

Do Voters Perceive Negative Campaigns as Informative Campaigns?

John Sides
Department of Political Science
George Washington University
jsides@gwu.edu

Keena Lipsitz
Department of Political Science
CUNY-Queens College
klipsitz@qc.cuny.edu

Matthew Grossmann
Department of Political Science
Michigan State University
matt@mattg.org

January 2009

forthcoming, *American Politics Research*

Abstract

We argue that citizens distinguish the tone of campaign from the quality of information that it provides, and that evaluations on each dimension respond differently to positive and negative political advertising. We test these claims using survey and advertising data from the 2000 presidential campaign and two 1998 gubernatorial races. In each race, citizens separate judgments about the tone of a campaign from judgments about the quality of information they have received. Furthermore, negative campaigning affects the former, but not the latter, set of evaluations. These results have implications for the debate over the impact of negative advertising and for how citizens perceive campaigns as political processes.

We thank Rick Lau, John Geer, an audience in the Department of Political Science at Temple University, and the anonymous reviewers for comments. The 2000 data are based on work supported by The Pew Charitable Trusts under a grant to the Brennan Center at New York University and a subsequent sub-contract to the Department of Political Science at the University of Wisconsin-Madison. The opinions expressed in this article do not necessarily reflect the views of the Brennan Center, the Campaign Media Analysis Group, the Wisconsin Advertising Project, or The Pew Charitable Trusts.

Introduction

Campaigns are key political moments because they bring citizens and leaders in closer contact and provide citizens opportunities to learn about and choose among those who want to represent them. Scholars have lauded campaigns for their ability to communicate information that helps citizens make voting decisions. At the same time, the conventional wisdom goes, Americans dislike campaigns, and negative campaigning—the criticisms that candidates make of one another—is a primary cause of this aversion.

Despite, or perhaps because of, the ubiquity of this conventional wisdom, we know little about how citizens view campaigns. Although it is easy to elicit derisive statements from citizens about campaigns, the mere existence of these sentiments does not mean that citizens' views are reducible to a simple and inevitable aversion, or that negative campaigning is to blame. The lack of scholarly consensus about campaign effects—such as the consequences of negative campaigning for voter participation—suggests that we have much to learn about how citizens view and respond to campaigns.

One promising way to address this gap is to broaden the question. Rather than investigating how campaign activities affect candidate evaluations or electoral participation, we investigate how citizens evaluate campaigns. In particular, we focus on how citizens evaluate the tone of campaigns and the information they provide. Are perceptions of tone and “informativeness” loosely connected in the minds of citizens, or are they more tightly linked, such that, for example, citizens who perceive the campaign as negative also perceive it as less informative? Moreover, how are perceptions related to the content of the campaign itself? As the campaign becomes more positive or negative, do citizens perceive that its tone has changed and, more importantly, do they reevaluate the quality of the information that the campaign is providing? Answers to these questions will

suggest an answer to a more fundamental question: can negative campaigns be informative campaigns?

These questions have both theoretical and normative implications. First, the question of how citizens feel about campaigns is arguably prior to the much-addressed, but largely inconclusive, relationship between negative campaigning and voter turnout. Knowing more about how citizens evaluate campaigns may help explain the lack of a conclusive answer. For example, although citizens may notice negative campaigning when it occurs, and label it as such, they may actually find it useful. If the information it conveys is perceived as relevant, then citizens' belief that the campaign is negative should not discourage them from voting. Our analysis thus speaks directly to a hypothesis that is often taken as fact, especially when scholars and journalists suggest that negative campaigning lowers turnout, but has rarely been subjected to direct empirical scrutiny: the more negatively candidates campaign, the less favorably citizens evaluate how they are campaigning.

Second, the results of our analysis speak to "revisionist" views of negative campaigning. In these accounts, negative campaigning provides citizens with valuable information about the candidates' views on key issues and with detailed evidence to back up these claims (Geer 2006). Negative campaigning thus facilitates comparisons between the candidates and, ultimately, a more informed choice on Election Day. If citizens exposed to negative campaigning actually find the campaign more informative, then these contrarian scholarly accounts will seem more persuasive than the conventional belief that citizens unequivocally disdain negative campaigning.

Finally, the content and sources of campaign evaluations will illuminate what citizens value in campaigns. While there is an extensive literature on people's attitudes toward government and governance (e.g., Hetherington 2005; Hibbing and Theiss-Morse 2002), campaigns have received far less attention in this respect, even though they are moments during which citizens pay unusually close attention to politics. Knowing how people evaluate campaigns is also important because those

evaluations have normative implications for how people feel about politics generally. Studying these evaluations may therefore suggest ways to improve campaigns as political processes.

The Structure of Campaign Evaluations

Little research speaks to the potentially numerous factors that might affect evaluations of campaigns; as West (2005: 43) notes, “Studies of the effects of ads have rarely paid much attention to the dimensions of evaluation.” This oversight may be explained by the fact that some influential research implies (but does not substantiate) a simple hypothesis: that citizens evaluate campaigns on a single dimension. In other words, a campaign’s many stimuli—ads, debates, media coverage, other events—are distilled into an overall global assessment. As Ansolabehere and Iyengar (1995: 110) write, “Viewers may learn from the mudslinging and name-calling that politicians in general are cynical, uncivil, corrupt, incompetent, and untrustworthy. Campaigns that generate more negative than positive messages may leave voters embittered toward the candidates *and* the rules of the game.” Thus, citizens’ distaste for negativity affects how they view the candidates, politicians, and perhaps politics in general, which leads to negative advertising’s demobilizing effect. Most important from our perspective, this research suggests that citizens’ (low) evaluations of the campaign’s tone “spill over” and affect their assessments of other aspects of the campaign.

We believe that citizens’ campaigns evaluations are more nuanced than this hypothesis suggests, and that they distinguish at least two dimensions of campaign communication: its tone and informativeness.¹ Tone captures how negative or positive the campaign is. Increasingly, commentary during elections—among candidates and journalists alike—revolves around whether and how much negativity exists (see Geer 2006). Candidates accuse each other of unfair attacks.

¹ To be sure, these are not the only possible dimensions on which citizens could evaluate campaigns, and it is not our goal to investigate all of the dimensions that citizens use of their own accord. Our claim is simply that tone and informativeness constitute two likely dimensions. More important to our argument is that citizens distinguish these two dimensions.

Journalists portray the campaign as a battlefield on which opposing sides lob slings and arrows. The increasingly extensive discourse about negativity has likely heightened citizens' attention to negativity in campaigns. In a 1996 Markle Foundation poll, "less negative campaigning" was the second most desired reform to campaigns. A bevy of other polling data suggests that citizens are more than willing to express opinions about negative campaigning, with most condemning it outright.² This suggests that citizens are taking note when candidates promote themselves or attack their opponents. The second dimension concerns an important function of campaigns: to provide information to citizens. Although most citizens approach elections with preexisting partisan leanings, issue positions, and often with at least some prior knowledge of the candidates, they express a desire to learn more (Lipsitz et al. 2005), and campaigns appear to help them do so (Arceneaux 2006; Freedman, Franz, and Goldstein 2004). In the Markle Foundation poll, the most desired reform was "more honesty and information from candidates." Thus, it seems likely that citizens are attentive to the quality of the information they receive in campaigns.

One might argue that citizens are not sophisticated or attentive enough to evaluate campaigns along more than one dimension. But such distinctions are evident with regard to other political objects. For example, citizens differentiate candidates in terms of multiple traits (Kinder 1986). This multi-dimensionality is facilitated by the second key fact: political attitudes are frequently characterized by ambivalence—that is, by the presence of both positive and negative dispositions toward the object of an attitude. Assessments of candidates are one important example (McGraw, Hasecke, and Conger 2003; Meffert, Guge, and Lodge 2004). If citizens can

² For example, in a February 2004 Pew survey, respondents were asked "How much does negative campaigning bother you?" Sixty-one percent said "very much," 20 percent said "somewhat," and only 18 percent said "not too much" or "not at all." In a June 2002 poll by the Institute for Global Ethics, 80 percent of respondents agreed that "negative attack-oriented campaigning is undermining and damaging our democracy" (with 46% saying "strongly agree"). In this same poll, vast majorities agreed that "negative, attack-oriented campaigning" is "making people less likely to go out to vote" (81%), is "unethical" (86%), and "produces leaders who are less trustworthy" (77%).

simultaneously believe that political figures have both good and bad qualities, then they may evaluate campaigns in a similar fashion.

How should assessments of the campaign's tone and informativeness be related? If people believe that negative campaigns are inherently lacking in substance, then a perception that the campaign is negative will be associated with a perception that it is uninformative. However, content analysis of campaign advertising suggests negative advertisements are in fact more likely than positive advertisements to focus on policy-related issues, to contain specific evidence for their claims, and to be "more *useful* than positive appeals in informing the public about the pressing concerns of the day" (Geer 2006: 63; see also Brader 2006; Mayer 1996; Jamieson, Waldman, and Sherr 2000; West 2005). Thus, a campaign might be simultaneously negative *and* informative. Citizens, and not just scholars, should be able to recognize this.

Moreover, some research on the processing of negative messages suggests that judgments about a campaign's tone and informativeness will be distinct. The theory of "affective intelligence" is particularly relevant (Marcus et al. 2000). Affective intelligence theory posits relationships between the tenor of messages, the emotional reactions those messages provoke, and the behavioral consequences of those emotions. More specifically, it argues that negative messages produce anxiety and fear, and that these emotions automatically activate the higher cognitive faculties involved in searching for more information.³ This search for information may extend well beyond the message that provoked the fear response to the larger information environment (Brader 2006: 135).

Thus, in the context of campaigning, the overall quality of a person's information environment, not simply the campaign's tone, should determine whether she believes the campaign is informative. If the broader environment provides useful information, then she may say that the

³ This argument fits well with other research that highlights the salience of negative information (Lau 1985) and how negative information stimulates cognition by better illuminating the potential costs and benefits of electoral choices (Kahn and Kenney 1999; Lau 1985).

campaign has been negative, but informative—exactly as content analyses of campaign advertising would predict. But if the information encountered is unhelpful, then she may find the campaign both negative and uninformative, a view that dovetails with conventional wisdom and contradicts accounts such as Geer's.⁴ If, over the course of a campaign, both outcomes result—varying, perhaps, over time or across messages or individuals—then in the aggregate voter judgments about tone and informativeness will be unrelated. Some research suggests such a null relationship because people are able to provide examples of both negative and positive advertisements that are informative and uninformative (Lipsitz et al. 2005; Freedman, Wood, and Lawton 1999). For this reason, these two dimensions of evaluation may not evince any systematic relationship.

Finally, how will each dimension of evaluation be related to the tone of the campaign? As the candidates “go negative” or take the high road, how do citizens’ evaluations respond? We first consider evaluations of the campaign’s tone itself. Some scholars have argued that citizens’ perceptions of negativity may not correspond to the actual tone of campaigns, at least as it is defined and measured by scholars (Sigelman & Kugler 2003). We expect a stronger relationship between the campaign’s tone and perceptions of negativity. Although citizens may disagree about whether a given piece of campaign information is negative, the pervasive discourse about negativity means they will more often agree about when the candidates are being negative. For example, it is difficult to argue that Lyndon Johnson’s infamous “Daisy” ad was positive or that Ronald Reagan’s “Morning Again in America” ad, which featured images of happy families moving into new homes and young couples getting married, was negative. People may disagree about whether negative campaigning is acceptable or fair (see Stevens et al. 2007), but that question is different from those addressed here.

⁴ Affective intelligence theory speaks more to the relationship between perceptions of negativity and informativeness than that between perceptions of positivity and informativeness. One possibility is that people exposed to positive messages feel more content with the information they already have because they have not experienced the anxiety that would lead to a search for additional information. They may say that the campaign has been informative because they have never found information lacking. Below, we examine positive and negative messages separately, which allows us to determine whether these types of messages have distinct effects.

We expect citizens' perceptions of the level of negativity to be systematically related to the actual level of negativity of campaign advertising.

It is less likely, however, that a campaign's tone will affect evaluations of informativeness. In general, although negative messages may be more informative than positive information, at least by certain criteria (Geer 2006), there remains considerable variability in the perceived informational value of negative and positive messages. To citizens, some negative messages may illuminate crucial shortcomings of the targeted candidate, but other negative messages may constitute shallow attacks. Some positive information may illuminate a candidate's agenda and issue positions, but other positive information may consist of vague statements about a candidate's values. Although citizens can distinguish between self-promotion and criticism, they will not find either systematically helpful or unhelpful. For example, Brooks and Geer (2007) find little difference in the informational value that subjects saw in positive ads and civil negative ads.

In sum, we expect people to evaluate campaigns along at least two dimensions—tone and informativeness—and the evaluations of these two dimensions should not be systematically related. Second, the actual level of negativity in a campaign should be clearly associated with perceptions of tone but not with perceptions of informativeness.⁵

Research Design and Data

To investigate these expectations, we need data that contain multiple indicators of campaign evaluations and that provide analytical leverage on the relationship between evaluations and campaign messages. We draw on public opinion data from three rolling cross-sectional surveys commissioned by the Annenberg School for Communication at the University of Pennsylvania. The

⁵ These expectations do not suppose that citizens pay close attention to the campaign or that they engage in a lot of effortful information processing. Instead, citizens are most likely merely forming impressions of the campaign from the limited information that they encounter. But even this kind of impression formation can still produce evaluations that differ across dimensions—much as it does when citizens are called upon to evaluate political candidates.

surveys were conducted during three campaigns: the 1998 California and Illinois gubernatorial races and the 2000 presidential race. The surveys in California and Illinois served as pilot studies for the considerably larger 2000 National Annenberg Election Study, or NAES (see Johnston, Hagen, and Jamieson 2004). Because the design of these surveys relies on daily interviews of separate cross-sections, their granularity is fine enough to track changes in opinion very closely and to capture the impact of campaign messages on opinion (see Johnston and Brady 2002; Romer et al. 2004). In California and Illinois, the Annenberg Survey sampled respondents in the San Francisco-Bay Area and Chicago media markets, respectively, from September 22 to November 2, 1998. The total sample size was 2,902 in California and 613 in Illinois. During the 2000 presidential race, the NAES was in the field for nearly a year. We rely on the subset of the NAES interviewed between July 18 and November 6, 2000 (N=34,391); this period of time captures the vast majority of general election campaigning.

Taken together, these surveys allow us to assess the structure and origins of evaluations across different types of elections (gubernatorial vs. presidential) and in different years (1998 vs. 2000). Even the two gubernatorial races were quite distinct. The California race featured the Democratic Lieutenant Governor, Gray Davis, and the Republican Attorney General, Dan Lungren. Davis won handily, garnering 60 percent of the vote. The Illinois race featured Democratic Congressman Glenn Poshard and Republican Secretary of State George Ryan. Ryan won the election with a narrow margin, garnering 52 percent of the vote. Ryan outspent Poshard nearly three to one but was hurt by a scandal. That the three Annenberg surveys were conducted during dissimilar races enhances the generalizability of our findings.⁶

⁶ A final question is whether results based on data from the 1998 and 2000 elections are generalizable to other election and more recent election in particular. We have no reason to suspect that our findings would prove unique to these elections.

All three surveys asked respondents to assess the campaign's negativity and informativeness. Before answering the questions on the CA and IL surveys, respondents were randomly assigned to one of three conditions. The first condition referred to "the candidates for governor" and did not specify either candidate by name. The second condition referred to the Republican (Lungren or Ryan), and the third to the Democrat (Davis or Poshard). The battery of questions began "Thinking about the [candidates for governor | governor's race], overall, would you say the [candidates for governor are | [name], the Republican, is | [name], the Democrat, is]..." The question wording of the specific measures was:

- "Useful Information": "...giving voters a great deal of useful information, some, not too much, or no useful information at all?"
- "Negative": "...conducting a campaign that is very, somewhat, not too, or not very negative?"
- "Criticizing": "How much time would you say the [candidates for governor are | Lungren is | Davis is] spending criticizing [their | his] opponent? A great deal of time, some, not too much, or aren't they doing this at all?"

All of the responses were recoded so that higher values indicate more favorable feelings toward the candidates—i.e., more useful information, less negative, and less time spent criticizing.

In the 2000 NAES, the last item (criticizing) was repeated verbatim. The other relevant items were not identical to those in the two pilot studies, but tap similar impressions:

- "Plans Once Elected": Thinking about what the candidates for president are saying, how much time would you say they are spending talking about what they themselves would do if elected? A great deal of time, some, not too much, or aren't they doing this at all?

- “Policy Issues”: Thinking about what the candidates for president are saying, how much time would you say they are spending talking about policy issues? A great deal of time, some, not too much, or aren’t they doing this at all?
- “Negative”: “Which of the major presidential candidates’ campaigns, if any, do you think has been too negative or nasty this year? George W. Bush’s, Al Gore’s, both or neither?”

The first two items speak to the potential informativeness of the campaign. “What they themselves would do if elected” and “policy issues” are the kinds of things that citizens deem most useful (Lipsitz et al. 2005). The last item—“negative or nasty”—provides another opportunity to observe whether citizens’ perceptions of the campaign tone derive from the actual campaign tone.⁷

To measure campaign messages, and in particular the campaign’s tone, we rely on televised political advertising. For California and Illinois, we purchased data originally collected by the Campaign Media Analysis Group (CMAG).⁸ For the presidential race, we draw on data from the Wisconsin Advertising Project, which collated, cleaned, and coded the CMAG data for public dissemination. For the California race, the CMAG data span the period from September 7 to November 2, 1998. During this period, Davis and Lungren aired a total of 3,609 advertisements in the San Francisco market.⁹ Davis aired the majority of these (1,847 vs. 1,762). For the Illinois race, the CMAG data span the period from July 7 to November 2, 1998. During this period, Ryan and Poshard aired 2,372 advertisements in the Chicago market, of which Ryan aired the vast majority

⁷ Each of these items was asked only of a random one-half of the sample. The “negativity” item was asked only beginning September 5.

⁸ Their technology tracks satellite transmissions from the major national networks and recognizes the digital “fingerprint” of various television programs and advertisements. With that fingerprint, CMAG then records when and where each political advertisement ran, and which candidate (or party or interest group) was its sponsor. Freedman and Goldstein (1999) were the first to demonstrate the power of this data source for political science.

⁹ In California, the CMAG data did not have clear labels for some ads and thus we could not code them for content. There were also several Spanish-language ads, for which CMAG did not record their content. Yet the vast majority of ads (83%) could be identified and coded. In Illinois, there were two Spanish-language ads whose content was not recorded. We were able to code 93 percent of the ads that aired in the Chicago market.

(1,735). During the 2000 presidential general election, Bush and Gore aired a total of 100,258 ads, with Bush airing somewhat more than Gore (55,286 vs. 44,972).¹⁰

To measure tone we coded each claim in the ad as referring to the candidate sponsoring the ad (a positive claim) or to the opponent (a negative claim). Coding the claim as opposed to the entire advertisement provides a more nuanced reading of the ad's content (see Geer 2006; Jamieson, Waldman, and Sherr 2000) and avoids slippery assessments of whether an entire ad can be categorized positive, negative, or some combination thereof (i.e., "contrast" ads).¹¹ Having coded the ads in this fashion, we counted the number of positive and negative claims that aired on each day (by multiplying the number of claims in each ad by the number of times the ad aired).

This combination of survey and advertising data is well-suited for exploring voter evaluations and their relationship to campaign information. First, the surveys ask respondents to evaluate the campaign in terms of several different attributes, allowing us to examine the relationship between evaluations of tone and informativeness. Second, the precise measurement of campaign advertising in respondents' media markets up to the day of their interview allows us to ascertain how evaluations are related to campaign messages. The time-series nature of the data makes it possible for us to observe campaign effects as they occur. Moreover, in the 2000 presidential race, markets in battleground states saw a great deal of advertising while markets in uncompetitive states saw very

¹⁰ In the presidential race, coordinated ads were considered candidate ads. We do not include party or independent ads because our focus is on perceptions of the candidates. In response to a reviewer's suggestion, we included cruder measures of the tone of party advertising, drawing on the Wisconsin Advertising Project's omnibus coding of each ad as positive, negative, or "contrast" and then assuming contrast ads were equally divided between positive and negative claims. Including these measures presumes that people conflate candidate ads and party ads and then evaluate the candidates based on both. When we included measures of positive and negative party advertising in the models in Table 2, we found two statistically significant effects: negative party ads are associated with less favorable evaluations of candidates on the "criticizing" dimension, while positive party ads are associated with less favorable evaluations on the "policy issues" dimension. The effects of candidate advertising remain similar to those reported in Table 2. The effects of the debates, candidate travel, and news coverage—which we discuss below—are generally weaker once party ads are included. Thus, the overall thrust of our results remains the same. Negative ads tend to be associated with evaluations of negativity but not evaluations of informativeness; positive ads have much weaker effects. In general, the effects of candidate ads are more numerous than those of party ads, suggesting that voters evaluate candidates based more on their own ads than those of their party.

¹¹ We each coded the ads separately. These initial assessments were all correlated at .90 or better. We then resolved any discrepancies to finalize the coding.

little, creating not only temporal but also geographic variation in the volume of campaign messages—a “natural experiment” of sorts (see Gimpel, Kaufman, and Pearson-Markowitz 2007; Johnston, Hagen, and Jamieson 2004). To our knowledge, no one has examined campaign evaluations and their relationship to advertising tone using comparable data.¹²

The Dimensionality of Campaign Evaluations

The dimensionality of campaign evaluations is first evident in whether the items tend to trend in similar directions over time. To track voter perceptions, we took daily averages of each of these indicators and then smoothed these averages to separate true opinion change from sampling fluctuation.¹³ Figure 1 presents the plots for the California and Illinois races. We include separate plots of the three conditions: the generic “candidates,” the Republican candidate, and the Democratic candidate. In each case, higher values indicate “better” campaign conduct—more informative and less negative or critical.

[insert Figure 1 about here]

Figure 1 first presents the trendlines for the “candidates.” At the survey’s outset, the trendlines are close to the mid-point of the scale (2.5), except for the perceived frequency of criticism. As the campaign progressed, respondents in both California and Illinois tended to believe that the candidates were becoming more negative and more critical. Yet the trend lines for

¹² Sigelman and Kugler (2003) analyze the California and Illinois races using the 1998 NES Pilot Study, which included a sample in each of three states, California, Illinois, and Georgia. Their measures of campaign content come from newspaper coverage and from a small sample of political ads from the University of Oklahoma’s Political Commercial Archive. The CMAG data provide a more complete collection of ads as well as information on when and, crucially, how often those ads aired (see Prior 2001). The NES included a single measure of voter perception (of the perceived negativity of the gubernatorial campaign) rather than the multiple measures included in the Annenberg data. Finally, the Annenberg data provide temporal leverage that a cross-section such as the 1998 NES Pilot Study cannot. As Sigelman and Kugler (2003: 154) note, candidate tone may vary significantly during a campaign, with consequences for voters’ perceptions. A potential disadvantage of the CA and IL data is that they were collected in select media markets, San Francisco and Chicago, that are not demographically representative of the rest of the state, although we do not believe that this would affect our hypothesis tests. Moreover, the 2000 NAES provides a nationally representative sample.

¹³ Specifically, we employed lowess procedure in Stata 9, with a bandwidth of 0.3.

informativeness (“useful information”) remain essentially flat in both states. Perceptions of Davis and Lungren display similar trends. Perceptions of Davis’ negativity and level of criticism became less favorable during the campaign, but perceptions of his informativeness do not become less favorable. Much the same conclusions can be drawn about perceptions of Lungren’s campaign.

The difference between perceptions of negativity and informativeness is evident in the over-time correlations among these trends, i.e., the correlations among the daily averages presented in Figure 1. For the California “candidates,” the correlation between perceptions of negativity and criticism ($r=.70$, $p<.001$) is much larger than the correlation between these measures and perceptions of informativeness ($r=.30$ and $r=.39$, respectively; $p<.001$). The results for evaluations of Davis and Lungren are even more striking. There are large correlations between perceptions of negativity and criticism (.70 for Davis; .68 for Lungren; $p<.001$), but smaller and statistically insignificant correlations between these perceptions and perceptions of informativeness (for Davis, -.01 and .10; for Lungren, $r=.19$ and .20).

A further test of the similarities and differences among these trends derives from a simple regression of the trendlines on the day of interview. These regressions reveal statistically significant declines for the measures of negativity and criticism in all three versions of the items. The “useful information” item exhibited no statistically significant decline in the case of the “candidates” and Davis, and only a very modest decline in the case of Lungren. In sum, as the 1998 campaign unfolded, Californians perceived the candidates as campaigning more negatively but not as providing less useful information.

Perceptions during the Illinois campaign tell a similar story. For the “candidates,” there is again a growing perception that the candidates are becoming more negative, but not necessarily that they are providing less useful information. In Poshard’s case, perceptions of negativity also decline somewhat but perceptions of informativeness follow a different pattern, particularly given the

increase at the campaign's end. In Ryan's case, perceptions on all three indicators initially become more favorable and then later less favorable. The trendline for informativeness manifests some of these trends, although in a more muted fashion. The correlations among these trendlines suggest, on average, stronger relationships between perceptions of negativity and criticism than between these perceptions and perceptions of informativeness. Similar regressions of the Illinois measures on day of interview again reveal that either the negativity item or the criticism item or both exhibit statistically significant declines, while the "useful information" item does not.

In Figure 2, we present similar plots for evaluations of the 2000 presidential candidates. In each plot, we also include vertical lines that denote the dates of the three presidential debates. Panel A presents trendlines for evaluations of the candidates on three dimensions: how much time they are talking about what they will do once elected, policy issues, and criticism of their opponent. Panel B presents the proportion of respondents who said that neither, both, or one or the other candidates were being negative.

Although these items are somewhat different than those included in the California and Illinois surveys, they still demonstrate the multi-dimensionality of evaluations. Panels A and B clearly show that over time respondents saw the candidates as more critical of each other and as more negative. As in California and Illinois, the vast majority believed that the candidates were spending some or a great deal of time criticizing each other, and this increased as the campaign wore on.¹⁴ This trend contrasts sharply with the *increasingly* favorable perceptions of the campaign's informativeness. Respondents saw the candidates as spending more, not less, time talking about their plans and about policy issues. Thus, in this race even more than in the two gubernatorial races, citizens saw no zero-sum trade-off between negative campaigning and informative discussion.

¹⁴ Interestingly, the proportion who believed Bush and/or Gore were being too negative is smaller than those who thought the candidates were spending some or much time criticizing each other. Of those who said that the candidates were spending some or a great deal of time criticizing each other, only 40 percent said that both candidates were being "too negative or nasty."

Another way to evaluate dimensionality is to examine each survey as a single cross-section and then examine relationships among these indicators across individuals, rather than over time. With only three indicators in the CA and IL surveys, a robust test of multi-dimensionality is not possible, but nevertheless the results of simple principal components models are instructive. In CA, separate models for each of the three sets of indicators (“candidates,” Davis, and Lungren) produce little evidence of a common underlying factor: the largest eigenvalues are less than 1 in the case of the candidates and Lungren, and only slightly over 1 for Davis. In IL, the same finding emerges, with an eigenvalue less than 1 for the “candidates,” and marginally greater than 1 for Poshard and Lungren. In the presidential data, the largest eigenvalue for the 4 indicators in Figure 2 is only .55.¹⁵ Thus, across these three races, these various evaluations of campaigns do not clearly tap a single underlying dimension.¹⁶

Ultimately, all of this evidence suggests that evaluations of campaign behavior are multi-dimensional. Over time, evaluations of campaign tone do not move in similar ways to those of campaign informativeness. Cross-sectionally, evaluations of a candidate’s negativity are distinct from evaluations of the information the candidate is providing. These dimensions appear to be separate constructs in citizens’ minds.

The Effects of Advertising Tone on Evaluations

To estimate the effect of advertising tone on perceptions, we combined the measures of positive and negative appeals with the survey data, and then constructed a model where perceptions are a function both of tone and relevant individual-level covariates. Our measures of tone are the

¹⁵ For this analysis, the item capturing perceptions of negativity is coded in three ordinal categories (both, Bush or Gore, and neither).

¹⁶ Further analysis shows that, as in the over-time analysis, the correlations between perceptions of negativity and criticism tend to be larger than those between these perceptions and those of informativeness (data not shown).

cumulative number of positive and negative appeals in the seven days before the respondent was interviewed.¹⁷

These models also include several individual-level controls. For perceptions of the “candidates,” our model includes partisanship (a “folded” version of the seven-category party identification scale), political information, political efficacy, as well as race and gender. Sigelman and Kugler (2003) demonstrate that perceptions of negativity are greater among those who are less efficacious, less partisan, but more politically attentive.¹⁸ Our model for perceptions of the specific candidates substitutes party identification for partisanship. To measure political information in the CA and IL data, we use recall of candidate names because these data contain no general measure of political information. In the presidential data, we use the interviewers’ assessment of the respondent’s political information. Efficacy is a scale combining responses to 3 items about one’s ability to affect government and comprehend politics.¹⁹

The models are OLS regressions, except in one case, perceptions of negativity in the presidential data, for which we employ an ordered probit model. In each model, we employ robust standard errors that reflect the “clustered” nature of the data (by day of interview in California and Illinois, and by media market in the presidential data). We estimate a model of each indicator of

¹⁷ In models of the “candidates” conduct, we combine the Republican and Democratic candidates’ advertising into a single measure. We divided these seven-day measures by 1,000 to make their scale more compact. Johnston, Hagen, and Jamieson (2004) find that the sum of the previous seven days’ advertising has the strongest impact on individual-level attitudes. Because the CMAG presidential data covered only the top 75 media markets, we include only respondents who lived in those markets, which is about 75% of the NAES sample.

¹⁸ The effect of advertising could be greater among voters who are more politically attentive. We investigated whether individual-level measures of political attentiveness and media consumption moderate the effects we describe below. We found little evidence that such measures play a robust moderating role. These results are available on request. Although party identification is associated with evaluations of the specific candidates—e.g., Democrats tend to evaluate Gray Davis’ campaign conduct more favorably—party identification does not moderate the effect of advertising messages in alternative model specifications. As the campaign wore on, Republicans and Democrats did not necessarily diverge in their assessments of how negatively the two candidates are campaigning. The trends in perceptions of negativity reveal no consistent pattern of increasing partisan polarization during the CA and IL campaigns. In fact, there is some evidence of partisan convergence in Illinois. These results are also available upon request.

¹⁹ In the presidential data, the items are “People like me have no say over party choices for presidential candidates”; “People like me have no say over who is president”; and “Presidential elections are too complicated to understand.” In the CA and IL data, the items are “State public officials don’t care much what people like me think”; “Sometimes state politics and government seem so complicated that a person like me can’t really understand what is going on”; and “People like me don’t have any say about what the state government does.”

candidate conduct, with separate sets of models in California and Illinois for the “candidates,” the Republican (Lungren or Ryan), and the Democrat (Davis or Poshard).

Table 1 presents the results of the models from the California and Illinois races. We report only the effects of the advertising measures but the full results are available in Appendices 1-2. Most generally, these models suggest a significant association between advertising tone and evaluations of the candidates. Negative appeals matter more than positive appeals, suggesting that negative information may be more salient to respondents. The effects of negative appeals, however, are not consistent across the different evaluations, as we hypothesized. In nearly every set of models, there is a significant relationship between the actual volume of negative appeals and respondent’s perceptions of negativity (either the “negative” or “criticizing” indicators, and sometimes both).

[insert Table 1 about here]

An illustration of the substantive magnitude of the relationships in Table 1 comes from the model of negativity for the CA “candidates.” Comparing respondents interviewed when the number of negative appeals was largest to those when interviewed when the number of negative appeals was smallest, we would expect their evaluations of the candidates’ negativity to become less favorable by approximately half of a standard deviation. Thus, respondents interviewed at a time when numerous ads included criticism of the opponent tend to believe that the candidates are spending more time criticizing each other and are conducting a more negative campaign. Public perceptions of negativity do in fact respond to reality.

Finally, as hypothesized, the tone of campaign messages is not consistently related to perceptions of informativeness. Most notably, there is no significant evidence that negative messages affect perceptions of informativeness. In only one case is the coefficient statistically significant (the Poshard model); here, it appears that respondents interviewed amidst a larger volume of negative appeals are less inclined to deem Poshard’s campaigning informative. A similar, though

marginally significant relationship, emerges in the models of Ryan's informativeness ($p=.11$), but a marginally *positive* relationship emerges in models of Davis' informativeness ($p=.11$, one-tailed). The other models show no relationship between campaign negativity and perceptions of informativeness. Taken together, these results confirm our expectation that negative campaign messages have no consistent effect on perceptions of informativeness.²⁰

Table 2 presents the results for the 2000 presidential campaign. The first set of models includes the two measures of advertising appeals as well as the individual-level controls (which are not shown; see Appendix 3). The findings from this campaign confirm those from the California and Illinois campaigns. First, the volume of negative appeals is strongly associated with the belief that the candidates are being negative and are criticizing each other. Curiously, the volume of positive appeals manifests a similar relationship, though one that is weaker in magnitude. The mere presence of campaign advertising—whether negative or positive—appears to produce less favorable perceptions of the candidates on this dimension. Negative advertising, however, has a more substantial effect.

[insert Table 2 about here]

Second, negative appeals do not lead respondents to perceive that the candidates are spending more or less time talking about policy issues or their plans once elected. The effect of negative appeals on these dimensions is not statistically significant. Unlike in the results from the California and Illinois races, however, positive appeals do affect evaluations of informativeness: the more positive advertising in the week before the interview, the more respondents feel that the candidates are talking about their issues and plans.

²⁰ Some research suggests the value in disaggregating negative appeals, such as by comparing appeals focused on issues to appeals focused on candidate traits (e.g., Fridkin and Kenney 2004). In this case, one might expect that trait-based attacks would be perceived as more negative and less informative than issue-based attacks. We investigated this distinction but found no systematically different effects of these two kinds of attacks. We also attempted to distinguish civil from uncivil attacks, using the definition proposed by Brooks and Geer (2007), but in these races no such attacks met their standards for incivility.

Of course, this model specification fails to capture the rich information environment surrounding the presidential campaign. Perceptions of the campaign may also derive from news coverage and candidate debates. To capture this information, we added these indicators to the models: first, the cumulated number of candidate visits to the respondent's state as of the date of interview, which proxies the amount of local news coverage of the campaign; second, the number of minutes devoted to the campaign on the three major network evening news broadcasts during the week before the date of interview; and third, a set of dummy variables to capture whether respondents were interviewed the week after one of the three presidential debates.²¹

The second model in Table 2 presents the results including these additional measures of campaign information. The central findings from the first model do not change: the volume of negative appeals is still associated with less favorable evaluations on two dimensions (criticizing, negative) but not with the others. The effects of positive appeals are weaker in magnitude and no longer statistically significant. The measures of visits and news have few effects of note, although the volume of news coverage is associated with the perception that the candidates are campaigning negatively. This dovetails with Geer's (2006) portrayal of the news as focused on negativity. Finally, respondents interviewed after the debates also manifest differences, relative to respondents interviewed earlier—although these differences do not conclusively prove that the debates themselves mattered. They are more likely to say that the candidates are negative and are spending more time criticizing each other. They are not consistently any more or less likely to say that the candidates are discussing their plans once elected or policy issues.²²

²¹ Richard Johnston graciously supplied the data on candidate visits. The news data were collected by the authors from the Vanderbilt Television News Archive. With regard to the debates, the NAES data do include self-reports of debate exposure, but these are afflicted with a large amount of over-reporting (Prior 2005). We prefer the dummy variables for this reason, and also because they help capture the “debate about the debate” that typically occurs after the debate itself. However, this measurement strategy has important limitations, as there was obviously other campaign activity that took place in the week after each debate. We must therefore be cautious in interpreting the effects of these variables.

²² We also attempted to distinguish trait- and issue-based negativity in presidential advertising. In the 2000 race, only about 11% of negative appeals mentioned candidate traits, limiting the usefulness of these data for parsing the effects of

In sum, the results from these three races confirm our expectations. Different evaluations of campaigns respond in distinct ways to positive and negative appeals. As the candidates air more negative appeals, citizens tend to perceive the campaign as more negative and the candidates as spending more time criticizing each other. Yet in almost every case there was no relationship between the volume of negative appeals and beliefs about whether the candidates were providing useful information or discussing policy issues. Negative campaigning does not lead citizens to evaluate the campaign less favorably on every dimension.

Discussion and Conclusion

Although our results derive from disparate campaigns, they generate a coherent set of conclusions. First, citizens separate evaluations of a campaign's tone from evaluations of how much useful information the campaign is providing. Second, the relationship between these dimensions of evaluation is modest at best. Citizens who believe a campaign is negative will not necessarily find it uninformative. Third, these multiple evaluations respond differently to shifts in the campaign's tone. Citizens discern when campaigns are negative; they are more likely to call the campaign "negative" when interviewed amidst a barrage of negative ads. Thus, the weak link between negative advertising and participation does not arise because citizens define negativity differently than do scholars. Instead, our results suggest that the weak link arises because negative campaigning does not lead citizens to evaluate the campaign as uninformative: evaluations of informativeness are only minimally and inconsistently tied to the tone of advertising. People's evaluations of negativity are thus somewhere in between the alarmist sentiments of some observers and the revisionist accounts of some scholars. On the one hand, citizens do not respond to negativity by disdaining campaigns

each type of negative appeal. Similar models to those in Table 2 suggest that issue-based negativity effects are much larger than those of trait-based negativity, but again we do not place too much stock in these results, given the paucity of trait appeals. Finally, as in our analysis of the CA and IL races, we found no evidence of incivility in these presidential ads, which prevents us from distinguishing appeals on this dimension.

generally, as suggested by some popular commentators. On the other, they do not consistently appreciate the benefits of negativity touted by some scholars.

These results have implications for research on negativity campaigning. The conventional story about the consequences of negative campaigning often takes this form: candidates criticize each other, which leads to “voter alienation” from the political process, which leads to lower turnout on Election Day (Ansolabehere and Iyengar 1995). This story is at odds, however, with a significant amount of empirical literature that finds either a positive effect of negative campaigning on turnout or, in the most comprehensive review, no consistent evidence of an effect (Lau, Sigelman, and Rovner 2007). Our findings suggest a reason why existing research on negative campaigning is so inconclusive. Although citizens do not necessarily “defend” negativity as much as some political scientists (Geer 2006; Mayer 1996), they sometimes find merit in negative campaigning. Thus, it makes sense that negative campaigning does not consistently produce alienation or lower turnout, or, for that matter, enthusiasm or higher turnout.²³

Extant research has devoted much less attention to how citizens evaluate campaigns and, more specifically, negative campaigning—a step arguably prior to “downstream” consequences such as turnout on Election Day. The multi-dimensionality of evaluations suggests that citizens make distinctions between helpful and unhelpful negative campaigning. Thus, a central task for future research is to “unpack” negativity to identify those aspects of negative campaigning that citizens find objectionable and those they find helpful. Recent research attempting to distinguish negative messages in terms of their relevance and civility is a promising step in this direction (Brooks and Geer 2007; Fridkin and Kenney 2006). Ultimately, future research could marry this approach and

²³ Unfortunately, the data we analyze here do not allow a credible exploration of turnout. In the pre-election survey, there is a question asking respondents whether they intend to vote, but this generates the usual socially desirable response—e.g., in the 2000 NAES data, 85% of respondents intended to vote. Analysis using this measure as a dependent variable uncovered finds no consistent statistically significant relationship between any measure of campaign evaluation and turnout intention, or between positive or negative campaign appeals and turnout intention. But we are hesitant to make much of these results given the dependent variable’s flaws.

ours, disaggregating *both* evaluations and messages into their constituent dimensions and looking for systematic relationships among these dimensions. A particularly important question is how citizens evaluate campaign message in terms of their informativeness. What constitutes useful information for citizens, and does that vary among citizens? Improving our understanding of campaign communication and the intricate ways in which citizens react to it will help us understand the true relationship between campaign practices and citizen satisfaction with the electoral process, which is an important outcome in its own right.

Our results also suggest several positive normative implications. We find it heartening that citizens evaluate campaigns in a more nuanced fashion than previously thought and change their evaluations of campaigns as they unfold. Citizens appear to respond to campaign advertising. They sometimes express aversion, sometimes acknowledge learning, and most often respond with a mixture of both. They do not always accentuate the negative in their evaluations and are thus not uniformly unfavorable toward campaigns, even though they are attentive to the tone of messages that candidates present to them.

Finally, our results also advise against categorical judgments for or against negative campaigning. Citizens may often express grievances about negativity in campaigns but these off-the-cuff complaints conceal a more nuanced evaluation of campaign conduct. In their own evaluations, scholars should not be any less nuanced. As we learn more about the multiple and often conflicting reactions that citizens have to campaigns, we may discover specific messages that both inform and engage citizens. Scholars should work to identify the kinds of messages that provide the information that citizens need along with the level of civility that they want.

Table 1. Models of Campaign Evaluations in the 1998 CA and IL Gubernatorial Races

	Not criticizing	Not negative	Useful Information
CA: Candidates			
Volume of positive appeals	0.03 (0.07)	0.01 (0.05)	0.06 (0.06)
Volume of negative appeals	-0.28** (0.08)	-0.45** (0.12)	-0.03 (0.06)
CA: Davis			
Volume of Davis' positive appeals	0.10 (0.08)	0.01 (0.08)	0.03 (0.09)
Volume of Davis' negative appeals	-0.47** (0.16)	-0.43** (0.13)	0.19 (0.15)
CA: Lungren			
Volume of Lungren's positive appeals	0.01 (0.10)	-0.02 (0.11)	0.10 (0.14)
Volume of Lungren's negative appeals	-0.15 (0.11)	-0.10 (0.10)	-0.02 (0.13)
IL: Candidates			
Volume of positive appeals	-0.04 (0.09)	0.06 (0.11)	-0.20 (0.14)
Volume of negative appeals	-0.08 (0.09)	-0.42** (0.10)	0.01 (0.09)
IL: Poshard			
Volume of Poshard's positive appeals	0.35 (0.54)	-0.45 (0.72)	0.63 (0.46)
Volume of Poshard's negative appeals	-1.16* (0.50)	-0.91 (0.67)	-1.07* (0.40)
IL: Ryan			
Volume of Ryan's positive appeals	0.23 (0.32)	-0.03 (0.18)	0.11 (0.17)
Volume of Ryan's negative appeals	-0.16 (0.16)	-0.40* (0.18)	-0.21 (0.14)

Cell entries are unstandardized regression coefficients, with robust standard errors in parentheses. Dependent variables are coded 1-least favorable to 4-most favorable. Each model also includes the individual-level control variables discussed in the text.

**p<.01; *p<05 (two-tailed).

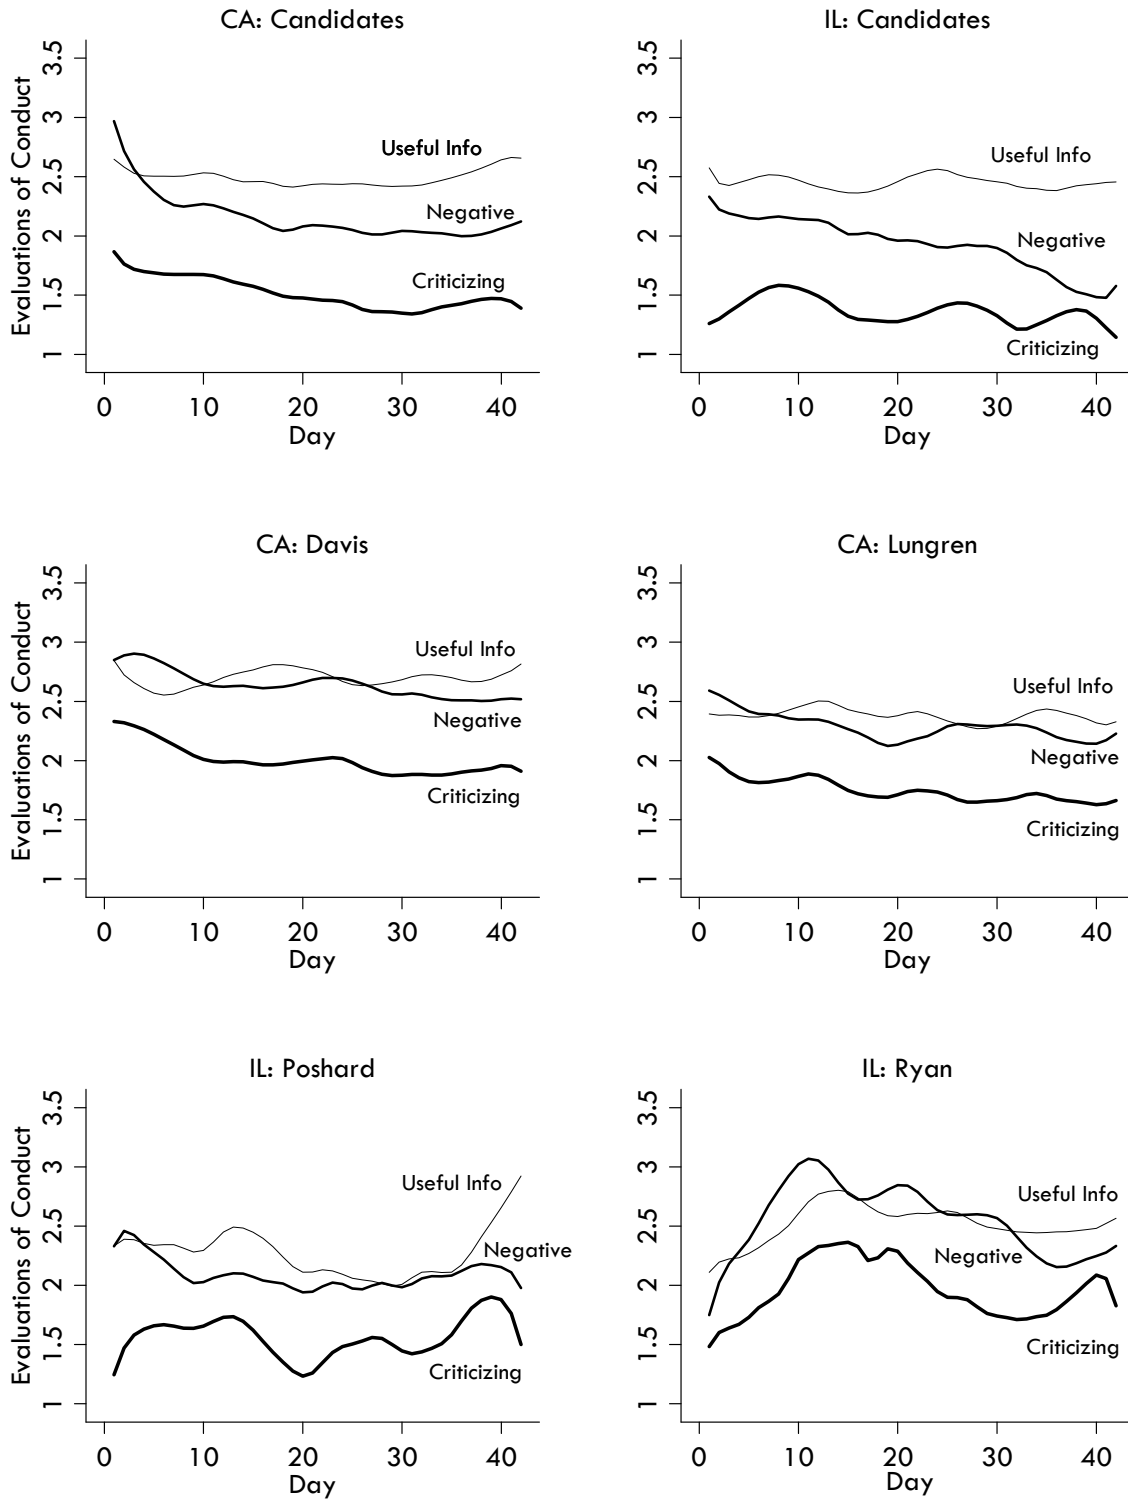
Table 2. Models of Campaign Evaluations in the 2000 Presidential Race

	Not criticizing	Not negative	Plans once elected	Policy issues
Model 1				
Volume of positive appeals	-0.002 (0.001)	-0.005* (0.002)	0.003* (0.001)	0.004* (0.002)
Volume of negative appeals	-0.019** (0.003)	-0.016** (0.005)	0.004 (0.003)	0.002 (0.004)
Model 2				
Volume of positive appeals	0.001 (0.002)	-0.004 (0.002)	0.003 (0.002)	0.003 (0.002)
Volume of negative appeals	-0.021** (0.003)	-0.015** (0.005)	0.005 (0.003)	0.004 (0.004)
Number of candidate visits	-0.012 (0.006)	-0.012 (0.014)	-0.003 (0.006)	-0.006 (0.007)
Number of minutes of news	0.001 (0.011)	-0.168** (0.048)	-0.011 (0.009)	-0.022 (0.012)
Week after first debate	-0.037 (0.023)	-0.070 (0.044)	0.066 (0.034)	0.057* (0.025)
Week after second debate	-0.057 (0.031)	-0.136* (0.060)	-0.097** (0.035)	0.029 (0.038)
Week after third debate	-0.144** (0.027)	-0.120** (0.044)	0.026 (0.022)	0.022 (0.027)

Cell entries are unstandardized OLS, ordered probit, or logit coefficients, with robust standard errors in parentheses. The dependent variables “criticize,” “promote”, and “policy” are coded 1-least favorable to 4-most favorable; these models are estimated using OLS. “Negative” is coded 0-both, 1-Bush or Gore, and 2-neither; this model is estimated using ordered probit. Each model also includes the individual-level control variables discussed in the text.

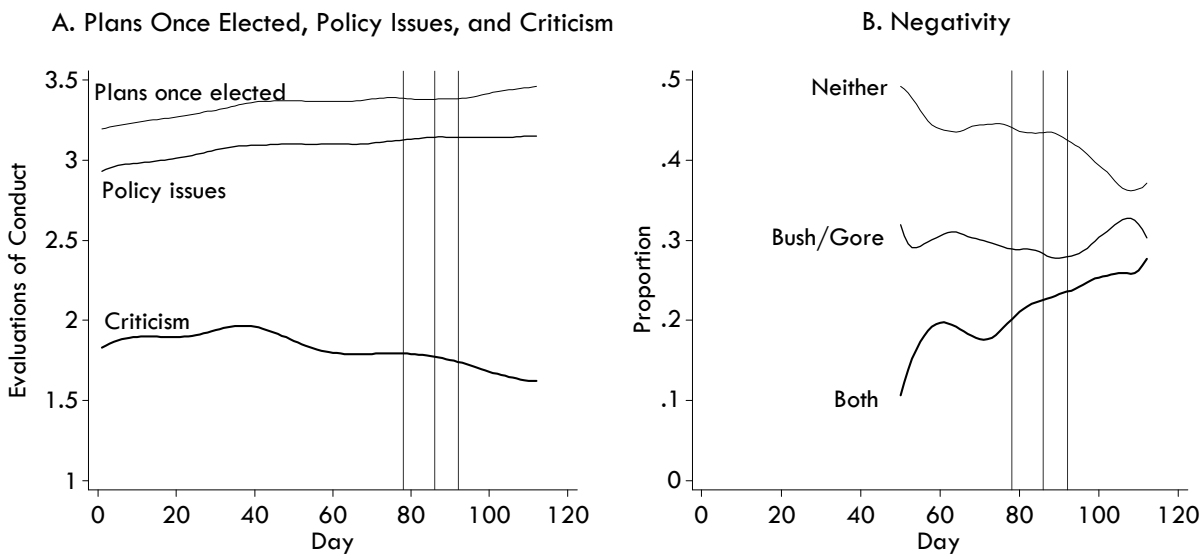
**p<.01; *p<05 (two-tailed).

Figure 1. Evaluations of Campaign Conduct in the 1998 CA and IL Gubernatorial Races



The lines are the smoothed trends in voter evaluations. These measures are coded 1 (negative) to 4 (positive). The data begin on September 22 and conclude on November 2.

Figure 2. Evaluations of Campaign Conduct in the 2000 Presidential Race



The lines are the smoothed trends in voter evaluations. These measures in Panel A are coded 1 (negative) to 4 (positive). The data begin on July 18 and conclude on November 6. The three vertical lines denote the presidential debates.

Appendix 1. Models of Campaign Evaluations in the 1998 California Gubernatorial Race

	Not criticizing	Not negative	Useful Information
Candidates			
Volume of positive appeals	0.03 (0.07)	0.01 (0.05)	0.06 (0.06)
Volume of negative appeals	-0.28** (0.08)	-0.45** (0.12)	-0.03 (0.06)
Strength of partisanship	-0.01 (0.03)	-0.06 (0.04)	0.12** (0.02)
Political information	0.01 (0.02)	0.06* (0.03)	0.06* (0.03)
Efficacy	0.06** (0.02)	0.05* (0.02)	0.07** (0.02)
Constant	1.66** (0.12)	2.48** (0.18)	1.97** (0.12)
R ²	0.03	0.05	0.05
N	937	930	943
Davis			
Volume of Davis' positive appeals	0.10 (0.08)	0.01 (0.08)	0.03 (0.09)
Volume of Davis' negative appeals	-0.47** (0.16)	-0.43** (0.13)	0.19 (0.15)
Party identification (high=Dem)	0.13** (0.01)	0.16** (0.02)	0.12** (0.02)
Political information	-0.01 (0.04)	-0.05* (0.03)	0.08** (0.03)
Efficacy	0.01 (0.02)	0.05* (0.02)	0.03 (0.03)
Constant	1.74** (0.15)	2.17** (0.13)	1.84** (0.17)
R ²	0.12	0.15	0.12
N	800	818	811
Lungren			
Volume of Lungren's positive appeals	0.01 (0.10)	-0.02 (0.11)	0.10 (0.14)
Volume of Lungren's negative appeals	-0.15 (0.11)	-0.10 (0.10)	-0.02 (0.13)
Party identification (high=Dem)	-0.10** (0.01)	-0.12** (0.01)	-0.13** (0.01)
Political information	-0.12** (0.03)	-0.17** (0.03)	-0.06 (0.04)
Efficacy	0.02 (0.02)	0.04 (0.02)	0.04 (0.03)
Constant	2.34** (0.13)	2.93** (0.11)	2.83** (0.12)
R ²	0.10	0.13	0.11
N	784	803	798

Cell entries are unstandardized regression coefficients, with robust standard errors in parentheses. Dependent variables are coded 1-least favorable to 4-most favorable. Each model also includes dummy variables for gender and for ethnicity (white vs. non-white).

**p<.01; *p<05 (two-tailed).

Appendix 2. Models of Campaign Evaluations in the 1998 Illinois Gubernatorial Race

	Not criticizing	Not negative	Useful Information
Candidates			
Volume of positive appeals	-0.04 (0.09)	0.06 (0.11)	-0.20 (0.14)
Volume of negative appeals	-0.08 (0.09)	-0.42** (0.10)	0.01 (0.09)
Strength of partisanship	0.06 (0.03)	-0.09 (0.05)	0.05 (0.05)
Political information	-0.10* (0.05)	-0.04 (0.04)	-0.09 (0.06)
Efficacy	0.07 (0.04)	0.02 (0.06)	0.03 (0.05)
Constant	1.32** (0.14)	2.35** (0.23)	2.50** (0.26)
R ²	0.08	0.10	0.02
N	211	209	207
Poshard			
Volume of Poshard's positive appeals	0.35 (0.54)	-0.45 (0.72)	0.63 (0.46)
Volume of Poshard's negative appeals	-1.16* (0.50)	-0.91 (0.67)	-1.07* (0.40)
Party identification (high=Dem)	0.05 (0.03)	0.03 (0.03)	0.06 (0.03)
Political information	-0.20* (0.07)	-0.08 (0.08)	-0.15* (0.07)
Efficacy	0.01 (0.06)	0.08 (0.05)	-0.07 (0.07)
Constant	1.84** (0.23)	2.16** (0.21)	2.56** (0.26)
R ²	0.10	0.10	0.12
N	164	168	169
Ryan			
Volume of Ryan's positive appeals	0.23 (0.32)	-0.03 (0.18)	0.11 (0.17)
Volume of Ryan's negative appeals	-0.16 (0.16)	-0.40* (0.18)	-0.21 (0.14)
Party identification (high=Dem)	-0.17** (0.02)	-0.14** (0.04)	-0.13** (0.03)
Political information	-0.01 (0.07)	-0.05 (0.07)	-0.01 (0.08)
Efficacy	0.01 (0.05)	0.03 (0.05)	0.02 (0.06)
Constant	2.36** (0.27)	3.23** (0.23)	2.78** (0.26)
R ²	0.17	0.13	0.13
N	170	172	168

Cell entries are unstandardized regression coefficients, with robust standard errors in parentheses. Dependent variables are coded 1-least favorable to 4-most favorable. Each model also includes dummy variables for gender and for ethnicity (white vs. non-white).

**p<.01; *p<.05 (two-tailed).

Appendix 3. Models of Campaign Evaluations in the 2000 Presidential Race

	Not criticizing	Not negative	Plans once elected	Policy issues
Volume of positive appeals	0.001 (0.002)	-0.004 (0.002)	0.003 (0.002)	0.003 (0.002)
Volume of negative appeals	-0.021** (0.003)	-0.015** (0.005)	0.005 (0.003)	0.004 (0.004)
Number of candidate visits in state	-0.012 (0.006)	-0.012 (0.014)	-0.003 (0.006)	-0.006 (0.007)
Number of minutes of news	0.001 (0.011)	-0.168** (0.048)	-0.011 (0.009)	-0.022 (0.012)
Week after debate 1	-0.037 (0.023)	-0.07 (0.044)	0.066 (0.034)	0.057* (0.025)
Week after debate 2	-0.057 (0.031)	-0.136* (0.060)	-0.097** (0.035)	0.029 (0.038)
Week after debate 3	-0.144** (0.027)	-0.120** (0.044)	0.026 (0.022)	0.022 (0.027)
Strength of partisanship	0.030** (0.007)	0.018 (0.015)	0.029** (0.008)	0.060** (0.009)
Political information	0.011 (0.007)	-0.023 (0.014)	-0.014 (0.007)	-0.027** (0.007)
Efficacy	0.109** (0.009)	0.157** (0.020)	0.013 (0.010)	0.054** (0.011)
Constant	1.449** (0.038)		3.290** (0.041)	2.950** (0.040)
Adjusted R-squared	0.02		0.01	0.02
N	12293	6829	12230	12124

Cell entries are unstandardized OLS, ordered probit, or logit coefficients, with robust standard errors in parentheses. The dependent variables “criticize,” “promote”, and “policy” are coded 1-least favorable to 4-most favorable; these models are estimated using OLS. “Negative” is coded 0-both, 1-Bush or Gore, and 2-neither; this model is estimated using ordered probit. Each model also includes dummy variables for gender and for ethnicity (white vs. non-white).

**p<.01; *p<05 (two-tailed).

Bibliography

- Ansolabehere, Stephen, and Shanto Iyengar. (1995). *Going Negative: How Political Advertisements Shrink and Polarize the Electorate*. New York: Free Press.
- Arceneaux, Kevin. (2006). Do Campaigns Help Voters Learn? A Cross-National Analysis. *British Journal of Political Science*, 36, 159-173.
- Brader, Ted. (2006). *Campaigning for Hearts and Minds: How Emotional Appeals in Political Ads Work*. Chicago: Chicago University Press.
- Brooks, Deborah Jordan, and John G. Geer. (2007). Beyond Negativity: The Effects of Incivility on the Electorate. *American Journal of Political Science*, 51, 1-16.
- Freedman, Paul, and Kenneth Goldstein. (1999). Measuring Media Exposure and the Effects of Negative Campaign Ads. *American Journal of Political Science*, 43, 1189-1208.
- Freedman, Paul, Michael Franz, and Kenneth Goldstein. (2004). Campaign Advertising and Democratic Citizenship. *American Journal of Political Science*, 48, 723-741.
- Freedman, Paul, William Wood, and Dale Lawton. (1999). Do's and Don'ts of Negative Ads: What Voters Say. *Campaigns and Elections*, Oct-Nov, 20.
- 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, 570-605.
- (2006). *Do Negative Advertisements Work: Exploring the Dimensions of Negativity*. Paper presented at the annual meeting of the American Political Science Association, Philadelphia, PA.
- Geer, John G. (2006). *In Defense of Negativity: Attacks Ads in Presidential Campaigns*. Chicago: University of Chicago Press.
- Gimpel, James G., Karen M. Kaufman, and Shanna Pearson-Markowitz. (2007). Battleground States versus Blackout States: The Behavioral Implications of Modern Presidential Campaigns. *Journal of Politics*, 69, 786-797.

- Hetherington, Marc J. (2005). *Why Trust Matters: Declining Political Trust and the Demise of American Liberalism*. Princeton: Princeton University Press.
- Hibbing, John R. and Elizabeth Theiss-Morse. (2002). *Stealth Democracy: Americans' Beliefs About How Government Should Work*. New York: Cambridge University Press.
- Jamieson, Kathleen Hall, Paul Waldman, and Susan Sherr. (2000). Eliminate the Negative? Categories of Analysis for Political Advertisements. In James A. Thurber, Candice J. Nelson and David A. Dulio (Eds.), *Crowded Airwaves: Political Advertising in Elections* (pp. 44-64). Washington DC: Brookings Institution Press.
- Johnston, Richard, and Henry E. Brady. (2002). The Rolling Cross-Section Design. *Electoral Studies*, 21, 283-295.
- Johnston, Richard, Michael G. Hagen, and Kathleen Hall Jamieson. (2004). *The 2000 Presidential Election and the Foundations of Party Politics*. New York: Cambridge University Press.
- Kahn, Kim Fridkin, and Patrick J. Kenney. (1999). Do Negative Campaigns Mobilize or Suppress Turnout? Clarifying the Relationship between Negativity and Participation. *American Political Science Review*, 93, 877-89.
- Kinder, Donald R. (1986). Presidential Character Revisited. In Richard R. Lau and David O. Sears (Eds.), *Political Cognition* (pp. 233-55). Hillsdale, NJ: Lawrence Erlbaum.
- Lau, Richard R, Lee Sigelman, and Ivy Brown Rovner. (2007). The Effects of Negative Political Campaigns: A Meta-Analytic Reassessment. *Journal of Politics*, 69, 1176-1209.
- Lipsitz, Keena, Christine Trost, Matthew Grossman, and John Sides. (2005). What Voters Want from Political Campaign Communication. *Political Communication*, 22, 337-354.

- Marcus, George E., W. Russell Neuman, and Michael MacKuen. (2000). *Affective Intelligence and Political Judgment*. Chicago: University of Chicago Press.
- Mayer, William G. (1996). In Defense of Negative Campaigning. *Political Science Quarterly*, 111, 437-455.
- McGraw, Kathleen M., Edward Hasecke, and Kimberly Conger. (2003). Ambivalence, Uncertainty, and Processes of Candidate Evaluation. *Political Psychology*, 24, 421-448.
- Meffert, Michael, Michael Guge, and Milton Lodge. (2004). Good, Bad, and Ambivalent: The Consequences of Multidimensional Political Attitudes. In Willem E. Saris and Paul M. Sniderman (Eds.), *Studies in Public Opinion* (pp. 64-92). Princeton: Princeton University Press.
- Prior, Markus. (2001). Weighted Content Analysis of Political Advertisements. *Political Communication*, 18, 335-345.
- (2005). *The Pitfalls of Self-Reported News Exposure*. Paper presented at the annual meeting of the American Association for Public Opinion Association, Miami Beach, FL.
- Romer, Daniel, Kate Kenski, Paul Waldman, Christopher Adasiewicz, and Kathleen Hall Jamieson. (2004). *Capturing Campaign Dynamics: The National Annenberg Election Survey*. New York: Oxford University Press.
- Sigelman, Lee, and Mark Kugler. (2003). Why is Research on the Effects of Negative Campaigning So Inconclusive? Understanding Citizens' Perceptions of Negativity. *Journal of Politics*, 65, 142-160.
- Stevens, Daniel, John Sullivan, Barbara Allen, and Dean Alger. (2008). What's Good for the Goose is Bad for the Gander: Negative Political Advertising, Partisanship, and Turnout. *Journal of Politics*, 70, 527-541.

West, Darrell M. (2005). *Air Wars: Television Advertising in Election Campaigns, 1952-1992*.
Washington DC: CQ Press.