their opponents are scurrilous. Although name calling and personal abuse have been fixtures of American politics dating back to the early days of the Republic, these modern rhetorical signals travel

[2]In Appendix Figure A2, we visualize some of these trends in recent election cycles.

more rapidly today and reach a much larger audience. Despite well-intentioned efforts to make campaign advertising more civil (e.g., by requiring candidates to appear in their ads and acknowledge their sponsorship of the message), the trajectory of advertising tone remains almost exclusively negative.

As prima facie evidence of the impact of exposure to negative campaigning on affective polarization, we first reanalyze evidence from a series of experiments conducted during the 1992 and 1996 election campaigns (Ansolabehere and Iyengar 1995). These studies manipulated participant exposure to negative and positive ads run by the presidential candidates, as well as ads from two California U.S. Senate races held in 1992. We also include a more recent experiment conducted during the 2012 presidential campaign, in which participants viewed positive or negative ads from the Obama and Romney campaigns. Although these experiments were administered on relatively small samples, the results show that participants exposed to negative ads expressed significantly higher levels of out-party animus than those exposed to positive ads.

We then test the causal impact of exposure to negative campaigning on population-level affective polarization using a series of large-scale cross-sectional surveys with representative samples that track the level of partisan animus over five election cycles. We extend previous designs used in the literature estimating the effect of campaign advertisements (Huber and Arceneaux 2007; Spenkuch and Toniatti 2018; Sides, Vavreck, and Warshaw 2022) by isolating within-county and within-election-cycle variation. We show that residents of media markets who are exposed to a greater barrage of negative messages tend to be more polarized than residents of markets less exposed to the ad campaign. The effect is substantively important: the average 7-day running volume of negative advertising during the last month of the general election accounts for approximately 10.4%, 12.3%, and 11.4% of changes in affective polarization in the 2000, 2004, and 2008 election cycles, respectively. However, this effect has substantially diminished by 2020 and 2024, a temporal pattern also detected in the experimental results, which we attribute to increased partisan sorting and ceiling effects in the indicators of partisan animus.

In what follows, we first survey the literature on the use of negative campaigning and its effects on voter attitudes. Next, we describe our experimental and observational research designs, sources of data, and analysis strategy. We then present our findings with an accompanying discussion of their implications for the study of affective polarization and American politics more generally.

2 Causes and Effects of Negative Campaigning

A large body of scholarship documents the growing use of negative appeals in American politics. Numerous studies examine both the underlying causes of this trend and its attitudinal and behavioral consequences for voters. We review each of these literatures in turn.

Proposed Explanations

Scholars have advanced several explanations for the prevalence of negative messaging in the political arena. One account is psychological: experimental research demonstrates that negative stimuli leave stronger imprints in human information processing. Individuals are more likely to attend to and retain negative, rather than positive, information (Baumeister et al. 2001; Bradley, Angelini, and Lee 2007; Soroka, Fournier, and Nir 2019). In addition, negative traits are weighted more heavily than positive traits in impression formation (Fiske 1980; Klein 1991). For campaigns seeking to break through a highly cluttered information environment, harsh attacks on opponents are therefore more likely to be noticed and remembered by target audiences.

A second factor underlying the rise of negative campaigning is candidates' need to attract news coverage. The audience for campaign communications includes not only voters but also journalists. In high-profile races, "free media" in the form of news coverage is at least as important—if not more so—than paid advertising as a channel of communication (see Iyengar

2022). Consistent with the journalistic maxim "if it bleeds, it leads," editorial judgments of newsworthiness exhibit a pronounced negativity bias (for evidence in the domain of economic news, see Soroka 2006). Accordingly, messages emphasizing conflict and controversy tend to generate more press coverage than those highlighting civility and cooperation on the campaign trail.

In fact, the available evidence indicates that negative ads are substantially more likely to receive press coverage than positive ads. During the 1988 campaign, for example, attacks launched by groups affiliated with George H. W. Bush against Michael Dukakis received front-page coverage in major newspapers. Coverage of presidential campaign advertising in the *New York Times* and the *Washington Post* reached approximately 200 stories in 1988——more than double the number in 1984 (Geer 2012).

By 2004, the newsworthiness of negative ads had increased by an additional 25 percent, yielding roughly 250 stories. Trends in broadcast news were similar: by 2008, network coverage of campaign advertising focused disproportionately on attack ads (Geer 2012). In short, negative ads are aired "not so much with an eye to influencing voters directly, but with the hope of altering the news media's narrative in the campaign" (p. 9).

A third explanation for the emergence of negative campaigning is strategic. Ever since the 1988 campaign, when the Dukakis team simply ignored the barrage of attacks that came their way, it has become ingrained in the consultants' playbook that the attacked candidate must immediately rebut the attack, typically in the form of a counter-attack. As a result, most ad campaigns converge to a negative equilibrium (for efforts to model this spiral of negativity, see Skaperdas and Grofman 1995; Ansolabehere and Iyengar 1995).

Consumer demand constitutes a final contributing factor. In an era of heightened polarization, partisans who strongly dislike their opponents are more likely to reward candidates who engage in attack messaging. Moreover, one consequence of intensified partisan animus is increased electoral loyalty. In the past three election cycles, rates of partisan defection have fallen into the single digits, with defection among strong partisans remaining especially rare (see Jacobson 2023). Consistent with this pattern, a large-scale analysis of partisan defection in the 2020 election found that a substantial majority of Republicans remained loyal—voting for Trump despite holding opposing positions on several salient issues (Tyler, Iyengar, and Wilkins, n.d.).

As the rarity of defection illustrates, in the current polarized environment, it is very difficult to persuade partisans (see Kalla and Broockman 2018 for evidence of null effects in contemporary campaigns). Campaigns therefore may be more inclined to pursue changes in turnout, either by mobilizing their base, or demobilizing their opponents. Raising doubts about a candidate's fitness for office may do little to dissuade supporters of that candidate but might be sufficient to discourage some of them from voting. It is this potential for campaign demobilization to which we turn next.

Effects of Negative Campaigns

In a series of carefully controlled experiments carried out during the 1990 and 1992 election cycles, Ansolabehere and Iyengar (1995) showed that exposure to a single negative ad was sufficient to depress intention to vote by around 4.5 percentage points among California voters. Participants exposed to the negative ad treatment also expressed more cynicism about the electoral process and their own ability to make a difference (Ansolabehere and Iyengar 1995, 103–104).

While the Ansolabehere-Iyengar experiments achieved strong control in the form of content equivalence across the positive and negative ad conditions, the experimental ads necessarily differed from actual ads in several respects. For instance, the treatment ads lacked background music and sound bites from the sponsoring candidate.

Studies that have investigated the effects of actual ads on voter engagement typically find null effects (for a review, see Lau, Sigelman, and Rovner 2007). The question of whether negative campaigns demobilize voters thus remains open. Moreover, given what we know about the level of out-party animus today, it seems quite plausible that negative messages will actually work to mobilize rather than demobilize partisans (for evidence that out-party animus is positively correlated with turnout, see Ahn and Mutz 2023).

Another set of studies has pursued the "backlash" hypothesis on the grounds that civility is normatively appropriate, leading voters to sanction candidates who attack. Tests of this hypothesis have yielded mixed findings. A handful of studies show some traces of backlash, but this evidence comes from overseas elections featuring multi-party systems that make it possible for a third-party candidate to gain support, when the two major parties engage in attacks and counter-attacks (for instance, Roese and Sande 1993; Walter and Eijk 2019; Galasso, Nannicini, and Nunnari 2023). Studies with American voters, however, generally find only weak traces of backlash (see Lau et al. 1999; Fridkin and Kenney 2004; Jasperson and Fan 2002).

Quite surprisingly, only a handful of studies have investigated the causal impact of negative campaigns on partisan animus. In the first study to document the phenomenon of affective polarization, the authors tracked changes in polarization over the course of the 2004 and 2008 campaigns. In both cases, they found that exposure to the campaign increased out-party animus, that voters living in states with higher levels of campaign advertising were more polarized, and that state-level exposure to negative ads was a contributory factor (Iyengar, Sood, and Lelkes 2012).

On the basis of time series data from 2000, 2004, and 2008, Sood and Iyengar (2016) reported similar results. Using both feeling thermometers and candidate trait ratings to measure affective polarization, the authors documented substantial increases in out-party animus over the course of the campaign. In both 2004 and 2008 (but not 2000), they found that voters in battleground states, who received significantly larger doses of negativity, registered significantly greater over-time increases in polarization. We replicate these analyses here as part of our observational survey databases.

In a cross-national extension, Martin and Nai (2024) found that in multi-party elections, supporters of parties that engage in or are targets of negative campaigning express higher levels of partisan animus. Their results held up across the introduction of multiple covariates including voter ideology and strength of partisanship.

Of course, the time series evidence is correlational, making it difficult to draw firm conclusions about the causal impact of campaign negativity. Small-scale experiments that manipulate the tone of campaign ads tend to bolster the observational evidence. In an experiment administered during the 2012 campaign (Lau et al. 2017) the researchers found increased polarization among participants assigned to view a negative ad, but the effect was conditional on participants' access to partisan news sources. The authors theorized that exposure to attack ads on polarization was indirect; supporters of the attacked candidate would be motivated to seek out supportive information refuting the attack from partisan providers and that this conjunction of exposure to campaign attacks and access to supportive partisan information would fuel out-party animus. In support of this expectation, they detected a significant interaction between access to partisan news and exposure to negative advertising (Lau et al. 2017).[3]

The most recent experimental study, based on a large-scale field experiment (Weiss, Green, and Willer 2025), tested the depolarizing effects of public service ads (sponsored by the National Governor's Association) that called for civility in political discourse. The treatment ads –drawn from five different states – featured prominent politicians from both parties calling for civility in campaigns. The experiment yielded mixed results. Although participants in the treated condition reported greater willingness to engage in bipartisan behaviors, their level of animus toward the opposing party remained unchanged. While this study addresses the potential for positive appeals to depolarize, it sheds no light on the polarizing effects of negative messages.

[3]We limit our attention in this literature review to negativity in advertising campaigns. There is a parallel literature addressing the effects of negative news coverage. Consistent with the demobilization hypothesis, this literature generally shows that exposure to incivility and conflict in news reports fosters cynicism and distrust toward politicians and government (see Patterson 1994; Cappella and Jamieson 1997; Mutz and Reeves 2005; Valentino, Beckmann, and Buhr 2001).

3 Experimental Designs

The "gold standard" design for estimating the effects of negative campaigning on out-party animus is a controlled experiment in which the only difference between treatment and control conditions is the tone of the message. This is the design implemented by Ansolabehere and Iyengar (1995) in a series of experiments that examined the effects of campaign negativity on voter turnout.

We reanalyze six in-person experiments conducted in Southern California during the 1992 and 1996 campaigns, and an additional online study conducted during the 2012 campaign. The 1992 studies covered two California U.S. Senate races and the presidential contest, the 1996 and 2012 experiments focused exclusively on the presidential race. Our reanalysis is limited to these experiments because they included the standard 101-point feeling thermometer rating of both candidates.

The 1992 Senate studies featured two races. The Democratic nominees were Barbara Boxer and Dianne Feinstein, running against Bruce Herschensohn (a conservative media commentator) and John Seymour (a state legislator), respectively. Feinstein won by a wide margin; Boxer narrowly defeated Herschensohn. The single-ad Senate treatments addressed crime, unemployment, and sexual harassment; the two-ad treatments focused on trustworthiness and economic stewardship.

In the 1992 presidential race, incumbent George H. W. Bush faced Bill Clinton. The single-ad treatment addressed sexual harassment; the two-ad treatments focused on the candidates' trustworthiness and economic policies. In the 1996 two ad study, the candidates were incumbent President Bill Clinton and Senator Bob Dole. The ads focused on the candidates' performance in office, differences in their tax and immigration policies, and questions of character. Finally, the ads used in the 2012 study dealt with the state of the economy and the candidates' (Obama's or Romney's) economic credentials as well as ads bearing on the candidates' personal backgrounds, trustworthiness, and capacity to lead.

The *Going Negative* experiments used two designs, distinguished by whether the treatment consisted of a single campaign ad or a pair of ads embedded into a videotape recording of a local newscast. The 1996 and 2012 studies exposed participants to a set of standalone campaign and product ads. In all designs, the researchers manipulated the number of negative ads to which participants were exposed (ranging from zero to three).

In the case of the 1992 single-ad studies, participants viewed nearly identical thirty-second ads in which the voiceover was altered to create either positive or negative messages. In the 1992 presidential study, for instance, the positive version of the Clinton ad read: "Anita Hill testifies before the Senate Judiciary Committee. [Ten-second clip.] The senators reject Hill's testimony. Bill Clinton believes that sexual harassment is a real problem. It is time to stand up for the rights of women." The negative version substituted George Bush as the target: "Anita Hill testifies... George Bush doesn't believe that sexual harassment is a real problem. It is time to stand up...." The close informational equivalence between experimental conditions makes the single-ad design a stringent test of the polarizing effects of negative advertising.

The 1992 two-ad studies embedded actual ads run by the candidates into the local newscast, meaning the video and soundtrack naturally differed across the positive and negative treatments. However, the issue area was held constant across conditions — in the 1992 presidential two-ad study, for example, both ads addressed either the candidates' trustworthiness or their ability to improve economic conditions.

The 1992 experiments were administered at two locations in the Greater Los Angeles area, chosen to represent areas with heavy concentrations of Democratic and Republican voters, respectively. Participants were recruited by telephone, flyers posted at shopping malls and parking lots, and announcements in employee newsletters. They were offered 15 dollars to participate in a study of "selective perception" of local news, a plausible cover story designed to conceal the study's true purpose. Although the sample was non-random, it closely matched the Los Angeles population on major demographic variables (for details, see Ansolabehere and Iyengar (1995), p. 29).

The facilities were identical at both sites: an office suite with two viewing rooms and a separate room for completing questionnaires. On arrival, participants completed a brief pretest survey before watching a fifteen-minute videotape of a recent local newscast. The treatment ad appeared in the middle slot of a three-ad commercial break midway through the newscast. Afterward, participants completed a post-test survey measuring their attitudes toward the candidates. They were then debriefed (told about the true purpose of the study) and paid.

The 1996 experiment was administered at multiple sites in the Greater Los Angeles area. The researchers' objective was to test the reinforcing effects of campaign advertising on partisanship and voting intention (see Iyengar and Petrocik 2000). The recruitment methods were unchanged – flyers posted in public locations, announcements in newsletters, and personal contact in shopping malls offering payment for participation in "media research."

Finally, the 2012 study was conducted online, with participants recruited from the YouGov online panel. This study was designed to monitor subjects' evaluations of ad content. Instead of asking for a summary evaluation of the ad, participants reacted continuously while the ad was playing. They were instructed (and given a practice task) on how to move a slider located immediately below the video in accordance with their feelings about the content of the ad (for details on the dials methodology, see Iyengar, Jackman, and Hahn 2016). Depending on the condition to which they were assigned, participants reacted (using the slider) to a pair of ads, one from each of the campaigns; the pairings manipulated the number of negative ads from one to three. After viewing the ads, participants completed the post-test survey which included feeling thermometer ratings of both candidates.

Analysis Procedure

As described above, the experiments were not intended to test how campaign ads influence affective polarization; instead they focused on turnout and candidate choice. Thus, reanalyzing the data from these experiments for purposes of testing the effects of advertising tone on partisan hostility requires making a number of choices. In the interests of transparency, we outline each of these below.

First, we standardized the experimental treatment conditions. Depending on the study, a subject was exposed to between zero and three political ads, some negative and some positive. We discard respondents who saw an equal number of positive and negative ads and focus on those who saw exclusively positive ads, exclusively negative ads, or no political ads at all.

The second decision concerns operationalization of the outcome variable. In the majority of cases, we define affective polarization as the feeling thermometer rating of the in-party candidate minus the rating of the out-party candidate. However, in a handful of 1992 experiments, the treatment condition was randomized across both Senate and Presidential races using the same control group of respondents who saw no political ads. For these studies, we average feeling thermometer ratings across the in-party Senate and Presidential candidates and subtract the corresponding average for the out-party candidates.

Finally, to aggregate treatment effect estimates across studies, we report OLS regressions with study fixed effects, such that respondents are only compared within the context of each study. We report regressions without controls and with covariates included via a Lasso double-selection procedure (Belloni, Chernozhukov, and Hansen 2014). In practice, Lasso selects only indicators for whether the respondent is a Democrat or a Republican; we incorporate these using the method proposed by Lin (2013). Given our directional hypotheses — we do not expect campaign advertisements to reduce polarization — we assess statistical significance using one-sided tests, reporting both robust standard errors and randomization inference p-values against the sharp null hypothesis of no individual-level treatment effect.

Results

Table 1 presents the aggregated treatment effect across the advertisement experiments. Respondents who saw a negative advertisement, relative to those who saw no political ads,

Table 1: Experimental Results

	No Controls	Lin Regression
Positive Ad Treatment	2.56	1.78
	(2.39)	(2.30)
Negative Ad Treatment	5.13*	4.25*
	(2.23)	(2.15)
Observations	2390	2390
R^2	0.093	0.206
Study FEs	✓	✓
Controls		✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization, pooled across all advertisement experiments. The treatment indicators capture whether the respondent was assigned to a positive or negative advertisement condition, relative to seeing no political ads. Column 1 regresses the outcome on treatment without controls. Column 2 replicates this specification with LASSO-selected covariates incorporated via the Lin estimator Lin 2013. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (one-tailed test)

increased the gap between their in-party and out-party evaluations by 5.1 points on a 101-point scale. The estimated effect is similar in magnitude with controls, though the point estimate decreases modestly to 4.3 points. The estimated effect of positive advertisements is substantively smaller and statistically insignificant. To get a sense of the effect size, consider the 1996 experimental sample, which is the largest. Respondents in the control group had an average affective polarization level of 31.3, meaning that exposure to negative ads elicit a roughly 12-16% increase in affective polarization relative to the control group mean.

Given the modest sample sizes, the asymptotic approximations underlying conventional regression-based p-values may be unreliable. As a complement, we report exact p-values from randomization inference under Fisher's sharp null hypothesis that the treatment has zero effect on every unit. The procedure repeatedly permutes the treatment vector consistent with the original assignment mechanism and computes the test statistic under each permutation, yielding a reference distribution that is exact in finite samples and does not rely on distributional assumptions (Imbens and Rubin 2015). Because randomization inference delivers an exact p-value in finite samples without relying on the asymptotic approximations behind conventional regression inference, testing the sharp null provides a check against the possibility that our main results are small-sample false positives.

For the effect of the negative ad treatment relative to the neutral condition, the one-sided p-value for the sharp null hypothesis is 0.008 without covariates and 0.018 with Lasso-selected covariates. The randomization inference p-values track our regression-based results closely, alleviating concerns that the latter are artifacts of small-sample distortions, and reinforcing the finding that the negative ad treatment increases affective polarization. Comparing the negative and positive ad treatments, the p-values for the sharp null hypothesis are marginally significant at 0.066 without controls and 0.063 with controls. These results, although failing to meet conventional levels of significance, are at least suggestive that the polarizing effects of campaign ads are distinctive to negative ads.

Another source of heterogeneity to examine is variation in the treatment effect over time. The treatment effect estimate for the 2012 experiment, 4.8 (SE = 3.6), is roughly 10% less than that in the 1990s studies, 5.4 (SE = 2.8). We caution against over-interpreting this pattern, however, as each individual experiment is underpowered given its effect size (the average sample size per study is 299). We revisit this trend when we turn to the observational evidence below.

4 Observational Research Design

While experimental studies provide a benchmark for internal validity, whether the effects generalize beyond the forced-exposure laboratory setting is questionable. In the real world, citizens may never encounter campaign ads, and when they do, may quickly forget what they saw.

Achieving tight control over the advertising message is not possible using real-world treatments. Observational designs are thus limited to comparisons of residents from geographic locations characterized by high or low levels of campaign advertising. In the U.S. context, given the state-by-state allocation of electoral votes, presidential campaigns limit their advertising to "battleground" states in which support for the major candidates is relatively even (Shaw 2006; Ashworth and Clinton 2007).

However, simply using residents of high exposure battleground media markets as the treated group is problematic. As prior studies on campaign advertisements have noted, areas targeted for advertising are also targeted for other forms of campaigning (e.g., appearances by the candidate, interviews on local media, and canvassing), which introduces any number of campaign-related confounded variables. More importantly, campaigns behave strategically in the allocation of their advertising budgets, targeting areas in which the audience for the ads is likely to respond favorably to their message. In 2024, Harris ads targeted Harris-leaning voters, while the Trump campaign tried to reach MAGA-inclined voters. Exposure to advertising is thus not only non-random, but also correlated with the political predispositions of the targeted audience. In the case of our outcome variable of interest—affective polarization—comparing residents of battleground and non-battleground states would confound exposure to advertising with voters' partisan preferences, the key driver of out-party animus.

Given these identification challenges, Huber and Arceneaux (2007) show that the closest observational approximation to an experimental treatment assigns counties that fall within a battleground media market, but lie across the state border and therefore outside of the battleground state itself, as treated units. Voters in these areas receive broadcast programming from battleground media, but are not exposed to broader campaign activities that typically accompany heavy ad spending. Moreover, unlike residents of battleground DMAs, they have not been targeted on the basis of their candidate preferences. In short, the design estimates an average effect for the treated counties in a non-battleground state that happen to belong to "saturated" DMAs by comparing them against counties in the same non-battleground state that belong to lower exposure markets.

Subsequent work has built and extended this approach in three ways. First, scholars have noted that substantial non-random variation in advertising intensity exists within non-battleground states, and have sought to further refine the control group by re-weighting observations (Urban and Niebler 2014) or restricting comparisons to geographically adjacent counties (Spenkuch and Toniatti 2018). Second, the original Huber and Arceneaux design is geographically limiting by construction, as it relies exclusively on counties in non-battleground states that happen to fall within battleground DMAs. Subsequent work has argued that the core identifying assumption—that campaigns do not explicitly target individual border counties—holds more broadly, even when the sample is expanded to include all counties along media market boundaries regardless of whether the DMA is in a battleground (Spenkuch and Toniatti 2018; Sides, Vavreck, and Warshaw 2022). For example, Spenkuch and Toniatti (2018) find that border counties constitute on average only 5% of a DMA's population, suggesting these areas are peripheral not only geographically but also to campaign strategy. Third, cross-sectional or single-election comparisons of narrow geographic regions risk conflating advertising differentials with long-term residential sorting or shifts driven by electoral coalition trends. To guard against this, Yamaya (2026) exploits within-election variation to study patterns in campaign contributions, using a difference-in-differences design that differences out election-specific geographic confounds. The present paper adopts a similar identification strategy, applying it to the question of affective polarization across five election cycles.

The design conducts paired comparisons across neighboring counties in the same state that straddle a DMA boundary. We illustrate this in Figure 1, using the state of Georgia. Counties in the sample are shaded in gray and largely fall outside of major metropolitan areas. Following previous work, we enumerate all possible pairs of counties that share a border (Spenkuch and Toniatti 2018; Sides, Vavreck, and Warshaw 2022; Yamaya 2026). These county pairs serve as the basic comparison groups, with variation in advertising within these pairs providing the identifying contrast.

Figure 1b shows how advertising intensity varies within the election cycle. Advertising

Figure 1: Research Design: Georgia DMA Boundaries and Variation in Ad Exposure

(a) Border counties in GA (b) Ad Airings Within Election Cycle

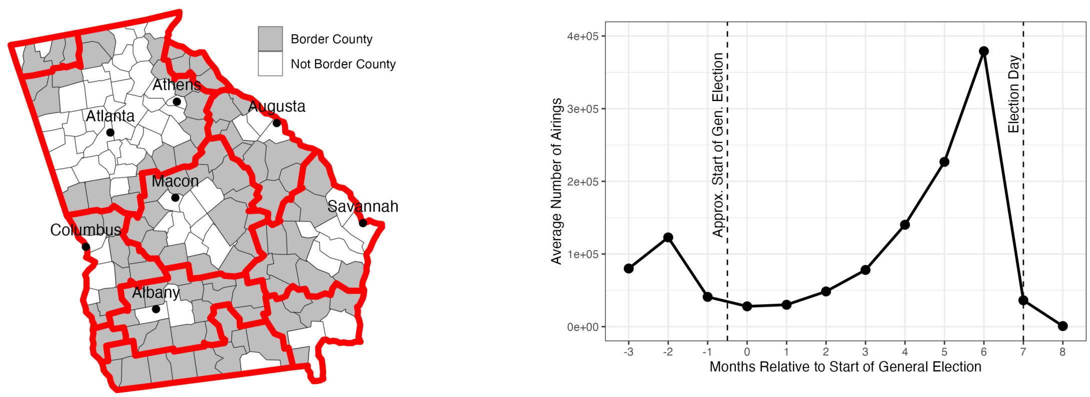

Note: Panel A displays a map of Georgia with DMA boundaries outlined in red. Counties that straddle a DMA boundary and enter our sample are shaded. Major cities are indicated for reference. Panel B plots the monthly average number of presidential campaign advertisement airings relative to the start of the general election campaign, averaged across the five election cycles in our sample.

in presidential campaigns intensifies as the election approaches. There is typically a modest bump during the primaries, followed by a sharp increase through the general election. We use this within-election change in advertising exposure as the treatment, and compare it against changes in partisan animus for our research design.

Implementing the design described above requires granular data on the volume of televised advertising across U.S. media markets over the course of a campaign, as well as repeated, large-scale measures of affect toward political parties and candidates drawn from a sufficiently broad set of designated market areas (DMAs). Such data are available for the 2000, 2004, 2008, 2020, and 2024 election cycles; comparable data are not available for 2012 and 2016. We next describe the three primary data sources used in our empirical analysis and then outline our estimation strategy.

Data

Advertising data: We draw our campaign advertising data from two sources. For the 2000, 2004, and 2008 presidential election cycles, we use the Wisconsin Advertising Project

(Goldstein, Franz, and Ridout 2002; Goldstein and Joel Rivlin 2007; Goldstein et al. 2011) to measure exposure to political advertising. These data cover all presidential campaign ads that aired on traditional network and cable broadcasts. In 2000, the dataset covers the 75 largest DMAs; coverage expanded to the 100 largest DMAs in 2004 and to all 210 DMAs in 2008. The Wisconsin Advertising Project database includes the precise number of airings per ad, the media market location, and human-coded categorization of ads based on their tone.

For the 2020 and 2024 presidential elections, we use advertising data provided by AdImpact, a private consulting firm that tracks political advertising across national TV and cable networks in all 210 DMAs. AdImpact monitors broadcasts airing weekdays between 7 a.m. and 1 a.m. and on weekends between 7 a.m. and midnight. The firm transcribes each ad and codes the transcripts for tone (as negative, contrast, or positive) and issue content. The data are validated in three ways: AdImpact cross-checks advertisements against scraped FCC public records and manually reviews any discrepancies, records are further verified against sources reporting directly from advertising agencies, and expected airing dates are continuously monitored with discrepancies manually reviewed. For each ad, AdImpact provides multiple indicators of exposure, including gross rating points (GRPs) and total number of airings.

Our analysis of advertising tone collapses contrast and negative ads into a single category. This decision is based on a content analysis of all presidential ads aired in 2024 using Stanford undergraduates as coders. The results showed that, on average, the ratio of negative to positive information in contrast ads was evenly negative and positive—around 50:50. Since negative information is more psychologically memorable than positive information, we argue that individuals are more likely to pick up on the negativity within contrast ads as opposed to the positive messaging. We are also more interested in measuring the effect of all negativity in political advertising which includes both negative and contrast ads. GRPs—the most precise indicator of ad exposure—are only available for the 2020 and 2024 cycles. For earlier cycles, we rely on total number of airings. The high correlation between weekly GRPs and airings in 2020 and 2024 validates using the number of airings as an exposure measure (Appendix A reports a linear relationship and correlations of .69 and .67, respectively). For consistency across all five election cycles, we use total airings as our primary measure of ad exposure, though we report results for both measures wherever relevant and find that differences in operationalization do not affect our substantive conclusions.

Survey data: Our survey data come from four sources covering the 2000 to 2024 election cycles: the National Annenberg Election Surveys (NAES), Nationscape Insights, the Polarization Research Lab's America's Political Pulse survey, and aggregated 2024 surveys from the YouGov-administered Stanford-Arizona State-Yale study. Across these sources, the survey instruments used to measure affective polarization differ modestly, as we document below. To facilitate comparisons of effect sizes across elections, we standardize all outcomes relative to their within-unit variation (Mummolo and Peterson 2018).

For the 2000, 2004, and 2008 elections, we use the National Annenberg Election Surveys (Annenberg Public Policy Center 2000, 2004, 2008)—large-scale rolling cross-sections of daily phone interviews with American voters (N = 100,626, 98,711, and 57,967, respectively). For all three iterations, the NAES selected respondents via random digit dialing, making it a nationally representative, probability sample. Respondents rate their affect toward presidential candidates on a 101-point feeling thermometer scale in 2000, and a 11-point scale in 2004 and 2008.

For 2020, we merge the advertising data with Nationscape (Vavreck and Tausanovitch 2021), a large-scale rolling cross-section online survey administered by researchers at UCLA (N = 494,796). Nationscape conducted weekly, nationally representative surveys via Lucid. The surveys measure affect toward Democrats, Republicans, and major presidential candidates on a 4-point Likert scale.

For 2024, we draw on two separate sources. First, we use the Polarization Research Lab's America's Political Pulse survey (Westwood and Lelkes 2024), which measures Americans' affect toward the two major parties on a 101-point feeling thermometer scale ($N =$ 157,954). Second, we use the Stanford-Arizona State-Yale (SAY) rolling cross-section survey from YouGov that interviewed approximately 797750 Americans at seven points during the campaign between December 2023 and the weeks immediately following the election. The SAY surveys record favorability toward Joe Biden, Kamala Harris, and Donald Trump on a bipolar 4-point scale, including an option of no opinion; we treat no opinion as the midpoint.

In all surveys, we compute respondents' average rating toward their in-party and subtract their average rating toward the out-party. Since the measure of affect differs across surveys, we utilize measures of both party and candidate affect. These measures also vary between a 4-point favorability and 101-feeling thermometer scale. We show correlation between measures in Appendix F to guard against concerns over measurement comparability. Because we focus on respondents residing in counties that straddle a media market boundary, we retain only approximately 40% of respondents in each dataset. We report sample characteristics and their differences from the general U.S. population in Appendix B. We also run balance tests on various demographic covariates from the surveys. While we retain a number of covariates across all surveys, such as race, party affiliation, and family income, some variables are not available.

Supplemental data: To match survey respondents to their DMA, we construct a county-to-DMA crosswalk using Capital IQ Pro's MediaCensus data, which reports all U.S. broadband and video operators by ZIP code, county, and DMA. We extract unique county-DMA pairings from the broadcaster geographies to build the crosswalk. Counties that fall into multiple DMAs are excluded from the analysis, since our public opinion surveys do not provide sufficiently precise location data to determine which portion of such counties respondents inhabit. We merge respondent locations with county geographies using shape files from the Census Bureau.

We run demographic balance tests between in-sample and not-sampled counties using IPUMS National Historical Geographic Information System data (Schroeder et al. 2025).

This dataset provides county-level demographic information on population, household income, race, gender, age, education, and more. We regress numerous demographic variables on a binary variable indicating whether the county was included in our sample for that particular election cycle.

To measure the correlation between our measures of affect, we leverage data from the 2016 George Washington University and 2020 Kent State University team module of the Cooperative Election Study (Gross 2019; Beam 2022). The George Washington University team module contains both feeling thermometer and favorability measures for party and candidates. The Kent State University team module contains feeling thermometers for both party and candidates. These modules combined allow us to calculate correlation statistics across measurements to ensure comparability.

Estimation

We estimate the design using a stacked two-way fixed effect model. In the model we interact the county and bi-monthly fixed effects with the aforementioned county border pair indicators, such that the identifying variation comes from comparing advertising intensity across geographically adjacent counties within a two-month period of time. The estimating equation is given below:

$$Y_{icp} = \theta_{c\times p} + \eta_{t\times p} + \tau_1 Negative_{icp} + \tau_2 Positive_{icp} + X\beta_{icp} + \varepsilon_{icp} \tag{1}$$

Where Y_{icp} is the level of affective polarization for respondent i, living in county c pair p. Location fixed effects (θ) are defined at the county-pair level and time fixed effects (η) are bi-monthly periods interacted with the county pair fixed effects. This specification compares respondents who take the survey in the same two-month campaign period and reside across the enumerated DMA boundaries. We also present results for alternative specifications with different treatment time-windows and time fixed-effect intervals. In some specifications, we add individual-level controls for party identification, race, age group, household income, and education levels.

Our primary treatment variable ($Negative_{icp}$) represents cumulative negative and contrast ad airings show in county c pair p in the previous 7-day window at which respondent i took the survey. This window choice is grounded in prior work (Gerber et al. 2011) showing TV ad effects dissipate within 1–2 weeks. For the 2000s, since the Wisconsin Ad Project contains specific airdates for each ad airing, we are able to calculate the precise 7-day running total for each individual respondent. In the 2020s, the AdImpact database aggregates ad airings at the weekly level. Since we lack precise airdates for each individual airing, we split respondents based on the day that they were surveyed. For respondents that were surveyed on Mondays, Tuesdays, and Wednesdays, they are assigned the previous week's level of ad exposure. Respondents that were surveyed on Thursday, Friday, Saturday, and Sunday are assigned the current week's level of ad exposure. This cut off is based on the empirical weekday distribution of ad airings in the 2000s data where we have specific airing dates and accounts for the fact that respondents at the beginning of the week may not have been exposed to a significant amount of ads in the current week. We find that weekdays average around 2000 ad airings per day whereas airings over the weekend diminish by 50%. We regress this with $Positive_{ctp}$ (cumulative 7-day positive ad airings) to isolate the effect of negative advertising from overall campaigning.

5 Results from the Observational Study

We begin by presenting election-specific estimates of the effects of exposure to negative advertising. Figure 2 plots the effect of exposure to negative ads on affective polarization in the five election cycles. To facilitate comparisons over years, both outcome and treatment variables are standardized, so coefficients can be interpreted as changes in standard deviation units. Each point corresponds to a separate regression for that year, with the treatment

Figure 2: Effect of Negative and Contrast Advertising on Affective Polarization Across Election Cycles

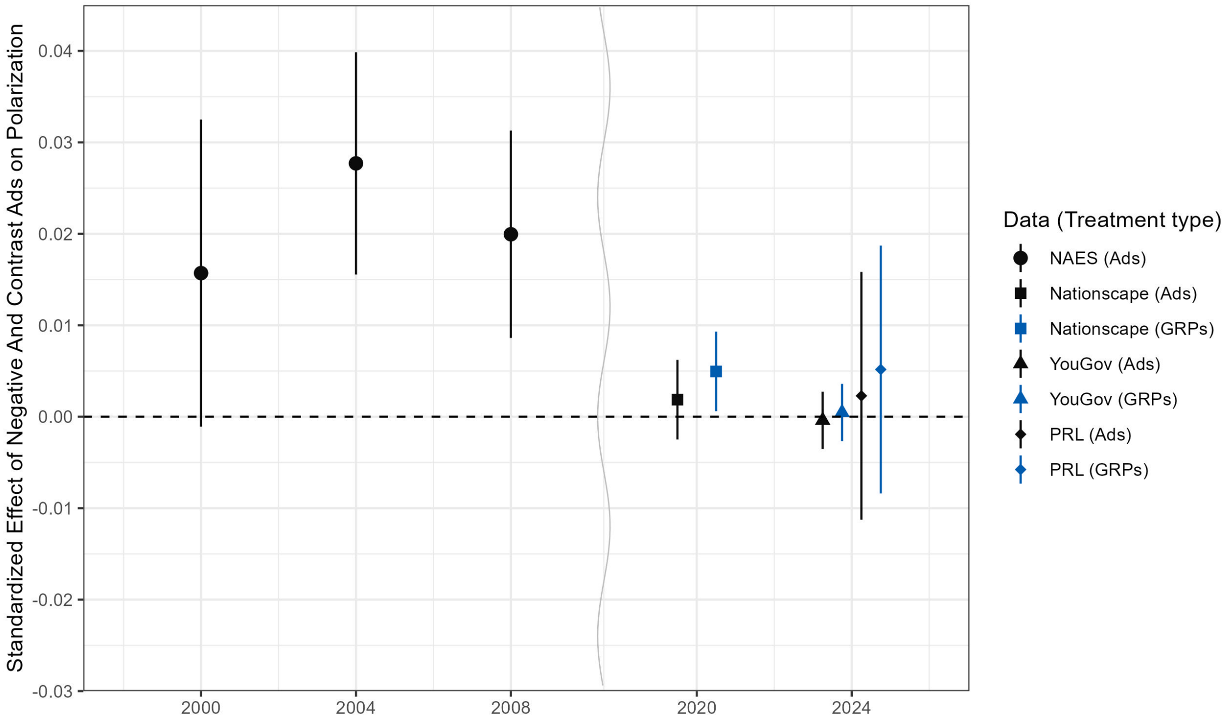

Note: Each point plots the estimated coefficient on negative advertising from Equation 1, with a separate regression estimated for each election year. The treatment variable is either number of ad airings or gross rating points (GRP), depending on data availability. The dependent variable is the level of affective polarization reported by the survey respondent. Both treatment and outcome are standardized by their within-unit standard deviation. Standard errors are clustered at the DMA level, and error bars represent 95% confidence intervals.

operationalized as either negative ad airings or GRPs, depending on data availability.

As shown in Figure 2, televised advertising in the 2000s generally shifted voters' evaluations of candidates and parties. A one-standard deviation change in the within-unit treatment leads to a 0.016, 0.028, and 0.020 standard deviation change in affective polarization for 2000, 2004, and 2008, respectively. The coefficients for 2004 and 2008 are statistically significant whereas the coefficient for 2000 was marginally significant (p-value = 0.067).

To put these numbers in perspective, first consider the year 2000. The 2000 NAES measures partisan affect on a 101-point feeling thermometer, and on that scale we estimate that every additional 1000 ad airings are associated with roughly 6.7 points of polarization. Affective polarization in our sample peaked at 35.4 points in December 2000 and bottomed at 19.4 points in June 2000. Our estimates imply that at the height of the election advertising season in October, the average weekly volume of negative ad exposure corresponds to approximately 10.4% of this bottom-to-peak swing. While we caution against additively extrapolating weekly effects across the full election season, we nonetheless view this as a substantial share of within-year movement attributable to negative advertising alone.

We can do the same effect quantification for the 2004 and 2008 elections. Starting in 2004, the NAES measures partisan affect on an 11-point feeling thermometer scale. On this scale, our estimates for both 2004 and 2008 imply that every additional 1,000 airings of negative ads are associated with approximately 0.50 points of polarization. For 2004 and 2008 respectively, the average weekly volume of negative ads in October accounts for approximately 12.33% and 11.4% of the bottom-to-peak swing in affective polarization.

By the 2020s, however, the polarizing effects of advertising have declined significantly. The 2020 election cycle is the first year where we have both GRPs and the number of airings. Number of airings yields a standardized coefficient of 0.002 while GRPs yields 0.005. Using the same back-of-the envelope calculations, this means that a typical 7-day window of negative ad airings and GRPs corresponds to approximately 1.79% and 7.05% of overall variation in affective polarization respectively. Regardless of whether we use negative ad airings or GRPs, negative advertising no longer elicits shifts in partisan animus of similar magnitude to earlier cycles.

The results are no different for 2024, when using the SAY dataset from YouGov (-0.0004 for negative airings and 0.0005 for negative GRPs) or the data from PRL (0.002 for negative airings and 0.005 for negative GRPs). The estimates using the PRL data are particularly useful because that survey also use a 101-point thermometer scale, facilitating comparisons to estimates in the 2000 NAES. The 2024 estimates suggest that an additional 1,000 negative airings are associated with an approximately 0.22 unit increase in polarization—over a 95% decrease from 2000.

To reinforce the diminishing effect of negative political advertising on affective polarization, we aggregate estimates across surveys and decades using the bootstrap. We randomly resample units at the border-pair level with replacement for each election year and take an equally weighted average across all survey-election year estimates over 2,000 simulations. The aggregate bootstrap estimate is 0.0108 (SE = 0.002), confirming that negative advertising has an overall polarizing effect. However, the effect of negative advertising diminishes significantly from the 2000s to the 2020s, with a decade-over-decade decline of -0.0195 standard deviations (SE = 0.005), suggesting that the overall polarizing effect is primarily driven by the earlier period.

This decline is not simply a function of advertising scale or volume; the same holds for proportional effort. To take into account the fact that total presidential advertising volume has increased almost seven fold between 2000 and 2024, we can compare changes in the raw scale against within-unit standard deviation changes in the treatment. A one-standard deviation change in ad volume in 2000 leads to a 0.482 change in the 101-point scale, while the same one-standard deviation change in 2024 leads only to a 0.066 change. Instead, as we will elaborate later, an important finding is that the movement in polarization over the entire election cycle has become much smaller over time. In 2024, polarization moved from its lowest point in February 2024 at 50.6 to a peak of 56.7 in September 2024, a mere difference of 6.1 points. Given the stability of partisan attitudes, fluctuations in advertising are less likely to contribute to polarization today.

Next, we compare the polarizing effects of negative advertising against the effects of positive advertising. As shown in Figure 3 below, positive ads do not polarize. The coefficients for positive advertisements are consistently small and statistically indistinguishable from zero, aside from the SAY dataset in 2024 specifically, underscoring that the polarizing effects of political campaigns derive primarily from negative advertising. The pattern is consistent with our broader argument that the decades-long rise in affective polarization is tied, in part, to the parallel growth in negative presidential advertising.

Figure 3: Effect of Positive Advertising on Affective Polarization Across Election Cycles

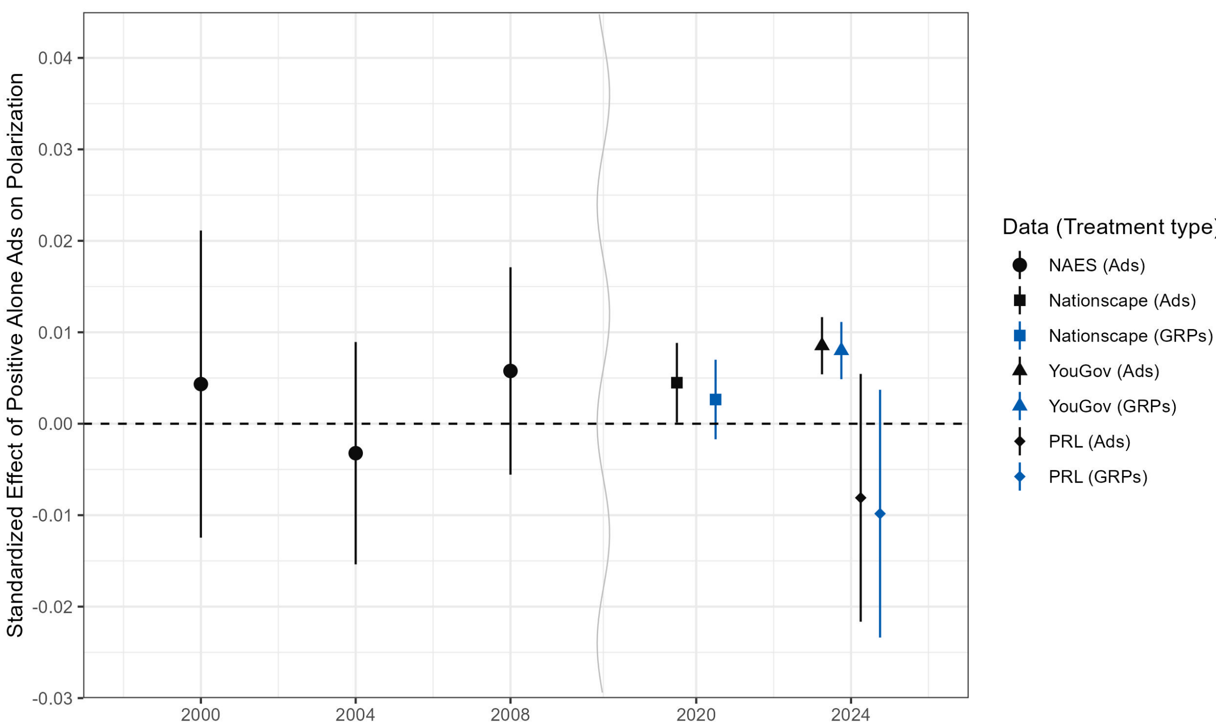

Note: Each point plots the estimated coefficient on positive advertising from Equation 1, with a separate regression estimated for each election year. The treatment variable is either number of ad airings or gross rating points (GRP), depending on data availability. The dependent variable is the level of affective polarization reported by the survey respondent. Both treatment and outcome are standardized by their within-unit standard deviation. Standard errors are clustered at the DMA level, and error bars represent 95% confidence intervals.

Why might negative advertisements be more polarizing? Most directly, they denigrate the opposing candidate by criticizing issue positions, past performance in office, and often personal character. In contrast, positive advertisements emphasize the sponsoring candidate's credentials and accomplishments, providing little basis for viewers to update their evaluations of the opponent. Negative ads may also elevate in-group evaluations by articulating substantive contrasts between the sponsoring and targeted candidates. Whereas positive appeals are one-sided, negative advertisements foreground differences between parties and candidates across salient policy domains and character traits.

We can assess the extent to which negative advertising influences both components of affective polarization (in-party approval, out-party disapproval). As shown in Figure 4, when we separate the in- and out-party evaluations, the effects of exposure to negative advertising primarily degrade evaluations of out-party candidates, with an in-party boost appearing only in 2004. By 2020, negative advertising no longer meaningfully moves either component of affective polarization.

The declining role of campaign advertising as a polarizing force likely reflects several interconnected features of contemporary American politics. Most importantly, the electorate has become increasingly sorted in the sense that their partisan and ideological affiliations now converge (Fiorina and Levendusky 2006). This has led to reduced cross pressures and a strengthened sense of partisan identity (Mason 2016). By contrast, in the early 2000s, partisan and ideological cleavages were less tightly aligned, producing greater identity cross-pressures that tempered hostility toward the out-party (see Cassese 2020; Brader, Tucker, and Therriault 2014). Under these less sorted conditions—and in the presence of comparatively weaker baseline animus—short-term campaign forces, such as negative advertising, had greater capacity to intensify affective polarization.

Today, however, given the level of social sorting and the highly elevated state of out-party animus (the modal thermometer score for the opposing party in the past two cycles of the ANES is 0), partisans have much stronger priors regarding their evaluations of the in and out

Figure 4: Decomposing Affective Polarization: In-Party Versus Out-Party Affect

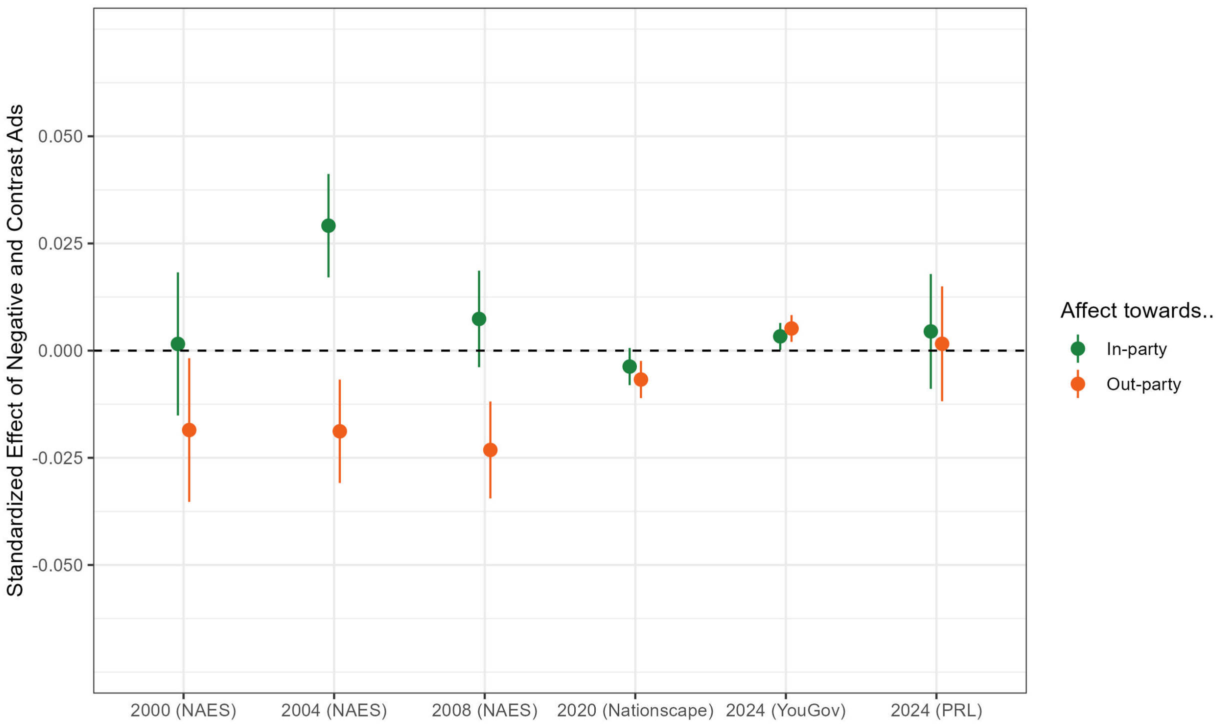

Note: Each point plots the estimated coefficient on negative advertising for either in-party or out-party affect, with a separate regression estimated for each outcome and election year. The treatment variable is the number of ad airings. Both treatment and outcome are standardized by their within-unit standard deviation. Standard errors are clustered at the DMA level, and error bars represent 95% confidence intervals.

party. The effects of short-term cues such as exposure to campaigns are therefore weakened.

Another development that has contributed to the strengthening of out-party animus is the development of new forms of journalism that inject partisan bias into news reporting. In the 1970s, Americans of all political stripes were exposed to the same news content – delivered by the three major television news network in "point - counterpoint" fashion, devoid of partisan slant (see Gaebler et al. 2023). In the post-Internet era, however, journalistic norms changed and multiple news providers now inject partisanship into their reporting. The availability of partisan media has enabled Democrats and Republicans to self select into "echo chambers" so as to encounter news that supports rather than challenges their political preferences (for evidence of selective exposure to news, see Iyengar and Hahn 2009; Broockman and Kalla 2024). Increased exposure to one-sided information flows has contributed to strengthened partisan predispositions. A more mechanical explanation for diminishing returns to negative advertising lies in the skewed distribution of partisans' prior attitudes. Contemporary partisans enter campaign season with out-party evaluations already at or near their lower bound, leaving limited scope for further decline (see Fasching et al. 2024). As a result, the polarizing effects of campaign communication are likely to operate primarily through shifts in in-party evaluations. However, given the extensive time-series evidence demonstrating the relative stability of in-party affect (Iyengar, Sood, and Lelkes 2012), variation in advertising tone is unlikely to be a principal driver of in-group favoritism.

A final possible explanation for the diminished effects of negative advertising is that the current media environment is highly saturated (Newell, Pilotta, and Thomas 2008) and voters are exposed to a greater amount of content and information. As a result, consumers have lower attention-spans and are less able to engage with the information they receive (Volk et al. 2025). Additionally, by being consistently inundated by negativity, viewers may already be highly polarized. This would reduce the marginal effect of each individual television advertisement.

We conclude the analysis of the survey data by documenting several robustness checks on our empirical approach, organized around three forms of threats to inference. The first is that the results are confounded by strategic placement of advertising. While we believe the media market border design addresses much of this concern, we take an additional step to rule out residual cross-sectional targeting effects. Because the bulk of presidential advertising intensity occurs during the general election, we test whether general-election ad volume predicts political outcomes such as voter registration, political interest, vote intention, and Democratic vs. Republican vote intention, measured during the earlier primary period. Since in the primary period, the majority of advertising has not yet aired, any such relationship would indicate that our results reflect underlying political differences across media markets which campaigns may use to infer where to place ads. Appendix Table D10 reports the estimates stacked across years. Across all outcomes, these estimates are not statistically significant.

The second threat to inference concerns changes in survey composition. Since our surveys consist of repeated cross sections, we want to ensure that our estimates are not driven by systematic changes in the types of people answering the survey. Appendix Figures E5a to E6b report a series of balance tests where we regress a set of demographic variables (characteristics that should not change within a year as a result of ad exposure) on ad volume with only unit fixed effects. This analysis shows that airings do not correlate with the super majority of survey demographics within a county, over time. To guard against the possibility that any imbalances drive our results, tables C3 to C8 in Appendix C reports two additional regression specifications for each year: one adding individual-level demographic controls from the survey data, and one interacting these controls with bi-monthly time period indicators to allow different demographic groups to follow different within-election time trends. The substantive interpretations of our results remain largely the same.

The same appendix also reports specifications that re-cluster standard errors at the state level and re-weight observations by the inverse of the number of times each county appears in the sample to account for the stacked structure of the data. Coefficient estimates are stable

Figure 5: Placebo Test: Pre-Election Trends in Affective Polarization (2024)

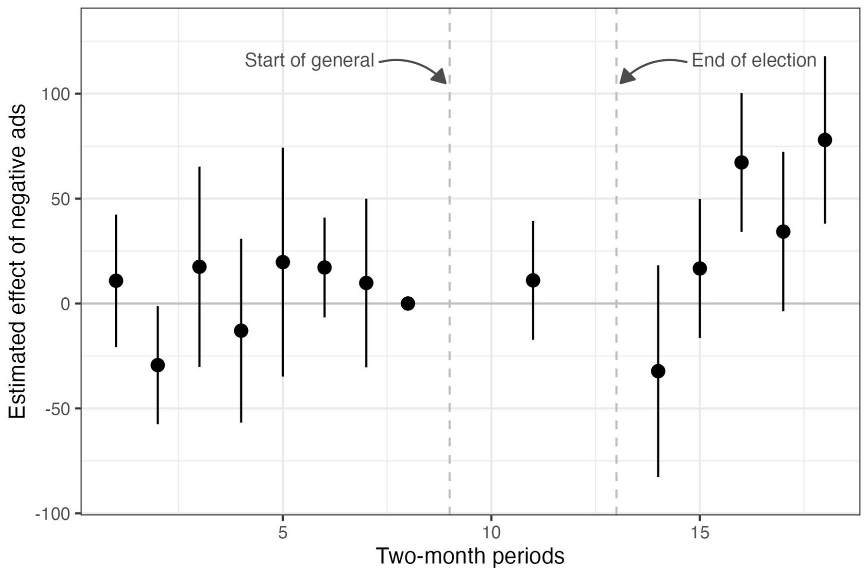

Note: This figure plots coefficients from a dynamic specification of the panel regression, in which the average cumulative negative advertising exposure in gross rating points (GRPs) during the general election period is interacted a set of two-month period indicators spanning the pre- and post-general-election periods; time periods within the general election are collapsed into a single interval. The omitted reference period is the period immediately preceding the general election. Standard errors are clustered at the DMA level and error bars represent 95% confidence intervals.

across these alternative specifications. For the 2000s surveys, where we do see an effect for negative ads, we also report alternative treatment time-windows (4-days and 10-days) and time fixed-effect specifications (weekly and monthly fixed-effects) in Table C9. We see that the results are robust to these alternative specifications.

The third concern involves unobserved, time-varying confounds that are specific to county-pairs within the election year. The standard diagnostic is the pre-trends test, but this is difficult to implement in most election years in our sample. Most of our surveys only span the calendar year of the election and do not interview respondents sufficiently far in advance of the election. Fortunately, one exception is the Polarization Research Lab's 2024 survey, which provides repeated measurements over an extended pre- and post-election period.

We estimate a dynamic model by dividing the sample into bi-monthly periods from January 2023 to December 2025 and interacting the average cumulative negative advertising

GRPs a respondent's county received during the general election with indicators for each two-month period. We use the immediate pre-general period as the reference category. Figure 5 plots the resulting coefficients. We find no evidence of differential pre-trends. Consistent with the general pattern across the 2020s, negative advertising produces only a small, positive, and statistically insignificant increase in affective polarization during the general election. However, the effect grows larger in the post-election period: counties near the media market border that received more negative advertising remain more polarized well after the election has ended.

6 Discussion & Conclusion

The experimental and survey results converge in showing that the conduct of political campaigns has meaningful repercussions for voter attitudes. Exposure to hateful speech from candidates can make partisan voters more hateful toward their opponents. The effect is stronger in the elections that occurred relatively early in the development of affective polarization and tapers off in recent campaigns. We surmise that this trend is attributable to more ingrained feelings of out-party animus caused by social sorting and other historical forces (e.g., the development of social media and providers of partisan news) that have contributed to the growth of affective polarization.

While our results are limited to presidential campaigns (and Senatorial campaigns for the experimental evidence), given the widespread adoption of negative advertising in statewide and local races, we anticipate similar polarizing effects in these campaigns. Consider the recent gubernatorial elections in Virginia and New Jersey. More than 70 percent of the ads that aired in these races offered negative rather than positive cues. Unsurprisingly, Virginia and New Jersey residents surveyed immediately after the election expressed very high levels of out-party animus with respect to the gubernatorial candidates(Wu and Iyengar 2026).

As we noted at the outset, the persuasive effects of elite rhetoric on public opinion are well documented. In a polarized environment, partisans will be more receptive to messages from copartisan candidates that demean the opposing candidate than to messages promoting the candidacy of the sponsor. Given the surge in the level of negative advertising over the past few cycles, it appears that campaign consultants are well aware of this generalization. Along with the factors we noted earlier, the greater persuasiveness of negative advertising is a strong incentive for candidates to adopt a harsh tone in their advertising.

The persuasion model is consistent with our evidence of lowered evaluations of the attacked candidate. However, we do not believe the persuasion account undermines our claim that ad campaigns are polarizing forces. In the first place, increased animus for the attacked candidate is in fact a polarizing effect at the level of candidate evaluation. More importantly, the survey evidence is overwhelming that evaluations of presidential candidates are proxies for evaluations of the national parties. In the ANES cumulative file, the correlation between the candidate-level affective polarization measure (the difference between in-party and out-party candidate ratings) and the standard party-level measure is .70. Moreover the over-time trends in candidate and party animus align almost perfectly (see Tyler and Iyengar 2024).

In closing, we note that our findings have implications for ongoing efforts to "depolarize" the electorate. The elevated level of out-party animus has set off alarm bells within the academy, American civil society organizations, and major foundations. The pressing concern appears to be that partisan animus can motivate anti-democratic behaviors such as denying the results of legitimate elections, approval of the use of violence as a means of advancing partisan interests, and the strategic spreading of misinformation. The fear is that polarization and such related changes in behavioral norms present an imminent threat to democratic institutions and norms.

These concerns have led to an outpouring of interventions designed to mitigate polarization. These interventions include facilitating interpersonal contact across party lines, correcting misperceptions of opposing partisans, and making non-partisan identities more salient (see Voelkel et al. 2024 for evidence of the efficacy of such treatments). Our results suggest that such interventions, by focusing on voter predispositions rather than elite behavior, are misplaced. To the extent that societal or contextual forces are responsible for the spread of polarization, treatments should seek to weaken these forces rather than focusing on the symptoms of the underlying problem.

One such potentially effective treatment is the "Disagree Better" initiative developed by Governor Spencer Cox of Utah and recently endorsed by the National Governor's Association. In the 2020 Utah gubernatorial election, Republican Cox and his Democratic opponent appeared in joint ads outlining their positions on major issues and making no critical reference to their opponent. Similar efforts to make campaigns more civil may break the spiral of negativity and animus so characteristic of the current state of American campaigns. However, given the powerful incentives underlying the strategy of "going negative," the prospects for such reform efforts seem, at best, less than promising.

More generally, we reiterate that understanding the origins of a societal condition is a necessary precondition for being able to treat that condition. Our findings suggest that affective polarization has worsened partly through the prevailing tone of elite communication. The fact that this effect has attenuated in recent cycles does not undermine the underlying relationship; if anything, the decline suggests that the electorate has become so polarized that sustained elite messaging can no longer move the needle. Effective interventions should therefore address the incentive structures that govern the content and tone of elite political communication rather than delivering behavioral or informational nudges aimed at curbing the symptoms of polarized individuals.

References

Ahn, Chloe, and Diana C Mutz. 2023. "The Effects of Polarized Evaluations on Political Participation: Does Hating the Other Side Motivate Voters?" *Public Opinion Quarterly* 87 (2): 243–266. https://doi.org/10.1093/poq/nfad012.

Allen, Jennifer, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts. 2020. "Evaluating the fake news problem at the scale of the information ecosystem." *Science Advances* 6 (14): eaay3539. https://doi.org/10.1126/sciadv.aay3539.

Annenberg Public Policy Center. 2000. *National Election Survey 2000.* Available from the Annenberg Public Policy Center website.

———. 2004. *National Election Survey 2004.* Available from the Annenberg Public Policy Center website.

———. 2008. *National Election Survey 2008.* Available from the Annenberg Public Policy Center website.

Ansolabehere, Stephen, and Shanto Iyengar. 1995. *Going Negative: How Political Ads Shrink and Polarize the Electorate.* New York, NY: Free Press. ISBN: 978-0-684-82284-6.

Ashworth, Scott, and Joshua Clinton. 2007. "Does Advertising Exposure Affect Turnout?" *Quarterly Journal of Political Science* 2. https://doi.org/10.1561/100.00005051.

Baumeister, Roy F., Ellen Bratslavsky, Catrin Finkenauer, and Kathleen D. Vohs. 2001. "Bad is stronger than good." *Review of General Psychology* 5 (4): 323–370. https://doi.org/10.1037/1089-2680.5.4.323.

Beam, Michael. 2022. *CES 2020, Team Module of Kent State University (KSU).* V. 1.0. https://doi.org/10.7910/DVN/BHBFKH.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen. 2014. "Inference on Treatment Effects after Selection among High-Dimensional Controls." *The Review of Economic Studies* 81 (2): 608–650. https://doi.org/10.1093/restud/rdt044.

Brader, Ted, Joshua A. Tucker, and Andrew Therriault. 2014. "Cross Pressure Scores: An Individual-Level Measure of Cumulative Partisan Pressures Arising from Social Group Memberships." *Political Behavior* 36 (1): 23–51.

Bradley, Samuel D., James R. Angelini, and Sungkyoung Lee. 2007. "Psychophysiological and memory effects of negative political ads: Aversive, arousing, and well remembered." *Journal of Advertising* 36 (4): 115–127. https://doi.org/10.2753/JOA0091-3367360409.

Broockman, David, and Joshua Kalla. 2024. "Selective exposure and echo chambers in partisan television consumption: Evidence from linked viewership, administrative, and survey data." *American Journal of Political Science* 69 (3): 847–865. https://doi.org/10.1111/ajps.12886.

Cappella, Joseph N., and Kathleen Hall Jamieson. 1997. *Spiral of Cynicism: The Press and the Public Good.* New York, NY: Oxford University Press. ISBN: 978-0-19-509064-2.

Cassese, Erin C. 2020. “Straying from the Flock? A Look at How Americans’ Gender and Religious Identities Cross-Pressure Partisanship.” *Political Research Quarterly* 73 (1): 169–183. https://doi.org/10.1177/1065912919889681.

Fasching, Neil, Shanto Iyengar, Yphtach Lelkes, and Sean J. Westwood. 2024. “Persistent polarization: The unexpected durability of political animosity around US elections.” *Science Advances* 10 (36): eadm9198. https://doi.org/10.1126/sciadv.adm9198.

Fiorina, Morris P., and Matthew S. Levendusky. 2006. “Disconnected: The Political Class Versus the People.” In *Red and Blue Nation? Characteristics and Causes of America’s Polarized Politics,* edited by Pietro S. Nivola and David W. Brady, 49–71. Washington, DC: The Brookings Institution.

Fiske, Susan T. 1980. “Attention and weight in person perception: The impact of negative and extreme behavior.” *Journal of Personality and Social Psychology* 38 (6): 889–906. https://doi.org/10.1037/0022-3514.38.6.889.

Fridkin, Kim Leslie, and Patrick J. Kenney. 2004. “Do Negative Messages Work?: The Impact of Negativity on Citizens’ Evaluations of Candidates.” *American Politics Research* 32 (5): 570–605. https://doi.org/10.1177/1532673X03260834.

Gaebler, Johann, Sean Westwood, Shanto Iyengar, and Sharad Goel. 2023. “No news is good news? The declining information value of broadcast news in America.” *PLOS One* 20 (10). https://doi.org/10.1371/journal.pone.0331607.

Galasso, Vincenzo, Tommaso Nannicini, and Salvatore Nunnari. 2023. “Positive Spillovers from Negative Campaigning.” *American Journal of Political Science* 67 (1): 5–21. https://doi.org/10.1111/ajps.12610.

Geer, John. 2012. “The News Media and the Rise of Negativity in Presidential Campaigns.” *PS: Political Science & Politics* 45 (3). https://doi.org/10.1017/S1049096512000492.

Gerber, Alan S., James G. Gimpel, Donald P. Green, and Daron R. Shaw. 2011. “How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment.” *The American Political Science Review* 105 (1): 135–150.

Gift, Karen, and Thomas Gift. 2015. “Does Politics Influence Hiring? Evidence from a Randomized Experiment.” *Political Behavior* 37 (3): 653–677.

Goldstein, Kenneth, Michael Franz, and Travis Ridout. 2002. *Political Advertising in 2000.* V. Final release. Madison, WI.

Goldstein, Kenneth, and Joel Rivlin. 2007. *Presidential Advertising 2003-2004.* Madison, WI.

Goldstein, Kenneth, Sarah Niebler, Jacob Neiheisel, and Matthew Holleque. 2011. *Presidential, Congressional, and Gubernatorial Advertising, 2008.* Madison, WI.

Gross, Kimberly. 2019. *CCES 2016, Team Module of George Washington University-B (GWU).* V. 2.0.

Huber, Gregory A., and Kevin Arceneaux. 2007. "Identifying the Persuasive Effects of Presidential Advertising." *American Journal of Political Science* 51 (4): 957–977. https://doi.org/10.1111/j.1540-5907.2007.00291.x.

Imbens, Guido W., and Donald B. Rubin. 2015. *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction.* Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781139025751.

Iyengar, Shanto. 2022. *Media Politics: A Citizen's Guide.* 5th. New York, NY: W W Norton & Co Inc. ISBN: 978-0-393-88777-8.

Iyengar, Shanto, and Kyu S. Hahn. 2009. "Red media, blue media: Evidence of ideological selectivity in media use." *Journal of Communication* (United Kingdom) 59 (1): 19–39. https://doi.org/10.1111/j.1460-2466.2008.01402.x.

Iyengar, Shanto, Simon Jackman, and Kyu Hahn. 2016. "Polarization in Less Than Thirty Seconds: Continuous Monitoring of Voter Response to Campaign Advertising." In *Political Communication in Real Time: Theoretical and Applied Research Approaches,* edited by Dan Schill, Rita Kirk, and Amy E. Jasperson, 171–195. Routledge.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22:129–146. https://doi.org/10.1146/annurev-polisci-051117-073034.

Iyengar, Shanto, and John R. Petrocik. 2000. "'Basic Rule' Voting: Impact of Campaigns on Party- and Approval-Based Voting." In *Crowded Airwaves: Campaign Advertising in Elections,* edited by David A. Dulio, Candice J. Nelson, and James A. Thurber, 113–148. Washington, DC: Brookings Institution Press.

Iyengar, Shanto, Gaurav Sood, and Yphtach Lelkes. 2012. "Affect, Not Ideology: A Social Identity Perspective on Polarization." *Public Opinion Quarterly* 76 (3): 405–431. https://doi.org/10.1093/poq/nfs038.

Iyengar, Shanto, and Sean J. Westwood. 2015. "Fear and Loathing across Party Lines: New Evidence on Group Polarization." *American Journal of Political Science* 59 (3): 690–707. https://doi.org/10.1111/ajps.12152.

Jacobson, Gary. 2023. "The 2022 Elections: A Test of Democracy's Resilience and the Referendum Theory of Midterms." *Political Science Quarterly* 138 (1): 1–22. https://doi.org/10.1093/psquar/qqad002.

Jamieson, Kathleen Hall. 1993. *Dirty Politics: Deception, Distraction, and Democracy.* New York, NY: Oxford University Press. ISBN: 978-0-19-508553-2.

Jasperson, Amy E., and David P. Fan. 2002. "An aggregate examination of the backlash effect in political advertising: The case of the 1996 U.S. Senate race in Minnesota." *Journal of Advertising* 31 (1): 1–12. https://doi.org/10.1080/00913367.2002.10673656.

Kalla, Joshua L., and David E. Broockman. 2018. "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments." *American Political Science Review* 112 (1): 148–166.

Klein, Jill G. 1991. "Negativity effects in impression formation: A test in the political arena." *Personality and Social Psychology Bulletin* 17 (4): 412–418. https://doi.org/10.1177/0146167291174009.

Lau, Richard R., David J. Andersen, Tessa M. Ditonto, Mona S. Kleinberg, and David P. Redlawsk. 2017. "Effect of media environment diversity and advertising tone on information search, selective exposure, and affective polarization." *Political Behavior* 39 (1): 231–255. https://doi.org/10.1007/s11109-016-9354-8.

Lau, Richard R., Lee Sigelman, Caroline Heldman, and Paul Babbitt. 1999. "The Effects of Negative Political Advertisements: A Meta-Analytic Assessment." *American Political Science Review* 93 (4): 851–875. https://doi.org/10.2307/2586117.

Lau, Richard R., Lee Sigelman, and Ivy Brown Rovner. 2007. "The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment." *Journal of Politics* 69 (4): 1176–1209. https://doi.org/10.1111/j.1468-2508.2007.00618.x.

Lin, Winston. 2013. "Agnostic Notes on Regression Adjustments to Experimental Data: Reexamining Freedman's Critique." *The Annals of Applied Statistics* 7 (1): 295–318.

Martin, Danielle, and Alessandro Nai. 2024. "Deepening the rift: Negative campaigning fosters affective polarization in multiparty elections." *Electoral Studies* 87:102745. https://doi.org/10.1016/j.electstud.2024.102745.

Mason, Lilliana. 2016. "A Cross-Cutting Calm: How Social Sorting Drives Affective Polarization." *Public Opinion Quarterly* 80 (S1): 351–377. https://doi.org/10.1093/poq/nfw001.

Mendelberg, Tali. 1997. "Executing Hortons: Racial Crime in the 1988 Presidential Campaign." *The Public Opinion Quarterly* 61 (1): 134–157.

Muise, Daniel, Homa Hosseinmardi, Baird Howland, Markus Mobius, David Rothschild, and Duncan J. Watts. 2022. "Quantifying partisan news diets in Web and TV audiences." *Science Advances* 8 (28): eabn0083. https://doi.org/10.1126/sciadv.abn0083.

Mummolo, Jonathan, and Erik Peterson. 2018. "Improving the Interpretation of Fixed Effects Regression Results." *Political Science Research and Methods* 6 (4): 829–835. https://doi.org/10.1017/psrm.2017.44.

Mutz, Diana C., and Byron Reeves. 2005. "The New Videomalaise: Effects of Televised Incivility on Political Trust." *American Political Science Review* 99 (1): 1–15. https://doi.org/10.1017/S0003055405051452.

Newell, Jay, Joseph J. Pilotta, and John C. Thomas. 2008. "Mass Media Displacement and Saturation." *International Journal on Media Management* 10 (4): 131–138. https://doi.org/10.1080/14241270802426600.

Patterson, Thomas E. 1994. *Out of Order: An incisive and boldly original critique of the news media's domination of America's political process.* New York, NY: Knopf Doubleday Publishing Group. ISBN: 978-0-679-75510-4.

Roese, Neal J., and Gerald N. Sande. 1993. "Backlash effects in attack politics." *Journal of Applied Social Psychology* 23 (8): 632–653. https://doi.org/10.1111/j.1559-1816.1993.tb01106.x.

Schroeder, Jonathan, David Van Riper, Steven Manson, Katherine Knowles, Tracy Kugler, Finn Roberts, and Steven Ruggles. 2025. *IPUMS National Historical Geographic Information System.* V. 20.0. Minneapolis, MN.

Shaw, Daron R. 2006. *The Race to 270: The Electoral College and the Campaign Strategies of 2000 and 2004.* Chicago, IL: University of Chicago Press. ISBN: 978-0-226-75134-4.

Sides, John, Lynn Vavreck, and Christopher Warshaw. 2022. "The Effect of Television Advertising in United States Elections." *American Political Science Review* 116 (2): 702–718. https://doi.org/10.1017/S000305542100112X.

Sinclair, Samantha, Artur Nilsson, and Jens Agerström. 2023. "Judging job applicants by their politics: Effects of target–rater political dissimilarity on discrimination, cooperation, and stereotyping." *Journal of Social and Political Psychology* 11 (1): 75–91. https://doi.org/10.5964/jspp.9855.

Skaperdas, Stergios, and Bernard Grofman. 1995. "Modeling Negative Campaigning." *American Political Science Review* 89 (1): 49–61. https://doi.org/10.2307/2083074.

Sood, Gaurav, and Shanto Iyengar. 2016. *Coming to Dislike Your Opponents: The Polarizing Impact of Political Campaigns,* 2840225. https://doi.org/10.2139/ssrn.2840225.

Soroka, Stuart, Patrick Fournier, and Lilach Nir. 2019. "Cross-national evidence of a negativity bias in psychophysiological reactions to news." *Proceedings of the National Academy of Sciences* 116 (38): 18888–18892. https://doi.org/10.1073/pnas.1908369116.

Soroka, Stuart N. 2006. "Good News and Bad News: Asymmetric Responses to Economic Information." *Journal of Politics* 68 (2): 372–385. https://doi.org/10.1111/j.1468-2508.2006.00413.x.

Spenkuch, Jörg L, and David Toniatti. 2018. "Political Advertising and Election Results." *The Quarterly Journal of Economics* 133 (4): 1981–2036. https://doi.org/10.1093/qje/qjy010.

Tyler, Matthew D., and Shanto Iyengar. 2024. "Testing the Robustness of the ANES Feeling Thermometer Indicators of Affective Polarization." *American Political Science Review* 118 (3): 1570–1576. https://doi.org/10.1017/S0003055423001302.

Tyler, Matthew D., Shanto Iyengar, and Arjun Wilkins. n.d. "Campaigns Reinforce Partisanship and Short-Term Forces: Evidence from a Large-Scale Panel Study of the 2020 US Presidential Campaign." *Political Science Research and Methods.*

Urban, Carly, and Sarah Niebler. 2014. “Dollars on the Sidewalk: Should U.S. Presidential Candidates Advertise in Uncontested States?” *American Journal of Political Science* 58 (2): 322–336. https://doi.org/10.1111/ajps.12073.

Valentino, Nicholas A., Matthew N. Beckmann, and Thomas A. Buhr. 2001. “A Spiral of Cynicism for Some: The Contingent Effects of Campaign News Frames on Participation and Confidence in Government.” *Political Communication* 18 (4): 347–367. https://doi.org/10.1080/10584600152647083.

Vavreck, Lynn, and Chris Tausanovitch. 2021. *Democracy Fund + UCLA Nationscape Project.*

Voelkel, Jan G., Michael N. Stagnaro, James Y. Chu, Sophia L. Pink, Joseph S. Mernyk, Chrystal Redekopp, Isaias Ghezae, et al. 2024. “Megastudy testing 25 treatments to reduce antidemocratic attitudes and partisan animosity.” *Science* 386 (6719): eadh4764. https://doi.org/10.1126/science.adh4764.

Volk, Sophia C., Anne Schulz, Sina Blassnig, Sarah Marschlich, Minh Hao Nguyen, and Nadine Strauß. 2025. “Selecting, avoiding, disconnecting: a focus group study of people’s strategies for dealing with information abundance in the contexts of news, entertainment, and personal communication.” *Information, Communication & Society* 28 (1): 21–40. https://doi.org/10.1080/1369118X.2024.2358167.

Walter, Annemarie S, and Cees van der Eijk. 2019. “Unintended consequences of negative campaigning: Backlash and second-preference boost effects in a multi-party context.” *The British Journal of Politics and International Relations* 21 (3): 612–629. https://doi.org/10.1177/1369148119842038.

Weiss, Chagai M, Don Green, and Robb Willer. 2025. *Politicians’ Bipartisan Appeals to Civility and Partisan Divides: A Field Experiment with U.S. Governors,* 5qxyw_v2. https://doi.org/10.31219/osf.io/5qxyw_v2.

Wesleyan Media Project. 2014. “Ad spending in 2014 elections poised to break $1 billion - Wesleyan Media Project.” https://mediaproject.wesleyan.edu/ad-spending-in-2014-elections-poised-to-break-1-billion/.

———. 2024. “Trump and Allies Launch Barrage of Negative Ads - Wesleyan Media Project.” https://mediaproject.wesleyan.edu/releases-082924/.

Westwood, Sean J., and Yphtach Lelkes. 2024. *America’s Political Pulse.* https://americaspoliticalpulse.com/data.

Wu, Victor Y., and Shanto Iyengar. 2026. “Exposure to Negative Advertising as a Polarizing Force: The Case of the 2025 New Jersey and Virginia Elections.”

Yamaya, Shun. 2026. “How Did the Internet Change Campaign Fundraising?” Working Paper.

Zaller, John R. 1992. *The Nature and Origins of Mass Opinion.* New York, NY: Cambridge University Press. ISBN: 978-0-521-40786-1.

Supplementary Information

A Descriptive of ads & Validation of Number of ads against GRPs

The current literature on political campaigning argues that advertising volume has gradually increased over the past several years (Geer 2012). At the same time, political campaigns have also grown increasingly negative in both tone and substance. To better illustrate this, figure A2 plots the movement of presidential TV advertising over the five election cycles in our study. We quantify ad volume using two metrics: total airings and estimated spending. Since 2000, total presidential TV advertising volume has increased roughly by a factor of 7 by the 2024 election. Campaigns have also become more negative by a little over 20 percentage points over the same time frame. Negative ads constituted approximately 60% of total ad spending and airings in 2000 but grew to account for over 80% of all ad volume by 2024.

Measurements for ad GRPs are only available for the 2020 and 2024 cycles from the AdImpact dataset. For the elections during the 2000s, we rely on total number of airings as a proxy for ad exposure. One common criticism of using ad volume for this purpose is that ad volume fails to account for viewer engagement with the advertisement. To further validate our usage of total ad airings as a proxy for exposure in the 2000s, we measure the correlaton between GRPs and airings in the 2020 and 2024 elections. Figure A2 implies that the relationship between weekly number of ads and GRPs is strongly positive and linear. This assuages concerns that our results are being primarily driven by a poor measurement of ad exposure, and is further evidenced by the similar effects of both measures on affective polarization.

While scholars have found ample evidence to support the cyclical nature of political advertising over the course of the campaign (see Geer 2012; Iyengar 2022, there is less documentation on the within-election trends in affective polarization. According to data from the ANES, levels of affective polarization has risen consistently across election cycles since the 1990s (Iyengar, Sood, and Lelkes 2012; Iyengar et al. 2019). We leverage data from the Polarization Research Lab's America's Political Pulse survey to plot within-election trends in affective polarization. In figure A3, we see that affective polarization gradually increases over the course of the general election: beginning at the trough near the start of the general election and peaking just before election day. Polarization decreases quite sharply after election day.

Figure A1: Trends in Presidential TV Advertising

(a) Total airings

(b) Estimated spending

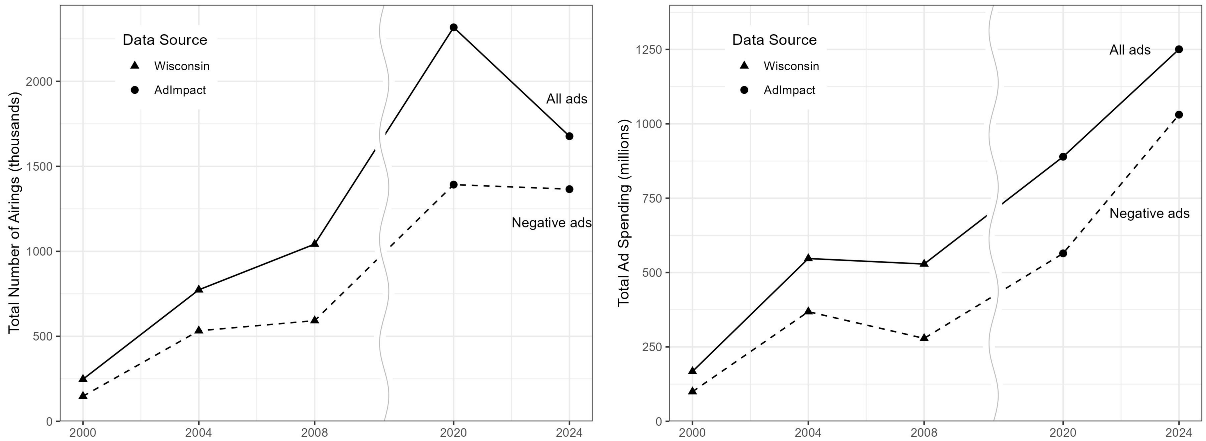

Note: This figure plots trends in presidential TV advertising across election cycles. Panel A shows the total number of airings for all ads and for negative ads separately, where negative ads combine the "negative" and "contrast" categories. Panel B plots estimated total spending for each election cycle.

Figure A2: Correlation Between Weekly Ad Airings and GRPs

(a) 2020

(b) 2024

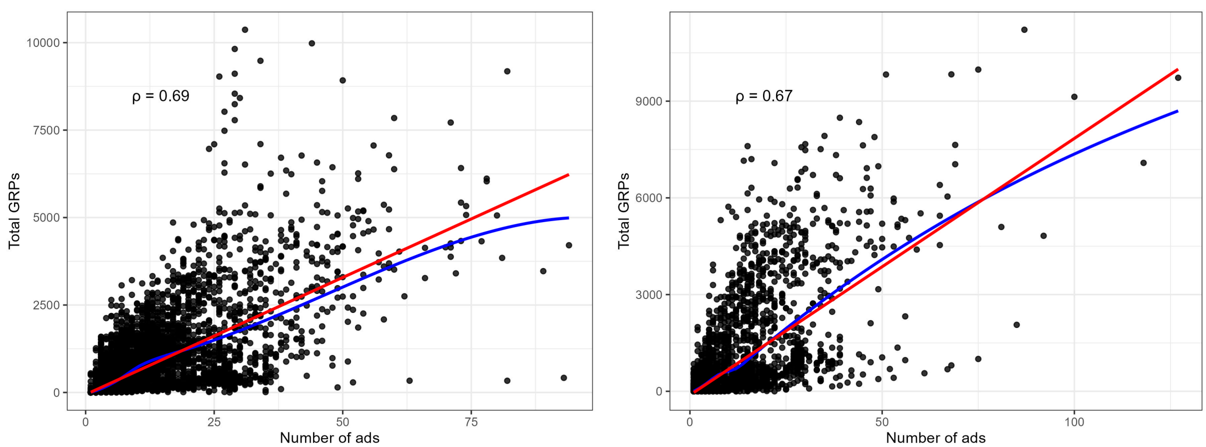

Note: This figure compares the weekly sum of campaign advertisement gross rating points (GRP) against the number of airings, separately for 2020 and 2024. Each observation is a DMA-week. The red line shows a linear fit and the blue line shows a LOESS curve.

Figure A3: Trends in Affective Polarization During the 2024 Election

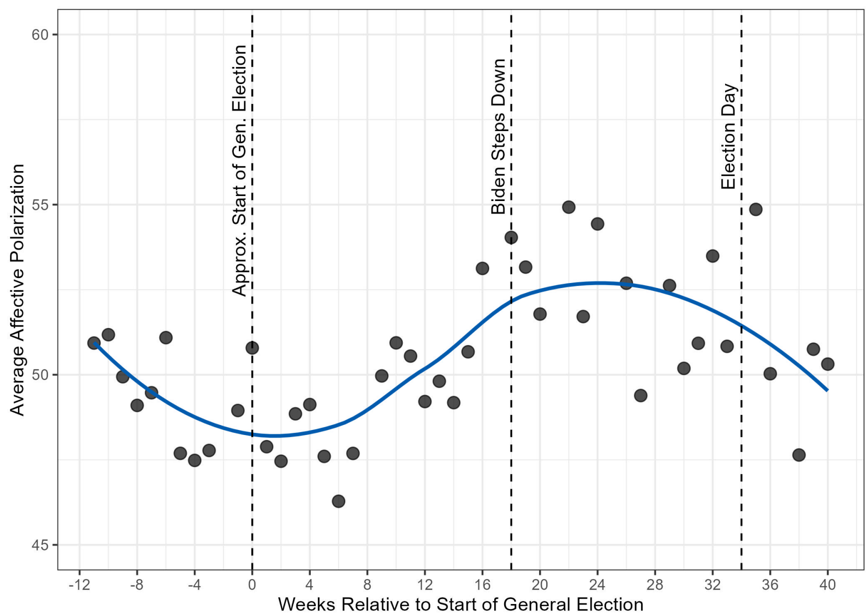

Note: Data are from the Polarization Research Lab's (PRL) America's Political Pulse surveys. Affective polarization is measured as each respondent's in-party thermometer rating minus their out-party rating. Each point plots the weekly average level of affective polarization across the 2024 election cycle. The blue line shows a LOESS curve.

B Sample Selection

In Figure B4, we depict which counties appear in our sample for each respective election cycle. Our sampling criteria proceeds in two steps. First, we filter for counties that are included in our border-pair design. Next, we take counties that have at least five observations during each election year. 2000 and 2004 are especially sparse since the Wisconsin Ad Project only covers the largest 50 and 75 DMAs respectively.

Tables B1 and B2 describes how our sampled counties differ from the rest of the counties in the United States. We regress an indicator variable for county sampling on various demographic variables from the U.S. census data. Table B2 includes state fixed effects to show how border counties differ relative to other counties in the same state. We can see that the sampled counties differ from the general population on a number of demographic variables. In particular, education and proportion of the county population being male are especially imbalanced across a majority of the datasets. These results caution us to consider potential demographic factors that could impact the generalizability of our results.

Figure B4: In-Sample Counties Across Election Years

(a) 2000

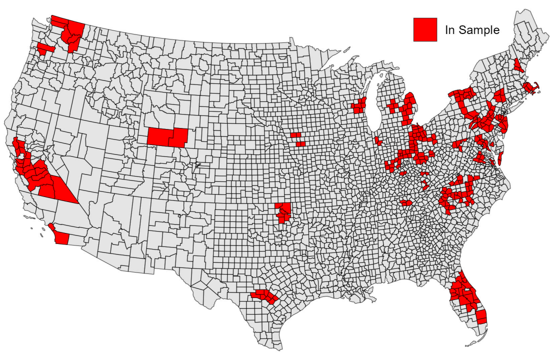

(b) 2004

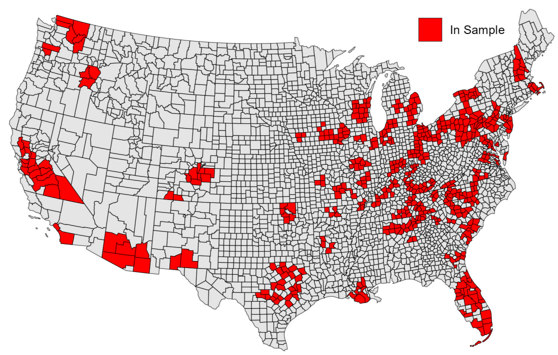

(c) 2008

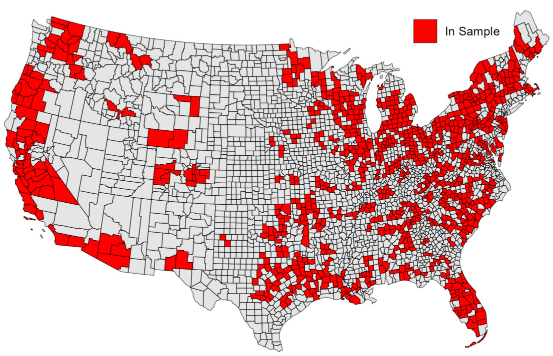

(d) 2020

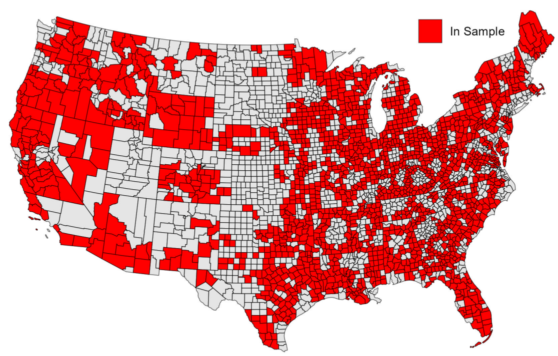

(e) 2024 – PRL

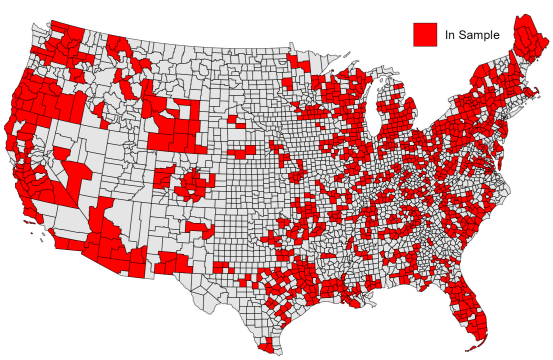

(f) 2024 – YouGov

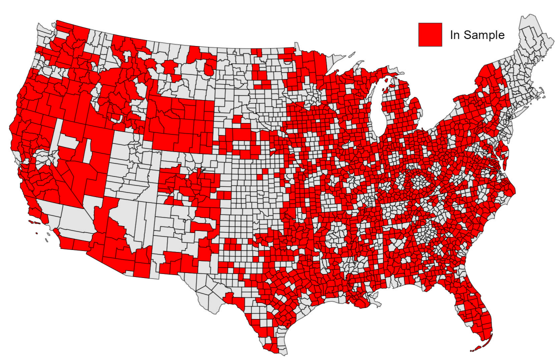

Note: This figure depicts in-sample counties across all five election years.

Table B1: Comparison of In-Sample and Out-of-Sample Counties (No State Fixed Effects)

	2000	2004	2008	2020	2024 – PRL	2024 – YouGov
Proportion Male	-1.19***	-1.81***	-2.00***	-1.52***	-2.25***	-0.80*
	(0.26)	(0.36)	(0.37)	(0.40)	(0.36)	(0.38)
Proportion White	-0.08	0.04	0.48***	0.63***	0.35***	0.53***
	(0.07)	(0.09)	(0.10)	(0.10)	(0.07)	(0.07)
Proportion Black	-0.17**	-0.12	0.10	0.51***	-0.17*	0.41***
	(0.06)	(0.09)	(0.10)	(0.10)	(0.08)	(0.08)
Proportion Over 18	1.18***	1.93***	0.95***	1.56***	1.08***	1.03***
	(0.19)	(0.25)	(0.27)	(0.28)	(0.24)	(0.26)
Proportion College Degree	-0.49***	-0.91***	-0.59***	-1.49***	-0.60***	-1.38***
	(0.09)	(0.13)	(0.14)	(0.14)	(0.10)	(0.11)
Proportion Been Married	-0.26*	-0.21	-0.99***	-0.62***	-1.13***	-0.39*
	(0.12)	(0.17)	(0.17)	(0.18)	(0.17)	(0.18)
Unemployment Rate	0.37	0.15	2.49***	0.29	1.86***	1.20**
	(0.24)	(0.33)	(0.31)	(0.39)	(0.39)	(0.41)
Proportion in Poverty	0.20	0.62**	-1.06***	-0.46*	-1.07***	-0.69***
	(0.17)	(0.23)	(0.20)	(0.23)	(0.17)	(0.18)
Intercept	-0.19	-0.58	0.86**	0.35	1.42***	0.49
	(0.22)	(0.30)	(0.32)	(0.33)	(0.27)	(0.29)
N	3,139	3,139	3,221	3,220	3,222	3,222

Note: This table presents regression results where a binary variable indicating whether a county was sampled is regressed on a number of county-level demographic varaibles. Each column represents a different election year cycle. We do not include fixed effects. We obtain all county demographic data from the NHGIS (Schroeder et al. 2025). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table B2: Comparison of In-Sample and Out-of-Sample Counties (State Fixed Effects)

	2000	2004	2008	2020	2024 – PRL	2024 – YouGov
Proportion Male	-0.75**	-0.97**	-0.83*	-0.22	-1.21***	-0.32
	(0.27)	(0.37)	(0.38)	(0.40)	(0.36)	(0.38)
Proportion White	-0.09	-0.09	0.28*	0.23	0.13	0.17
	(0.09)	(0.12)	(0.12)	(0.12)	(0.09)	(0.10)
Proportion Black	-0.19*	-0.38**	-0.25	0.09	-0.47***	-0.08
	(0.08)	(0.12)	(0.13)	(0.13)	(0.11)	(0.12)
Proportion Over 18	0.29	0.63*	-0.62*	0.53	-0.19	0.85**
	(0.22)	(0.30)	(0.30)	(0.31)	(0.28)	(0.29)
Proportion College Degree	-0.24*	-0.36*	0.18	-0.65***	-0.20	-0.88***
	(0.11)	(0.15)	(0.16)	(0.16)	(0.11)	(0.12)
Proportion Been Married	-0.05	-0.04	-0.38*	0.20	-0.59***	-0.06
	(0.14)	(0.19)	(0.18)	(0.19)	(0.18)	(0.18)
Unemployment Rate	-0.16	-0.33	0.89**	0.11	0.61	0.58
	(0.26)	(0.36)	(0.34)	(0.39)	(0.39)	(0.41)
Proportion in Poverty	0.11	0.11	-0.46	0.76**	-0.00	0.37
	(0.18)	(0.25)	(0.24)	(0.26)	(0.20)	(0.21)
N	3,138	3,138	3,220	3,219	3,221	3,221
State FE	✓	✓	✓	✓	✓	✓

Note: This table presents regression results where a binary variable indicating whether a county was sampled is regressed on a number of county-level demographic varaibles. Each column represents a different election year cycle. We include state-level fixed effects for each model. We obtain all county demographic data from the NHGIS (Schroeder et al. 2025). * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

C Year by year analyses

In appendix C, we present the raw coefficients of alternative models on the effect of campaign ads on affective polarization. We present five different models. First, we present the coefficients for the model using the non-standardized, original scale for advertising volume and polarization. Second, we cluster standard errors at the state level. Third, we re-weight the model by the number of times each county appears in the sample. Fourth, we include individual-level demographic controls in the model. Finally, we interact these demographic controls with bi-monthly period indicators to account for demographic within-election time trends. We find that the statistical and substantive interpretations of the coefficients for each election year remains largely the same. For 2020 and 2024, we also find that number of airings and GRPs have similar directional effects on polarization.

We also present alternative treatment time-windows and time-fixed effects for the models in the 2000s. We manipulate both the treatment time-windows and time-fixed effects. First, we present results changing the cumulative treatment windows to 4-day and 10-day. Next, we also present results changing the fixed effect structure to weekly and monthly fixed-effects. We aggregate across the 2000s surveys by bootstrapping the stacked data. We see that the results for the alternative treatment and fixed effect time windows do not meaningfully change the substantive and statistical implications of our results.

Table C3: The Effect of Campaign Ads (2000)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	6.71	6.71	2.44	13.31*	13.36*
	(4.74)	(4.96)	(4.46)	(5.60)	(5.61)
N Positive Ads	2.53	2.53	7.45	4.35	4.22
	(8.14)	(6.60)	(7.63)	(8.56)	(8.59)
N	15,179	15,179	15,179	13,392	13,392
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2000 election cycle. The treatment variables are the number of negative and positive presidential campaign advertisement airings (in thousands). Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C4: The Effect of Campaign Ads (2004)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	0.60***	0.60***	0.67***	0.49***	0.48***
	(0.15)	(0.16)	(0.15)	(0.13)	(0.13)
N Positive Ads	-0.16	-0.16	0.22	-0.15	-0.15
	(0.53)	(0.27)	(0.41)	(0.48)	(0.48)
N	29,478	29,478	29,478	25,298	25,298
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2004 election cycle. The treatment variables are the number of negative and positive presidential campaign advertisement airings (in thousands). Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C5: The Effect of Campaign Ads (2008)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	0.45*	0.45**	0.45**	0.40	0.39
	(0.18)	(0.16)	(0.16)	(0.24)	(0.24)
N Positive Ads	0.12	0.12	0.20	0.09	0.08
	(0.14)	(0.16)	(0.13)	(0.13)	(0.13)
N	35,599	35,599	35,599	30,948	30,948
SE Clustered by	DMA	State	DMA	DMA	DMA
County $\times$ Pair FE	✓	✓	✓	✓	✓
Bi-month $\times$ Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month $\times$ Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2008 election cycle. The treatment variables are the number of negative and positive presidential campaign advertisement airings (in thousands). Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C6: The Effect of Campaign Ads (2020)

Panel A: Number of Ad Airings

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	0.01	0.01	0.00	0.02	0.03
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
N Positive Ads	0.03	0.03	0.01	0.02	0.02
	(0.03)	(0.02)	(0.03)	(0.03)	(0.03)
N	218,547	218,547	218,547	208,346	208,346
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Panel B: Gross Rating Points (GRPs)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
GRP Negative Ads	0.22	0.22	0.18	0.30	0.31
	(0.22)	(0.16)	(0.20)	(0.21)	(0.21)
GRP Positive Ads	0.14	0.14	0.00	0.13	0.13
	(0.23)	(0.16)	(0.19)	(0.21)	(0.21)
N	218,547	218,547	218,547	208,346	208,346
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2020 election cycle. The treatment variables are the number of negative and positive presidential campaign advertisement intensity. Panel A measures treatment by ad count and Panel B by GRP exposure. Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C7: The Effect of Campaign Ads (PRL 2024)

Panel A: Number of Airings

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
GRP Negative Ads	3.46	3.46	0.78	2.46	1.85
	(4.08)	(4.29)	(4.28)	(4.52)	(4.38)
GRP Positive Ads	-28.32	-28.32	-15.82	-30.68	-27.08
	(31.90)	(23.28)	(30.38)	(32.88)	(32.41)
N	26,028	26,028	26,028	26,028	26,028
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Panel B: Gross Rating Points (GRPs)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	0.22	0.22	-0.05	-0.04	-0.09
	(0.57)	(0.60)	(0.55)	(0.67)	(0.64)
N Positive Ads	-2.89	-2.89	-1.10	-2.58	-2.29
	(4.02)	(2.93)	(3.25)	(3.98)	(3.91)
N	26,028	26,028	26,028	26,028	26,028
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2024 election cycle, using data from the Polarization Research Lab. The treatment variables are the number of negative and positive presidential campaign advertisement intensity. Panel A measures treatment by ad count and Panel B by GRP exposure. Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C8: The Effect of Campaign Ads (YouGov 2024)

Panel A: Number of Ads

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
N Negative Ads	-0.00	-0.00	-0.00	0.01	0.01
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
N Positive Ads	0.17***	0.17***	0.17***	0.14***	0.14***
	(0.03)	(0.04)	(0.02)	(0.03)	(0.03)
N	409,811	409,811	409,811	406,762	406,762
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Panel B: Gross Rating Points (GRPs)

	Original scale	Clustered	Re-weighted	Controls	Interacted controls
GRP Negative Ads	0.01	0.01	0.01	0.08*	0.08*
	(0.05)	(0.04)	(0.04)	(0.03)	(0.03)
GRP Positive Ads	1.33***	1.33***	1.42***	1.22***	1.26***
	(0.26)	(0.30)	(0.25)	(0.27)	(0.28)
N	409,811	409,811	409,811	406,762	406,762
SE Clustered by	DMA	State	DMA	DMA	DMA
County × Pair FE	✓	✓	✓	✓	✓
Bi-month × Pair FE	✓	✓	✓	✓	✓
Controls				✓	✓
Bi-month × Controls					✓

Note: This table presents regression results where the dependent variable is respondent-level affective polarization in the 2024 election cycle, using data from the Stanford-Arizona State-Yale rolling cross-section survey from YouGov. The treatment variables are the number of negative and positive presidential campaign advertisement intensity. Panel A measures treatment by ad count and Panel B by GRP exposure. Column 1 shows the baseline specification from Equation 1 with the outcome on its original scale. Column 2 adjusts the clustering structure. Column 3 reweights observations by the number of times each enters the dataset, reflecting the stacked structure of the data. Column 4 adds individual-level demographic controls, including party identification, race, age group, household income, and education. Column 5 interacts these controls with the time indicators to allow for demographic-specific time trends. The unit of observation is a respondent-county pair. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

Table C9: Bootstrap Alternative Treatment Windows and Time Fixed-Effects

	4-Days	10-Days	7-Days	7-Days	7-Days
Cumulative Negative Ads	0.0136***	0.0206***	0.0208***	0.0085*	0.0167***
	(0.0039)	(0.0042)	(0.0041)	(0.0041)	(0.0048)
Cumulative Positive Ads	0.0081	0.0015	0.0024	0.0038	-0.0079
	(0.0044)	(0.0042)	(0.0044)	(0.0042)	(0.0062)
Bi-monthly Cohort FE	✓	✓	✓		
Monthly Cohort FE				✓	
Weekly Cohort FE					✓

Note: This table presents bootstrapped estimates of campaign advertising effects, across different treatment time-windows and time fixed-effect specifications. We randomly resample units at the border-pair level with replacement for each election year in the 2000s and take an equally weighted average across all survey-election year estimates over 2,000 simulations. In column 1, we aggregate ads on a running 4-day window. In column 2, we aggregate ads on a running 10-day window. In columns 3 to 5, we aggregate ads on a running 7-day window. Bi-monthly time fixed-effects are used for columns 1 to 3, monthly fixed-effects for column 4, and weekly time fixed-effects for column 5. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test).

D Placebo tests

In Appendix D, we test whether advertisements are strategically aired by campaigns. We first compute, for each county, the average 7-day running volume of advertising survey respondents were exposed to during the general election. We then restrict the sample to survey respondents interviewed during the primary period and, within each border county pair, compare outcomes by bi-month period to test whether they are correlated with the difference in general-election ad volume that will air later in the county. We stack estimates across the NAES, Nationscape, and PRL survey data, including election-year fixed effects interacted with the border-pair by bi-month indicators. Across outcomes such as registration rates, political interest, and vote intention, we find no meaningful relationship between future general-election ad volume and primary-period political outcomes across border county pairs.

Table D10: Cross Sectional Targeting: Ad Treatment on Political Outcomes

	Registration	Political Interest	Vote Intention	Pres. Vote Int. (D)	Pres. Vote Int. (R)
Average Negative Ads	-0.01	-0.07	0.00	0.00	0.01
	(0.01)	(0.04)	(0.03)	(0.01)	(0.01)
Average Positive Ads	0.01	0.13	0.02	-0.00	-0.01
	(0.04)	(0.10)	(0.08)	(0.01)	(0.01)
N	105,817	120,408	92,085	512,127	512,127
SE Clustered by	DMA	DMA	DMA	DMA	DMA
Year × Bi-month × Pair FE	✓	✓	✓	✓	✓

Note: This table presents regression results where a number of politically relevant variables are regressed on the average amount of negative and positive ads aired during the general election. We stack observations from the 2000, 2004, 2008, 2020, and 2024 election years together (SAY data lacks relevant outcomes so only PRL data included). Column 1 use a binary variable to indicate voter registration, Column 2 uses political interest, and Columns 3-4 use presidential vote intention for the the Democratic and Republican candidate respectively. The unit of observation is a respondent in a county pair in a election year. Standard errors are clustered at the DMA level. We use year by bi-month by county pair fixed effects. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$ (two-tailed test)

E Covariate Balance

In Appendix E, we test whether any changes in survey composition are correlated with advertising volume in ways that could confound our results. To do so, we use demographic characteristics that are largely stable and should not be causally affected by advertising exposure, so any relationship between them and ad volume would instead indicate imbalance over time. We regress ad volume on a series of demographic variables for each survey, where we binarize categorical individual-level demographic variables. In order to compare within the same border county pairs over time, we include only the county-pair unit fixed effects. Figure E5 and E6 plot these estimates over all outcomes, with p-values adjusted using False Discovery Rate correction. Across all six surveys, demographic composition is largely unrelated to advertising volume. A small number of coefficients remain statistically significant after correction, particularly in the 2020s, but their magnitudes are substantively small, suggesting these do not reflect meaningful imbalance that could confound our main results.

Figure E5: Covariate Balance on Residualized Treatment

(a) 2000

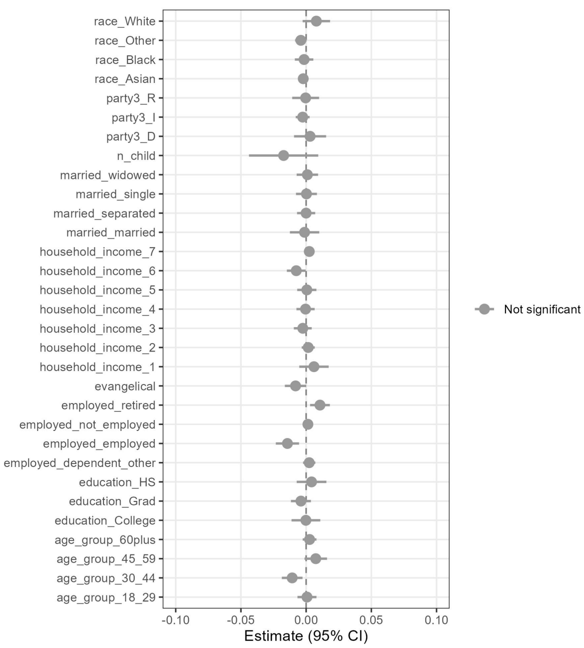

(b) 2004

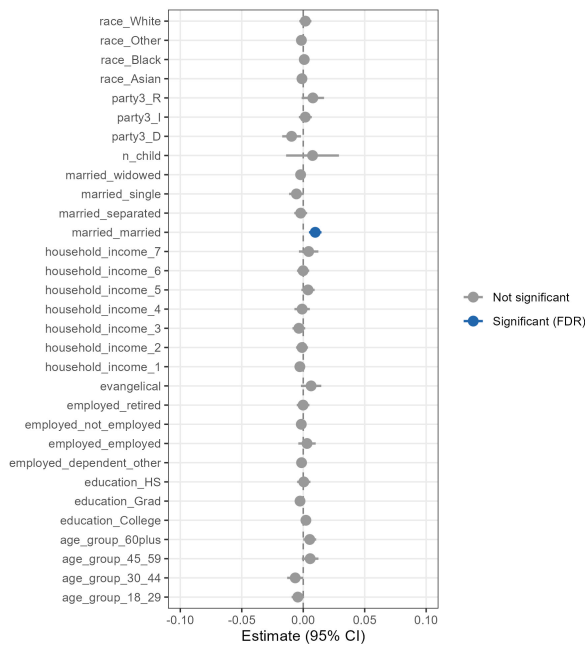

(c) 2008

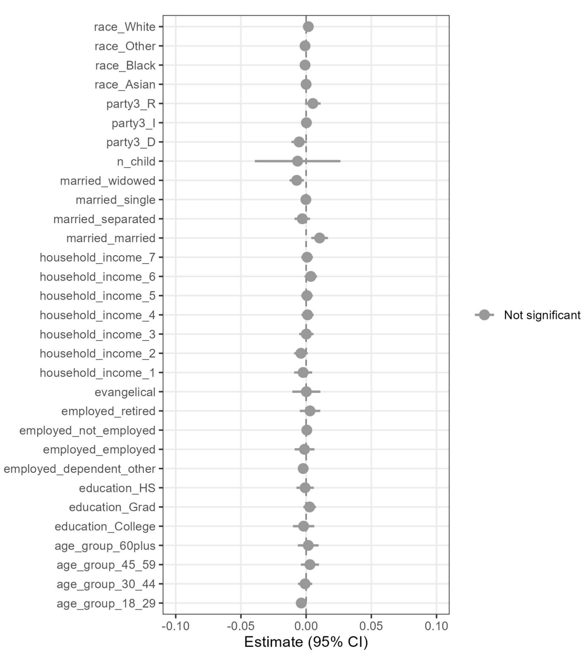

(d) 2020

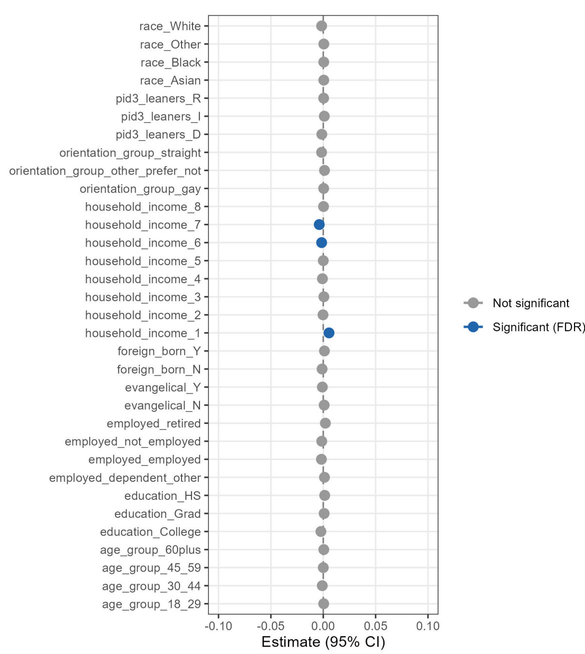

Note: This figure plots covariate balance tests across elections. Each panel indicates the residualized coefficient of the treatment on various demographic variables. We highlight covariates that are significant after correcting the p-values.

Figure E6: Covariate Balance on Residualized Treatment Cont.

(a) 2024 – PRL (b) 2024 – YouGov

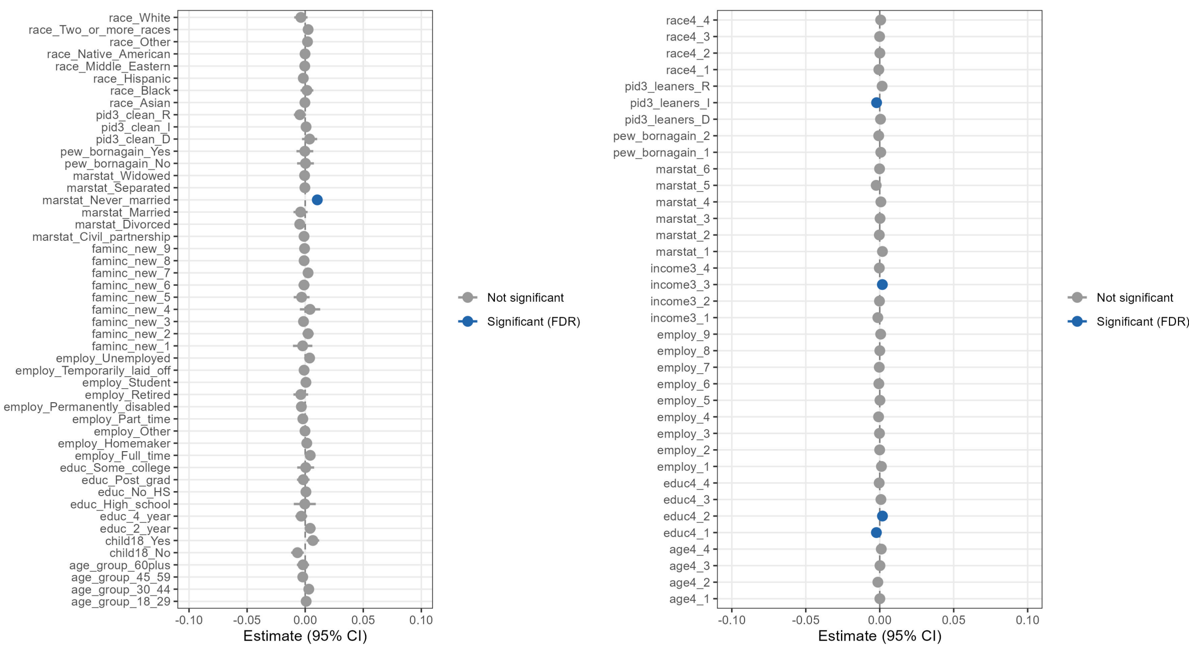

Note: This figure plots covariate balance tests across elections. Each panel indicates shows the residualized coefficient of the treatment on various demographic variables. We highlight covariates that are significant after correcting the p-values.

F Validation of Affect Measures

One potential critique of our analysis is that our measures of affective polarization varies across data sources. The NAES utilizes 101-point and 11-point feeling thermometers, Nationscape and The SAY YouGov surveys have 4-point favorability measures, and the PRL study employs a 101-point feeling thermometer. One concern is that the differences in effect sizes across time are being driven by differences in how affect is measured.

To guard against this concern, we calculate the correlation between our measures of partisan affect and plot their relationship. We utilize data from the 2016 George Washington University Team Module and 2020 Kent State University Team Module from the Cooperative Congressional Election Study (Gross 2019; Beam 2022). We make three comparisons: First, we find the correlation between the 101-point thermometer ratings and 5-point favorability measure for party. Second, we find the relationship between the 5-point favorabilty measure for party and candidate. Finally, we find the correlation between the 101-point thermometer ratings for party and candidate. This serves to validate our substantive comparisons between measures of affect in our main results.

We find that the correlation between all of our measures of affect are highly correlated with each other. Looking across measurement scales, the correlation coefficient between the feeling thermometers and favorability measures are 0.85 and 0.73 for Democrats and Republicans respectively. Looking at within-scale comparisons, measures of party and candidate affect also have strong relationships. We run a chi-square test on the 5-point favorability measures for party and candidates which produce a substantively large and statistically significant chi-square statistic. This suggests a strong relationship for both Democrats and Republicans. The same is true for the 101-point feeling thermometers. The correlation coefficient between party and candidate measures are 0.9 and 0.85 for Democrats and Republicans respectively. This adds credibility to our analysis where we make comparisons in effect sizes across survey studies where the specific measure of affect differs.

Figure F7: Correlation Between Party 101-Point Thermometer and 5-Point Favorability

(a) Democrats (b) Republicans

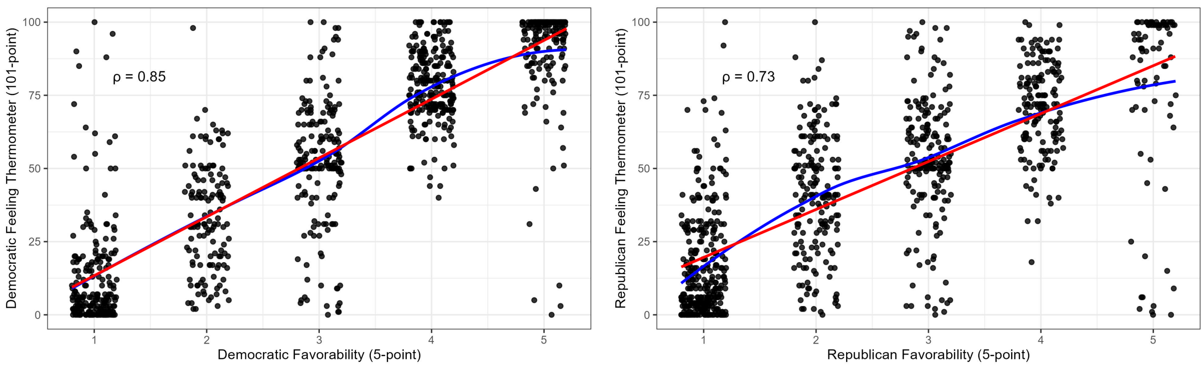

Note: This figure plots the correlation between the 101-point feeling thermometer and 5-point favorability measures for party. Each observation is an individual respondent. The red line shows a linear fit and the blue line shows a LOESS curve.

Figure F8: Correlation Between Party and Candidate 101-Point Thermometer

(a) Democrats (b) Republicans

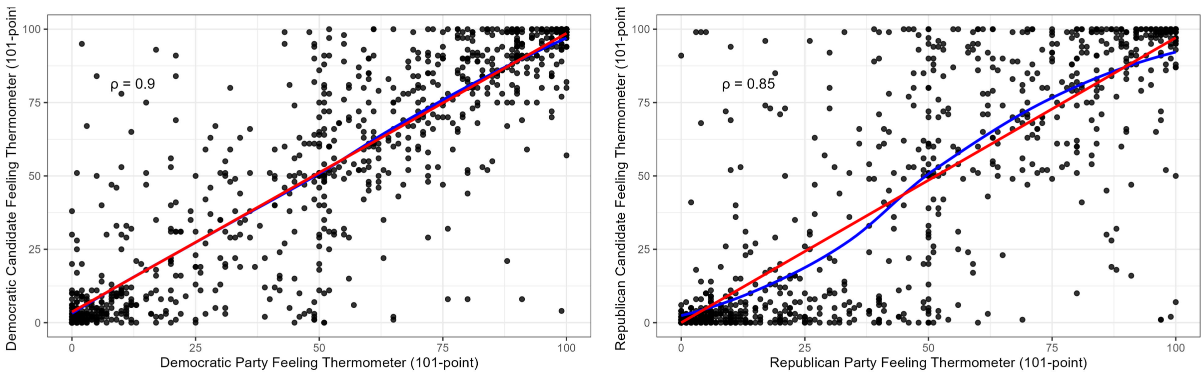

Note: This figure plots the correlation between the 101-point feeling thermometers for presidential candidate and party. We match Joe Biden to the Democratic Party and Donald Trump to the Republican Party. Each observation is an individual respondent. The red line shows a linear fit and the blue line shows a LOESS curve.

Figure F9: Correlation Between Party and Candidate 5-Point Favorability

(a) Democrats

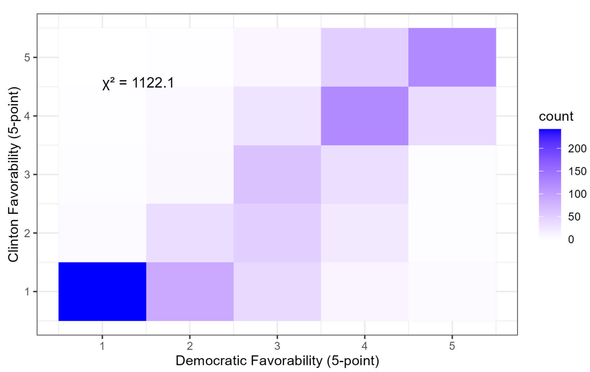

(b) Republicans

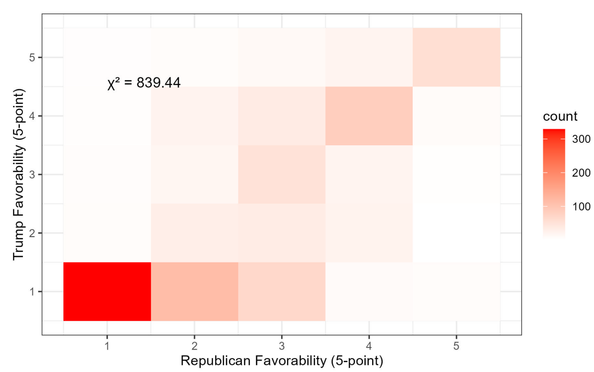

Note: This figure plots the correlation between the 5-point favorability measure for presidential candidate and party. We match Joe Biden to the Democratic Party and Donald Trump to the Republican Party. Each tile represents the total amount of respondents with that particular response combination.