

“Targeted” Advertising and Voter Turnout: An Experimental Study of the 2000 Presidential Election

Joshua D. Clinton
Princeton University

John S. Lapinski
Yale University

Scholars disagree whether negative advertising demobilizes or stimulates the electorate. We use an experiment with over 10,200 eligible voters to evaluate the two leading hypotheses of negative political advertising. We extend the analysis to examine whether advertising differentially impacts the turnout of voter subpopulations depending on the advertisement’s message. In the short term, we find no evidence that exposure to negative advertisements decreases turnout and little that suggests it increases turnout. Any effect appears to depend upon the message of the advertisement and the characteristics of the viewer. In the long term, we find little evidence that the information contained in the treatment groups’ advertisements is sufficient to systematically alter turnout.

Political pollsters, media consultants, members of the press, and scholars hold various opinions on whether “negative” advertising is destabilizing democracy in America. The idea that negative campaigns decrease voter turnout and increase voter apathy was advanced by academics and warmly received by the news media. While perhaps intuitive and certainly eye-catching, this idea has recently come under attack.

In the late 1980s and throughout the 1990s, scholars focused their attention on showing how campaigns matter. This work represents a direct attack on the minimal effects hypothesis associated with the Columbia and Michigan schools of voting (Lazarsfeld, Berelson, and Gaudet 1948; Berelson, Lazarsfeld, and McPhee 1954; Campbell et al. 1960; Campbell 2000). Of particular interest is the question of whether political advertisements—especially negative advertisements—demobilize or stimulate voters.

Interest in the relationship between campaigns and voting has persisted for over fifty years, but only recently have researchers argued that political advertising has contributed heavily to the disappearance of voters. Although scholars have used many methods to test if negative advertisements increase or decrease turnout, much remains unsettled.

This article contributes to the conversation regarding advertising effects in several ways. First, we employ a large scale, controlled, randomized experiment to evaluate the two leading theories of advertising effects. In doing so, our work provides a valuable, but to date absent, experimentally based examination of both the demobilization and stimulation hypotheses. This is important since the only previous large-scale experimental study is the highly influential and controversial *Going Negative* study of Ansolabehere and Iyengar (1995)—a study that shows that negative advertising is causally related to lower levels of voter turnout. Second, we design our experiments to test for both the *immediate effects* of political advertising as well as for possible *lasting effects*. Ours is the first experimentally based study to test for long-term effects. This move is important since researchers are ultimately concerned with whether advertising has an effect that persists beyond the laboratory.

A third contribution is more theoretical in nature. Relying upon interviews with political consultants, we offer a third theoretical possibility: the “differential effects” hypothesis. While it is too strong to suggest that the differential effects hypothesis represents a new theory of political advertising, its articulation and examination is valuable because it has implications that differ sufficiently from existing frameworks. Fourth, our experimental design includes several features that attempt to improve external validity.

The article proceeds as follows. Section one outlines the theoretical questions and issues we address using both existing literature and interviews with the campaign consultants of Bush and Gore. In addition to sketching briefly the demobilization and stimulation hypotheses, we also explain the differential effects hypothesis. Section two describes the experimental research design we use to examine how political advertisements affect turnout. Section three investigates the immediate impact of political advertising on viewers’ self-reported likelihood of voting and evaluates the evidence in favor of each of the three hypotheses. Section four uses the panel design we employ to analyze the persistence of treatment effects on turnout, and section five concludes.

Theory and Hypotheses

Two theories dominate political advertising research—the demobilization and stimulation hypotheses. The names are appropriate as each indicates the expected impact of negative advertising on voter turnout. The primary claims of the demobilization and the stimulation hypotheses address the influence of an advertisement’s tone and not its message. To understand the focus on advertising tone, consider three post-WWII trends. First, voter turnout has continued to decline (Rosenstone and Hansen 1993; Sigelman et al. 1985). Second, political campaigns have grown in scope, breadth, and cost (Sorauf 1988). Third, campaigns have become predominantly negative (Jamieson 2001; West 1997).

The demobilization hypothesis argues that a causal relationship exists between the (negatively correlated) trending variables of voter turnout and campaign negativity (Ansolabehere and Iyengar 1995; Ansolabehere et al. 1994). Specifically:

H1A (DEMOBILIZATION): Negative advertising decreases voter turnout.

Although Ansolabehere et al. note that campaigns can be “mobilizing or demobilizing events, depending upon the nature of the messages they generate” (1994, 829), they argue that negative advertisements in presidential, congressional, and gubernatorial races increase voter apathy and thereby decrease turnout. They also report that exposure to positive advertisements has the opposite effect.

The behavior of political Independents is of particular interest to the demobilization hypothesis given their increasing size in the electorate and their status as being “up for grabs.” The demobilization hypothesis claims “to the extent that Independents are responsive to any form of advertising, it is negative advertising” (Ansolabehere and Iyengar 1995, 94). Negative advertising is believed to have a pernicious effect on Independents since attack campaigns “heighten the partisan flavor of political discourse” and drive “the Independent voter from the active electorate” by breeding “distrust of the electoral process and pessimism about the value of an individual’s own voice” (112–13). This claim implies:

H1B (DEMOBILIZATION COROLLARY): Negative advertising decreases the turnout of Independents more than those identifying as Democrats or Republicans.

Empirical support for the demobilization hypothesis is mixed (see Lau et al. 1999). Nonetheless, the demobilization hypothesis has survived despite a large amount of contradictory evidence primarily because it is supported by the experimental work of Ansolabehere, Iyengar, and their coauthors.¹ Ansolabehere et al.’s (1994) research is the only large-scale, empirical project that relies upon controlled, randomized experiments across time and different types of races. Consequently we believe the claims of the demobilization hypothesis have persisted because Ansolabehere and Iyengar’s experimental work has not been adequately replicated.

An alternative theoretical framework is the stimulation hypothesis credited to Finkel and Geer (1998), although many scholars have built upon and refined the framework (Freedman and Goldstein 1999; Lau and Pomper 2002; Wattenberg and Brians 1999). As initially posited by Finkel and Geer, a weak version of the hypothesis is essentially the null hypothesis of the demobilization hypothesis. A stronger interpretation is:

H2A (STIMULATION): Exposure to negative advertising may increase turnout.

The stimulation hypothesis draws upon insights from several literatures and suggests that negative information is more helpful than positive information in making political decisions (Hamilton and Zanna 1972; Kanouse and Hanson 1972; Lau 1982, 1985). Lau shows that, all else equal, negative information is more influential than positive information because “negative stimuli somehow

¹Note that some survey-based work also supports the demobilization hypothesis (Ansolabehere, Iyengar, and Simon 1999; Kahn and Kenney 1999).

stand out against a generally positive background” and because “people are more strongly motivated to avoid costs than to approach gains” (Lau 1985, 132). The logic of the stimulation hypothesis is: (1) negative advertising is equal or more informative than positive advertising, (2) information increases interest, (3) interested voters are more likely to vote, leading to (4) an increase in turnout (or at least not a decrease) in response to exposure to negative advertising. In other words, negative advertising increases turnout because negative advertisements result in more informed voters (or at least not less informed voters).²

Empirical work investigating the stimulation hypothesis also reveals a relationship between the amount of attention paid by voters to the campaign (media exposure) and advertising’s effect. Specifically, “the effects of advertising tone are statistically insignificant among all media exposure groups, except the ‘high media group’” (Finkel and Geer 1998, 587)—presumably because this group possesses the necessary context with which to process the information.

H2B (STIMULATION COROLLARY): Exposure to negative advertising increases turnout among individuals who are attentive to the campaign.

A considerable amount of empirical work, the bulk originating from work with National Election Studies (NES) survey data, supports the stimulation hypothesis (Finkel and Geer 1998; Freedman and Goldstein 1999; Lau and Pomper 2002; Wattenberg and Brians 1999). Most suggests that negativity slightly increases turnout in presidential, congressional, and gubernatorial races.

Interviews with seven leading political consultants of the Republican National Committee (RNC) and the Democratic National Committee (DNC) as well as with individuals who worked directly for the Bush and Gore campaigns lead us to believe that perhaps neither the demobilization nor stimulation hypotheses expresses the relationship between political advertising and voter turnout. Consultants who produce and screen test the political advertisements that researchers study firmly believe that advertising tone is not the crucial factor in influencing turnout or vote choice. Instead, they believe that what matters most to voters in determining whether or not they vote are the issues and personalities involved in the race. Alex Castellanos, a Republican media consultant who worked for the RNC in the 2000 election, reflects the thinking of most of the campaign consultants we interviewed. Claiming that knowing an advertisement’s tone is insufficient to understand an advertisement’s effect on voter turnout, Castellanos notes:

Does negative advertising depress turnout? As with most things in life, yes and no. Oakland Raiders football is hard, smash-mouth football. And sometimes more people watch the game when it is mean, smash-mouth football. But sometimes mean football means less people will watch (Castellanos 2001).

² Although the hypothesis is vague as to when turnout should increase or remain constant in response to exposure (although presumably it has to do with the information level of the voter [Zaller 1992]), in either case the prediction offers a stark contrast to that of the demobilization hypothesis.

Castellanos follows his remark with a discussion/description of the issues important to the demographic groups vital to the Republican Party’s electoral success and how advertisements are used to mobilize and persuade prospective voters. He emphasizes how issues and values are important to voters and suggests that although the tone of an advertisement may be important, the advertisement’s message is more important.

Although issues and values appear to matter most to political consultants interested in creating effective advertisements, neither the demobilization nor the stimulation hypotheses fully account for this possibility. Instead, the two frameworks focus primarily on advertising tone. Of course, it should be noted that some important work—particularly the work of Kahn and Kenney (1999)—argues that tone and issue content interact with each other. This point is considered in further depth below.

The alternative to the demobilization and stimulation hypothesis we suggest involves the possibility that advertisements attempt to prime voters by activating (potentially latent) predispositions—some of which, once activated, lead citizens to increase their likelihood of voting. Since not all issues in a campaign are of equal interest/consequence to all voters, the impact advertisements have on turnout is conditional on both the message of the advertisement and the interests of the audience. For example, the political debate about prescription drug plans for the elderly is arguably most consequential to the elderly. Therefore, advertisements about prescription drug plans are arguably likely to be highly relevant to the elderly’s decision of whether and for whom to vote. The “differential effects” hypothesis suggests:

H3 (DIFFERENTIAL EFFECTS): Advertisements affect turnout if the advertisement addresses an issue of potential concern to the voter.³

The differential effects hypothesis follows the stimulation hypothesis in assuming that the turnout decision of voters depends upon the amount of personally relevant information they possess. It distinguishes itself from the stimulation hypothesis in that it locates an advertisement’s influence on voter participation primarily in its message or issue content rather than its tone. Although this does not preclude the possibility that tone and issue content interact, issue content is the mechanism by which turnout is affected.⁴

The idea that issues matter to political campaigns is not a new perspective. However, empirical support for the claim focuses solely on the effect of issues on vote choice. In other words, the focus is on *who* prospective voters choose and

³ Finkel and Geer provide part of the motivation for our differential effects hypothesis when they write: “And future research should explore more extensively the potentially *differential effects* of trait and issue negativity on the electorate, and the potentially differential effects of the press coverage and campaign advertisements on aggregate and individual level turnout” (1998, 591–92).

⁴ Kahn and Kenney (1999) argue that tone and issue content both matter, but their emphasis is on creating a more nuanced categorization of negative campaigning. Their analysis is also at the campaign (not advertisement) level.

whether voters are “primed” so that certain issues or considerations are made more salient—not whether prospective voters actually vote or not (e.g., Carsey 2000; Iyengar and Kinder 1987; Iyengar and Simon 1993; Johnson et al. 1992; Krosnick and Kinder 1990; Mendelberg 2001; Milburn 1991; Pan and Kosicki 1997). The differential effects hypothesis draws upon the priming literature by arguing that just as not all prospective voters are equally susceptible to priming (Zaller 1992), so too is it possible that not all issues affect voters in the same way.

Since the differential effects hypothesis holds that an advertisement influences the participation of a viewer only insofar as it addresses an issue of concern to the viewer or exposes the viewer to a new issue of interest, it is necessary to identify the issues of (potential) relevance to voter subpopulations in order to evaluate the validity of the hypothesis. Consequently, we rely heavily on the ideas of political consultants to identify which issues mattered for distinct demographic groups in the 2000 presidential election.

Experimental Design and Methodology

Existing support for the demobilization and stimulation hypotheses is ambiguous because the strongest support for each depends upon different methodologies. Support for the demobilization hypothesis largely stems from the large-scale (and unreplicated) experimental work of Ansolabehere and his colleagues. By utilizing randomized, controlled experiments, they are arguably able to determine causation (Gerber and Green 2000). In contrast, support for the stimulation hypothesis is based upon the results of observational studies. Given the conflicting evidence from the different methodologies, it is not entirely clear which framework best describes how advertisements affect political participation. To resolve this controversy, we reexamine both hypotheses in an experimental setting.

Although experimental methods occupy a privileged place in social science due to their ability to determine causation, they are not without their limitations. Two criticisms are particularly applicable to experimental work on the effect of advertisements. First, because of the “one-shot” manner in which advertising experiments are conducted, they usually have the ability to measure only short-term effects (Bartels 1993; Finkel and Geer 1998). Consequently, it is unclear whether the measured effects are real lasting effects or merely short-term reactions to the provided stimuli (treatments). Second, experiments “have their own considerable limitations, primarily with respect to external validity” (Bartels 1993, 267). In particular, it is unclear the extent to which results performed on a convenience sample recruited through advertisements in local newspapers, shopping malls, and other public venues generalize to the national population (Smith 1983). It is also ambiguous how results from an experimental condition generalize to real-world conditions, as experimental conditions necessarily consider only a small subset of advertisements in an environment sterile of outside informa-

tion. We seek to alleviate some of the limitations in the experimental design we adopt.

We conduct our experiments using the panel of Knowledge Networks (KN). KN's panelists are randomly selected using list-assisted RDD (Random Digit Dialing) sampling techniques on a quarterly-updated sample frame consisting of the entire United States telephone population who fall within the Microsoft Web TV network (87% of the United States population fall within this network). Our sample respondents are therefore very close to a national RDD sample. KN Panelists are given an interactive television device (a Microsoft WebTV) and a free Internet connection in exchange for taking surveys. We use randomly selected eligible voters (i.e., panelists over the age of 18) from the KN panel for the study.

Using the KN panel enables us to improve incrementally upon the experimental design of previous research in three ways. First, although our surveys are conducted over the Internet, the selection mechanism used by KN to recruit panelists ensures that our subjects are close to a random probability sample of the United States adult population.⁵ Second, because we use a preexisting panel, recontact costs are trivial. Consequently, we contact respondents after the election to determine if they report having voted. This enables us to examine, for the first time, whether advertisements' effects last beyond the interview containing the experimental exposure and the extent to which reported effects of previous experiments may be inflated due to artificial stimulation created by experimental conditions. A third benefit is that panelists took self-administered experimental surveys in their own homes using an interactive television device that instantly displays high-quality video content on their television.

The experiment took place in the course of a KN respondent's weekly survey (KN respondents are asked to take one survey a week in exchange for free Internet service and a WebTV). Participants in our study were sent an email to their WebTV account informing them that their next survey was ready to be taken. A hyperlink (to the KN servers) in the email led to our experimental survey. Individuals could complete the survey at their own convenience.⁶ As all respondents must take the survey on the WebTV device, the survey was designed knowing that all participants have an identical survey experience (Dillman 2000). The survey asked a variety of political and nonpolitical questions (e.g., their interest in NASCAR); however, our questions were always asked first. During the survey, those participants randomly selected to be in a treatment group were shown a

⁵ All telephone numbers have an equal probability of selection, and the sampling is done without replacement. All telephone numbers not in Microsoft's WebTV network (around 13% at the time the experiments were fielded (9/2000–11/2000)) were excluded. Selected numbers are called and the head of household is informed that the household has been selected to join the Knowledge Networks panel. The household cooperation rate during this time averaged 56%.

⁶ The median response date was 2.7 days for both waves of our study. This means that half of the respondents who took the survey completed it within 2.7 days of being sent the survey invitation email. The response rate was 68%, which is the number of people taking (numerator) the survey divided by the number assigned (denominator).

full-screen advertisement and asked a few follow-up questions. The participants had no knowledge of their participation in an experiment, although they were certainly aware of the fact that they were taking a survey with some political content.⁷

The experimental design we employ has three features. First, we vary the exposure of participants to advertisements. Two of the five treatment groups are exposed to only one advertisement. In doing so, we replicate some of the strongest experimental findings of Ansolabehere and Iyengar (1995). Although the “one ad” approach isolates an advertisement’s effect, it does not match well with the realities of a presidential election campaign, since an advertisement revealed to be effective when shown in isolation may no longer be so when an opponent’s response is also shown. To measure the effect of advertisements in a more realistic environment, three of the treatment groups expose respondents to competing advertisements.

Second, we use advertisements that were produced and aired by the competing candidates (or their respective parties) in the 2000 presidential election. Using real advertisements was a conscious decision: we want to understand the effect of actual political advertisements on turnout. While the advertisements deal with several different policy issues, strictly speaking this should not be too important, given that only an advertisement’s *tone* matters for the demobilization and stimulation hypotheses. Finally, we administer a post-election survey to every respondent in our study to determine whether they report having voted in the presidential election. The experimental design is summarized in the appendix.

Treatment groups A (Gore negative and positive), B (Gore positive), and C (Gore negative) were in the field from October 10, 2000 to November 7, 2000. Due to a technical difficulty (known prior to the administration of the surveys), a (essentially) random subset of 4,614 respondents in Wave I saw no videos. This group serves as our control group for Wave I.⁸ Wave II (consisting of treatment groups D (Gore negative; Bush positive), E (Gore negative, Bush negative) and Control Wave II) was in the field from October 30, 2000 to November 6, 2000. The median respondent in Wave I completed her survey on October 12, 2000 (i.e., two days after being sent the survey), and the median respondent in Wave II completed her survey on November 1, 2000 (i.e., three days after being sent the survey). Respondents were randomly assigned to treatment groups, with 2,850 respondents assigned to groups A, B, and C, 2,500 assigned to groups D and E, and 1,500 assigned to group Control II.

By design, the goals of Wave I and Wave II differ. The experimental manipulations of Wave I attempt to quantify the effect of being shown a single advertisement (or a pair of advertisements) from Gore on the likelihood of voting. In other words, it tries to determine the treatment effect when presented with a single

⁷ Unlike the experiments of Ansolabehere and Iyengar (1995), where the respondents were unaware of the goal of the interview, our advertisements are embedded within a political survey. It is unknown both whether this difference introduces any bias or the direction of such a bias.

⁸ To control for the possibility that the group that saw no advertisements systematically differs from those who did, we employ statistical controls.

(Democratic) perspective. Wave II investigates the effect of advertisements when exposure is not limited to the views of a single candidate. Specifically, Wave II investigates the effect of a Bush (R) positive or negative advertisement *conditional on having seen a Gore (D) negative advertisement*. Thus, Wave II attempts to mimic the dynamics of a campaign and assess the impact of advertisements expressing different (partisan) viewpoints. Cost concerns resulted in our focus on selected ad pairings.

Testing the demobilization and stimulation hypotheses requires classifying the tone of an advertisement as being predominantly “negative” or “positive.” Some have argued that a simple negative versus positive categorization is insufficient and that a more nuanced scheme is needed. Since there is some controversy on what the “correct” schema is, and since the stimulation and demobilization hypotheses do not make such distinctions, we create a simple dichotomous coding scheme informed by the political consultants interviewed as well as our own content analysis. Since investigating the differential effects hypothesis requires identifying both the advertisement’s message as well as its target population, Table 1 describes the content, tone, and subpopulations targeted by each advertisement in our study according to the interviewed consultants.⁹

As is to be expected from random assignment, there is no evidence of significant sampling variability in the Wave I treatment groups. However, relative to the national population, every Wave I group overrepresents males and underrepresents blacks and those without a high-school education. This is not uncommon in RDD phone survey research and is reflective of the KN panel.¹⁰ Although the demographics of Wave II are closer to the national population than Wave I (despite the random assignment of respondents), there are some significant differences in the demographic composition of the Wave II treatment groups (e.g., the proportion over the age of 60). Consequently, Wave II analysis controls for the demographic differences. The appendix presents a breakdown of sample demographics.

Results of the Pre-Election Experiments

In this section we examine the experimental data to determine the support for the three hypotheses in terms of the immediate (short-term) effects of advertisements on voter turnout in presidential elections. The dependent variable is self-reported voter turnout. Although self-reported turnout is clearly an imperfect measure of actual turnout given the tendency of potential voters to overreport

⁹ Although researchers have begun to investigate the possibility that voters may perceive advertisements differently than experts in consequential ways (e.g., Freedman and Goldstein 1999; Lawton and Freedman 2001), our interest is only in whether there are any systematic effects of objectively classified advertisements, not the more nuanced question of whether different voters perceive advertisements differently.

¹⁰ A detailed research note on Lapinski’s web page (<http://www.yale.edu/newmedia>) describes the sample in full detail along with comparisons to the Census and the 2000 National Election Study.

TABLE 1
Description of Advertisements in Study

Name	Sponsor	Consultant	Tone	Issue	Run Dates/Context	Targeted Groups
"Phyllis Hunter"	Bush	Russ Schriefer of Stevens & Schriefer Group	+	Education Record in Texas	10/18/00 to 10/30/00 Part of the original convention film, Phyllis Hunter is a black TX educator whose best line is "we've got George leading the way" (Schriefer, 2001).	Suburban women
"Gore-gantuan"	RNC	Alex Castellanos National Media	-	Gov't Spending, Budget Surplus	10/10/00 to 11/6/01 Al Gore is a bigger spender than Bill Clinton.	Older, whiter suburban men & "good ole mom at home" (Castellanos, 2001). Seniors, Women
"Bean Counter"	Gore	Pollsters Hickman/Brown Consultant Bill Knapp of Squier, Knapp & Dunn	+	Prescription Drugs	8/30/00 to 9/24/00 More than just selling prescription drugs, it was also about "selling a person" (Brown, 2001).	Seniors
"Siding"	DNC		-	Prescription Drugs	8/29/00 to 9/3/00 Very short life and limited buy. Bush sides with the drug companies.	Seniors
"Trickle Down"	Gore	Pollsters Hickman/Brown Consultant Bill Knapp of Squier, Knapp & Dunn	-	Tax Cuts	10/6/00 to 10/20/00 Aimed at talking about the tax cut, but was "too much of a detour" (Knapp, 2001).	Lower income men

voting (see Shaw, de la Garza, and Lee 2000), this is not a concern for us. Respondents in the treatment and control groups of our randomized experiments are drawn from the same pool of respondents (i.e., KN panelists) and asked the same questions. Consequently, there is no reason to suspect that the overreport bias differentially affects the treatment or control groups—implying that the estimated treatment effects are unbiased. In other words, since respondents in the control group are as likely to overreport as those in the treatment groups, we can attribute any difference (after employing adequate statistical controls) to the effect of advertisements. Thus, although the estimation of *level effects* (i.e., actual voter turnout) is biased due to voter overreporting, estimation of the *treatment effects* associated with showing people advertisements is unaffected.¹¹ An important consequence of this limitation is that it is not possible to forecast from our results the net effect of advertisements on overall actual turnout. However, the results do allow us to determine if positive and negative advertisements increase or decrease turnout.

Wave I Results

Respondents in Wave I were asked the following question after being shown an advertisement(s):¹²

“Please rate the chances that you will vote in the presidential election in November. Are you: Certain to vote, Probably vote, Chances 50/50, Less than 50/50, Probably won’t vote”

A χ^2 test for independence of the Wave I turnout question shows that one cannot reject the (null) hypothesis that the probability of voting is unrelated to the treatment.¹³

The inability to distinguish statistically between the voting intentions of the treatment groups indicates that exposure to a single Gore negative advertisement (Group C), a single Gore positive advertisement (Group B), or exposure to both a Gore negative and positive advertisement (Group A) does not affect the self-reported probability of voting. This constitutes evidence against the demobilization hypothesis, since respondents exposed to a negative advertisement do not systematically differ in their likelihood of voting relative to those who were not exposed. Although a single negative advertisement may seem insufficient to provoke voter apathy, recall that previous support for the demobilization hypothesis was found in experiments using a single advertisement *hidden in a newscast*

¹¹ Although one might hope that it is possible to test the claim that any bias that exists affects all treatment groups equally, this is impossible because it would require the simultaneous recovery of the unknown treatment effect and the unknown bias associated with the treatment effect.

¹² Immediately prior to the administration of the survey, we knew that some respondents would not be able to see the advertisements because of a temporary technical difficulty. We were unable to identify these respondents prior to assignment. This group was intended to serve as a control group along with a smaller, fully randomized control group. Unfortunately, the smaller control group was never fielded in Wave I, making it necessary to rely upon multivariate methods to analyze Wave I.

¹³ A test of independence yields a test statistic of $\chi^2(12) = 4.26$.

(Ansolabehere and Iyengar 1995). It may also be posited that the lack of a systematic effect is due to a fault in the advertisement. However, this would imply that something other than tone influences an advertisement's effectiveness.

Interpreting the results in terms of the stimulation hypothesis is more difficult due to the inherent ambiguity in the predictions made by the hypothesis. In other words, because the stimulation hypothesis can predict a null finding (if the advertisement is sufficiently uninformative), it is unclear whether the lack of treatment effects is evidence against the stimulation hypothesis or evidence that the advertisements tested were sufficiently uninformative so as to fail to stimulate voters.

To investigate fully the amount of support for the stimulation and differential effects hypotheses and to test the demobilization hypothesis' corollary (H2B) that party affiliation matters, we examine the effect of the advertisements' conditioning on both respondent characteristics and treatment-group assignment. This examines the possibility that although there are no unconditional and systematic effects of exposure, advertisements may affect the voting behavior of specific population subgroups. To test for this possibility, we estimate a probit model for the dichotomous dependent variable of whether the respondent reports that they are "Certain to vote." The covariates used, described in the appendix, include measures of respondents' personal and political characteristics, as well as their exposure to the campaign. We attempt to control for exposure to the campaign by including (the log-transformation of) the amount of money spent by the Democratic Party campaigning for the presidential election in the respondent's state, which was provided by Kirk Brown, pollster for Vice President Al Gore. We also control for the respondent's interest in the campaign and the presence of a Senate or gubernatorial campaign. As others have argued that turnout decreases as the "cost" of voting rises (Wolfinger and Rosenstone 1980), we also control for registration ease (the number of days before an election the state requires to register).

Table 2 reports the estimation results of three model specifications. Every model specification improves upon the 73% fit of the naive model (which simply predicts that every respondent is "Certain to vote"). Specifications Model 1 and Model 2 examine the possibility that there is a systematic effect on the likelihood of voting once respondents' demographic and political characteristics are controlled.¹⁴ The specifications differ in that Model 1 list-wise deletes the 14% missing data in the household income variable and Model 2 uses the multiple imputation procedure outlined in King et al. (2001).¹⁵ Table 2 reports the model specifications and results, with coefficients attaining standard levels of significance (i.e., 5% probability of making a two-tailed Type I error listed as (**) and 10% as (*)).

¹⁴In other words, these two specifications examine the possibility that the null result of the χ^2 test reported above is the result of offsetting compositional differences in the treatment groups.

¹⁵Technically, the multivariate normal imputation algorithm of King et al. (2001) cannot sensibly be used for categorical data. This criticism is muted for our case given that household income is measured on a 17-point scale and therefore is "roughly continuous."

TABLE 2
Wave I Probit Turnout Model

Variable	Model 1	Model 2 (Income Imputed)	Model 3: Main (Income Imputed)	Model 3: Group A Interactions (Gore negative, Gore positive)	Model 3: Group B Interactions (Gore positive)	Model 3: Group C Interactions (Gore negative)
Constant	-2.39** (.234)	-2.29** (.217)	-2.43** (.256)			
Group A [Gore negative, Gore positive]	-0.02 (.083)	.037 (.076)	-1.31* (.751)			
Group B [Gore positive]	.131* (.074)	.141** (.068)	1.68** (.630)			
Group C [Gore negative]	.147* (.078)	.148** (.073)	.738 (.708)			
Log of Democratic Spending	-0.03 (.003)	-0.04 (.003)	-0.005* (.0029)			
Registration Rules	-0.03 (.003)	-0.005 (.003)	-0.005* (.0028)			
Senate Race	.055 (.058)	.054 (.054)	.054 (.054)			
Gubernatorial Race	.151** (.072)	.151** (.067)	.156** (.068)			
Registered to Vote	1.52** (.067)	1.49** (.061)	1.51** (.062)			

TABLE 2 continued

Variable	Model 1	Model 2 (Income Imputed)	Model 3: Main (Income Imputed)	Model 3: Group A Interactions (Gore negative, Gore positive)	Model 3: Group B Interactions (Gore positive)	Model 3: Group C Interactions (Gore negative)
Black	.283** (.092)	.313** (.086)	.315** (.087)			
Education	.094** (.024)	.086** (.022)	.069** (.027)	.148* (.086)	.019 (.070)	.011 (.074)
Attention to Campaign [low]	-.887** (.064)	-.899** (.059)	-.905** (.072)	.344 (.220)	-.005 (.176)	-.270 (.198)
Attention to Campaign [high]	1.04** (.053)	1.03** (.049)	1.12** (.060)	.008 (.186)	-.259* (.159)	-.445** (.163)
Age	.008 (.008)	.010 (.008)	.018* (.010)	.010 (.031)	-.078** (.028)	-.008 (.029)
Age ²	.00003 (.00008)	-.00002 (.00008)	-.00007 (.00009)	.0001 (.0003)	.0008** (.0003)	.00003 (.0003)
Female	.159** (.047)	.143** (.043)	.064 (.052)	.122 (.161)	.325** (.140)	.237 (.148)
Democrat	.235** (.061)	.229** (.056)	.248** (.068)	.191 (.208)	-.015 (.173)	-.276 (.184)
Republican	.374** (.064)	.366** (.059)	.334** (.071)	.197 (.220)	.156 (.185)	-.012 (.199)
HH Income	.040** (.007)	.039** (.006)	.042** (.008)	.001 (.024)	-.007 (.023)	-.011 (.021)
N	5640	6499	6499			
% Corr. Pred.	85%	85%	85%			
Pseudo R ²	.43	.43	.44			

The similarity of the results for Models 1 and 2 reveals that multiple imputation for the missing values makes little substantive difference. Model 2’s results reveal that there is some systematic effect on respondents’ self-reported certainty of voting when treatment effects are assumed to be constant across individuals conditional on the included covariates. The coefficients of interest in Model 2 are those of *Group A*, *Group B*, and *Group C*, as they represent the difference in propensity of voting relative to the control group (which is the omitted category). Under the demobilization (stimulation) hypothesis, *Group B* should be positive (positive) and *Group C* should be negative (positive and statistically greater than that of *Group B*). The results indicate that exposure to either has a positive effect on turnout.¹⁶ This result is inconsistent with the demobilization hypothesis. Although model 2 does reveal that exposure to the advertisements increases the certainty of voting in *Groups B* and *C*, the negative advertisement is not significantly more influential than the positive advertisement. That exposure to both advertisements (*Group A*) has no effect is surprising given the individual exposure effects (i.e., *Group B* and *Group C*).

Model 3 permits treatment effects to vary by respondent characteristics to explore the corollaries to the demobilization and stimulation hypotheses and to test the differential effects hypothesis. The column labeled “Model 3: Main” provides the main effect coefficients and standard errors while the columns labeled “Model 3: Gore negative and positive,” “Model 3: Gore positive,” and “Model 3: Gore negative” provide the results for the interactions of the treatment effects with the reported variables. In other words, the coefficients in the “Model 3: Gore negative and positive” column represent the interaction terms of *Group A* with the main effects listed in the first column. A likelihood ratio test supports the (joint) inclusion of the reported interactions.¹⁷

The predictions of the demobilization and stimulation hypotheses are identical to the predictions for Model 2 in terms of the main effects. The fact that only the coefficient on *Group B* (Gore positive) is positive at traditional significance levels (although *Group A* (Gore positive and Gore negative) is positive at a weaker level) is damaging to both the demobilization and stimulation hypothesis—as the only advertisement that has a nonzero effect on all viewers is a positive advertisement once we permit the effect to covary with respondent characteristics.

¹⁶This finding contradicts the χ^2 analysis above, indicating that the distribution of respondents with the measured characteristics across the treatment groups is somewhat unequal. Consequently, once the respondent characteristics are accounted for, a significant difference emerges.

¹⁷Additional interactions were not included in the specification of Model 3 because they have no testable implications in terms of the three hypotheses. Note that models including additional interactions do not affect these substantive results. A likelihood ratio test (LRT) between the specifications of Model 2 and Model 3 reveals that one can reject the null hypothesis that the included interaction terms are jointly zero with less than a 1% chance of making a Type I error. The test yields a test statistic of $\chi^2(27) = 53.40$.

The demobilization corollary H1B predicts that Independents are most susceptible to negative advertising. If true, the interaction of *Group B* with *Republican* and *Democrat* should be positive—accounting for the fact that the impact of the advertisement is estimated relative to the omitted category of political Independents. Inspection reveals that party identifiers are not more or less likely to report being “certain to vote” relative to political Independents for any of the treatment groups. Consequently, there is no support for either the primary or corollary predictions of the demobilization hypothesis in Wave I.

The stimulation hypothesis’ corollary H2B claims that individuals who are highly interested in the campaign and primed with negative information (advertisements) are more likely to vote. Examining the interaction of *Attention to Campaign (high)* and *Group C* in the “Model 3: Gore negative” column shows that the interaction is significant and negative—implying that respondents who pay a great deal of attention to the campaign and who are exposed to a single negative advertisement are *less* likely to vote than those high attention respondents in the control group.

Having found faint support for the traditional demobilization and stimulation hypotheses, we now move to test the predictions of the differential hypothesis. Testing the hypothesis is inherently more ambiguous due to its conditional nature. However, since both advertisements deal with the issues of prescription drugs and HMO reform, we expect the largest impact on older and poorer individuals. Consequently, we would expect to observe positive interactions between the net effect of *Age* and *Age*² and *Household Income* and the treatment groups (*Group A*, *Group B*, and *Group C*). Our interviews with the consultants reveal a few additional testable implications that we empirically evaluate.

In discussing the Gore positive “Bean Counter” advertisement (shown to treatment groups A and B), Kirk Brown, the pollster responsible for Al Gore’s advertising testing in the 2000 campaign, noted that when the advertisement was shown (8/30/00 to 9/24/00) Gore was in need of a “bump,” and “when we ran ‘Bean Counter,’ which was after the convention, we had generally gotten a bounce out of the convention and had improved with women, so we tried to solidify those gains” (Brown, 2001). Referencing the testing and strategy behind the “Bean Counter” advertisement, Brown remarked that “‘Bean Counter’ worked better with women than men, but it still was pretty good with men—especially those against the interests of big companies.” Assuming the advertisement was effective, we should expect a positive interaction effect between *Female* and *Group A* and *Group B*.


In contrast, the consultants noted that the “Siding” advertisement was not very effective. As a result, it received only limited exposure—being shown for only five days (as opposed to the 26 days for the “Bean Counter” advertisement). Consequently, this advertisement may be ineffective in targeting the elderly with its HMO reform message.

Model 3 provides limited, but not unqualified, support for the differential effects hypothesis. In terms of the predictions for the positive “Bean Counter”

FIGURE 1

Group B (Gore Positive) Treatment Effect By Respondent Characteristics

The line denotes the predicted probability of a respondent of average characteristics being “Certain to Vote” using the specification of Model 3. The shaded area indicates the 95% confidence interval.



advertisement in Group B, the results are precisely as expected among women and the elderly (i.e., the interactions of *Group B* with the net effect of *Age* and *Age*² and *Female* reported in the “Model 3: Gore positive” column are positive). Given the difficulty in interpreting probit coefficients due to their conditional nature, Figure 1 depicts the treatment effect of Group B (Gore positive) relative to the control group for a respondent with average characteristics.

The upper left and right figures show that women and men in the control group are equally likely to vote and that the probability of voting increases with age. The lower left and right figures report the probability of voting for male and female respondents in Group B (Gore positive). Two items are of interest. Comparing the lower-left and lower-right figures shows that females exposed to Group B are more likely to vote than males—precisely as is expected. Comparing the upper and lower figures reveals that the elderly (and young) are more likely to vote after being exposed to the “Bean Counter” advertisement than those in the control group. This is expected given the importance of the issue to the elderly (who would benefit) and the young (who would pay).

Support for the differential hypothesis is less promising in treatment groups A (Gore negative and Gore positive) and C (Gore negative). There is no evidence that respondents of any description in Group C are any more or less likely to report being “certain to vote” relative to the control group despite being exposed to a negative advertisement dealing with HMO reform. Although it may be that the “no effect” is because the “Siding” advertisement is as ineffective (as the consultants suggest), such an interpretation cannot explain the results of Group A. Assuming that the “Siding” advertisement is ineffective, we would expect the treatment effects for Group A to be identical to those for Group B. However, the interaction effects for treatment *Group A* and *Female*, *Age* and *Age*² are not different from zero.

The results of Wave I are clearly not favorable to the predictions of either the demobilization or stimulation hypotheses, as very little support is found for either the primary or corollary predictions. The predictions of the differential hypothesis fare somewhat better, although the support is not unambiguous.

Wave II Results

The second set of experiments differs from those of Wave I in three respects. First, Wave II aims to determine the effect of an advertisement conditional on previous exposure to a competitor’s advertisement. Second, Wave II exposes respondents to advertisements only a week before Election Day. Third, the question used to assess the probability of voting differs from Wave I. Wave II respondents were asked:

“Using a ten-point scale, please indicate how likely you are to vote in this year’s elections for president and Congress. If you are certain that you will vote, pick a number closer to “8,” “9,” or “10.” If it is less likely that you will vote, use a number closer to “1,” “2,” or “3.” You may choose any number between one and ten.”

Although subjects were randomly assigned to treatment and control groups, ex post differences in the treatment groups’ demographic composition indicate the need to control for demographic characteristics. Consequently, we proceed directly to the ordinary least-squares regression of the three specifications given in Table 3. Recall that Group D consists of a negative Gore and a positive Bush advertisement, and Group E involves exposure to both a negative Gore and a negative Bush advertisement. Since Group E contains two negative advertisements, it represents our strongest test of the demobilization and stimulation hypotheses and their corollaries. The demobilization hypothesis predicts a negative coefficient for *Group E*, and the stimulation hypothesis predicts that the coefficient for *Group E* is positive and greater than *Group D* (which may also be positive).

As in the Wave I experiments, the results for Wave II reveals some evidence of systematic and positive treatment effects. Model 4—which handles the 14% nonresponse of household income using list-wise deletion—does not find either the treatment effects *Group D* or *Group E* to be statistically significant. Although Model 5, containing the multiple imputations for household income, indicates positive treatment effects for *Group E*, the magnitude of the effect is unimpressive—amounting to an increase of .2 in the likelihood of voting on a 10-point scale (and only 5% of the effect that being registered to vote has). Consequently, the findings of Model 4 and 5 are contrary to the predictions of the demobilization hypothesis and not particularly encouraging for the stimulation hypothesis.¹⁸

Model 6 permits the effect of advertisements to vary with respondent characteristics.¹⁹ As in Wave I, there is no evidence that the interactions of *Group D* and *Group E* with *Democrat* and *Republican* are positive. Contrary to the corollary prediction of the demobilization hypothesis (H1B), there is no support that political independents are less likely to vote relative to party identifiers after being exposed to two negative advertisements. Support for the stimulation hypothesis’ corollary is also elusive, as respondents who pay a great deal of attention to the campaign are actually less likely to vote after being exposed to the two negative advertisements of Group E.

Support for the differential effects hypothesis can be examined by determining whether the advertisements differentially affect the probability of voting among the targeted groups. As both advertisements in Group E (Gore’s negative “Trickle Down” and Bush’s negative “Gore-gantuan”) deal with tax cuts and government spending, we expect a positive interaction between *Group E* and *Household Income*. The interaction of *Group E* and the net effect of *Age* and *Age*² should also be positive—reflecting the concern of the elderly with government spending (on the elderly).

¹⁸Strictly speaking, a null finding cannot reject the stimulation hypothesis, as it might be the case that the advertisements are uninformative. However, so long as the theory is falsifiable, exposure to two negative advertisements represents our strongest case under which we might expect to observe an increase in turnout.

¹⁹A LRT for the null hypothesis that the interactions are jointly zero is rejected with a test statistic of $\chi^2(18) = 31.44$.

TABLE 3
Wave II OLS Turnout Model

Variable	Model 4	Model 5 (Income Imputed)	Model 6: Main (Income Imputed)	Model 6: Group D Interactions (Gore negative, Bush positive)	Model 6: Group E Interactions (Gore negative, Bush negative)
Constant	4.17** (.427)	4.12** (.393)	3.70** (.496)		
Group D [Gore negative, Bush positive]	.130 (.106)	s.145 (.099)	1.37 (.890)		
Group E [Gore negative, Bush negative]	.132 (.109)	.203** (.102)	.959 (.893)		
Log of Democratic Spending	.005 (.006)	.002 (.005)	.003 (.005)		
Registration Rules	-.006 (.006)	-.012** (.005)	-.012** (.005)		
Senate Race	-.040 (.108)	-.050 (.101)	-.038 (.101)		
Gubernatorial Race	.179 (.130)	.159 (.121)	.145 (.122)		
Registered to Vote	3.82** (.12)	3.75** (.107)	3.78** (.107)		
Black	-.066 (.156)	-.002 (.146)	.001 (.146)		
Education	.073* (.044)	.071* (.041)	.051 (.055)	-.131 (.100)	.243** (.102)
Attention to Campaign (low)	-2.66** (.136)	-2.71** (.125)	-2.55** (.163)	-.567* (.307)	-.117 (.309)
Attention to Campaign (high)	1.04** (.103)	1.05** (.097)	1.21** (.129)	-.254 (.239)	-.494** (.243)
Age	-.023 (.015)	-.017 (.013)	.005 (.018)	-.044 (.035)	-.072** (.035)
Age ²	.0003* (.0001)	.0002* (.0001)	.000006 (.0017)	.0004 (.0003)	.0007** (.0003)
Female	.017 (.085)	.054 (.080)	.055 (.107)	.080 (.199)	-.117 (.204)
Democrat	.332** (.118)	.351** (.108)	.476** (.145)	-.430** (.263)	-.102 (.274)
Republican	.457** (.125)	.522** (.115)	.655** (.152)	-.602** (.284)	.017 (.292)
HH Income	.053** (.013)	.055** (.012)	.030* (.017)	.062** (.032)	.055* (.032)
N	2711	3137	3137		
Adj. R ²	.60	.60	.60		

Deriving the expectations with respect to Group D (Bush positive “Phyllis Hunter” and Gore negative “Trickle-Down”) is less straightforward given the different issues the advertisements target. Table 1 reveals that consultants designed the advertisements in Group D to target women (with education in the “Phyllis Hunter” advertisement) and low-income men (with economic issues in “Trickle Down”). Consequently, it is not clear what the marginal effect of being female should be relative to the control group given that one advertisement targets each group. However, the focus on tax cuts can be predicted to increase the probability of voting conditional on income (i.e., the interaction of *Group D* and *Household Income* should be positive).

Model 6 reveals some support for the predictions of the differential hypothesis. The interaction terms of *Group D* and *Group E* with *Household Income* are positive (although only at 10% probability of making a Type I error for Group E). The net effect of income on the likelihood of voting for those in either treatment groups is roughly three times the effect relative to those in the control group. Contrary to predictions the net effect of the interaction of *Group E* on *Age* and *Age*² is negative for all respondents—indicating that the elderly reported a lower probability of voting after being exposed to negative advertisements against tax cuts and government (over)spending. Although the treatment effect also varies with respect to education in Group E, this finding is unanticipated.

In summary, the findings from the Wave I and Wave II experiments provide no evidence for either the primary or the corollary predictions of the demobilization hypothesis. There is no evidence that negative advertisements systematically decrease the reported likelihood of voting. Although support for the stimulation hypothesis is evident in the specifications of Model 2 and Model 5, permitting the effect of advertisements to depend on respondent characteristics reveals little evidence that negative advertising increases the turnout of viewers. Only the predictions of the differential hypothesis remain unrejected, although the support is ambiguous given that only some of the anticipated effects are realized.

The Lasting Effects of Advertising: Results from the Post-Election Survey

A unique component of our research design is that we reinterview respondents after the election to determine whether or not they report having voted. This departure from a traditional single-interaction design allows us to examine whether advertisements have a lasting effect. While measuring for persistence is significant, it is important to note a priori that it is unlikely that an effect will be found. Relative to the amount of information respondents are exposed to after taking the survey and prior to the election, the amount of information contained in the treatment advertisements is presumably relatively small (especially for Wave I respondents). It would therefore be surprising to find that exposure to the information contained in the treatment groups is decisive in the respondent’s deci-

sion of whether or not to vote. However, this is an empirical question we are in a position to answer.

Although 29% of the respondents who participated in Wave I and II failed to complete our post-election survey, conducting a test of selection effects for a binary second-stage equation reveals that there is no evidence that this nonresponse is correlated with a respondent's decision to vote (Dubin and Rivers 1989/1990).²⁰ Consequently, we condition only on those who complete the post-election survey.

Overall, 82% of the respondents in the post-election study reported voting. Although this is certainly a biased measure of actual voting behavior, insofar as the bias affects the treatment groups equally our treatment effects remain unbiased estimates of the advertisement(s) effect. As respondents in Wave I and II were administered identical post-election surveys, we pool their responses for efficiency reasons. Table 4 reports the estimates of analysis of the post-election survey using the dichotomous dependent variable of whether or not the respondents reported voting in the 2000 presidential election using list-wise deletion (Model 7) and multiple imputation (Model 8).

The results of Table 4 are not encouraging for any of the three hypotheses. The cleanest predictions from the demobilization (stimulation) hypothesis are that the treatment effects associated with negative advertisements (i.e., **Group C** and **Group E**) should be negative (positive). Model 8 reveals that only the treatment effect associated with **Group D** is nonzero. Although the lack of main effects is entirely consistent with the differential effects hypothesis, there is no evidence that the effect of exposure depends upon viewer characteristics. Contrary to the predictions of the differential effects hypothesis, we cannot reject the hypothesis that the demographic and campaign interest interactions in previous specifications are jointly zero ($\chi^2(51) = 42.99$). Necessarily, this also implies that there is also no evidence for the corollary predictions of the demobilization (H1B) and stimulation (H2B) hypotheses.

While we believe that our panel experimental design is very important, it is important to place the null results in the proper context. The fact that exposure to the treatment groups does not influence the likelihood of respondents voting in systematic (and predictable) ways does not imply that advertisements do not influence whether or not respondents vote. Instead, it demonstrates that exposure to the one (or two) advertisements used in the treatment groups does not systematically affect whether respondents decide to vote conditional on all other information about the race (or process) they possess. Although we attempt to control for exposure to campaign information, it is certainly possible—if not

²⁰The selection equation consists of the full set of interactions and the following instruments: an indicator of whether the respondent is the household member who initially agreed to participate in the KN panel, the number of days elapsed between the assignment of the respondent's first survey (KN profile survey), and the entry into our pre-election survey and whether or not the respondent is married. Despite the presence of 73 covariates, the model fit predicts nonresponse on the post-election survey little better than the naive model and achieves a pseudo R^2 of .03.

TABLE 4
Post-Election Probit Voting Model

Variable	Model 7	Model 8 (Imputed Income)
Constant (Std. Err.)	17.19* (9.13)	14.17* (8.52)
Group A [Gore negative, Gore positive]	.094 (.069)	.040 (.063)
Group B [Gore positive]	.033 (.063)	.034 (.059)
Group C [Gore negative]	.056 (.067)	.079 (.063)
Group D [Gore negative, Bush positive]	.138** (.064)	.122** (.060)
Group E [Gore negative, Bush negative]	.046 (.064)	.032 (.060)
Log of Democratic Spending	.0002 (.002)	-.0007 (.002)
Registration Rules	-.001 (.002)	-.002 (.002)
Senate Race	.003 (.044)	.003 (.04)
Gubernatorial Race	.022 (.053)	.054 (.051)
Registered Vote	.919** (.049)	.951** (.046)
Black	.004 (.069)	.030 (.065)
Education	.083** (.018)	.086** (.017)
Attention to Campaign (low)	-.413** (.054)	-.407** (.050)
Attention to Campaign (high)	.119** (.041)	.114** (.038)
Age	.010 (.006)	.011* (.006)
Age ²	-.00008 (.00006)	-.00009 (.00006)
Female	.020 (.035)	.023 (.032)
Democrat	.129** (.048)	.140** (.045)
Republican	.155** (.051)	.158** (.047)
HH Income	.007 (.005)	.007 (.005)
N	6170	7072
% Corr Pred.	74%	74%
Pseudo R ²	.12	.13

plausible—that exposure to information exterior to the treatment groups is more influential than that of the few advertisements in our treatment groups. The positive finding of our post-election analysis is that it appears necessary to expose respondents to more than two advertisements in order to create any lasting effect if we are to determine how respondents weigh various sources of information when deciding whether or not to vote. The lack of persistence we document might provide a reason as to the increasing frequency of political advertisements as the election approaches—candidates must continually advertise to remain relevant.

Conclusion

Our study attempts to settle a long-standing empirical controversy and provide some additional theoretical direction for the future study of advertising effects. The longitudinal tests we conduct provide the first large-scale experimental follow-up study to Ansolabehere and Iyengar's seminal work on negative advertising. In so doing, we subject the predictions of both the demobilization and stimulation hypotheses to tests using a common methodology in the context of the 2000 presidential election.

We find no evidence to support the predictions of the demobilization hypothesis. It is never the case that exposure to negative advertising decreases either the reported probability of voting or the actual voting. Furthermore, no evidence is found to support the claim that political Independents are most susceptible to negative advertisements.

The stimulation hypothesis fares slightly better. Our results indicate that advertisements increase (or at least do not decrease) the propensity of voting in the short term regardless of tone. However, once treatment effects are permitted to vary according to respondent characteristics, this finding vanishes—leaving no support for the claim that respondents exposed to either a single negative advertisement (Group C) or two negative advertisements (Group E) are more likely to intend to vote, even for the highly interested.

Relying on interviews with political pollsters and media consultants involved with the Bush and Gore campaigns of 2000, we also consider a third possibility: that perhaps tone is less important than the intersection of issues viewers care about and the message of the advertisement. The differential effects hypothesis we outline attempts to account for the possibility that advertisements “target” (and are most effective among) specific voter subpopulations. Although unsurprising to the political consultants, this possibility has not been widely examined by academics.

Evidence for the differential effects hypothesis is stronger than that for either the demobilization or stimulation hypotheses. Several treatment groups reveal effects that depend on voter characteristics precisely as predicted. However, support is not universal, as some predictions are not borne out.

Our results suggest several avenues for future research. First, there appears to be very little reason to continue to investigate the demobilization hypothesis in the context of presidential elections. Second, additional research into the role that information plays in determining whether a potential voter participates (and the impact on vote choice) is needed. In particular, how do voters weigh information contained in advertisements relative to other sources of information about the candidates? Our post-election survey reveals that other information is clearly relevant to whether or not voters vote and that the impact of exposure does not persist for very long outside the experimental setting. Only by understanding the context within which voters receive and analyze the information contained in political advertisements will it be possible to determine the effectiveness of advertisements. We hope our work contributes to the pursuit of a better understanding of the causal mechanisms underlying how advertising affects the electorate.

Appendix

Experimental Design

Group	Sample Size	Content	Advertisement(s)
A	2223	Gore negative Gore positive	“Siding” “Bean Counter”
B	2236	Gore positive	“Bean Counter”
C	2204	Gore negative	“Siding”
D	1381	Gore negative Bush positive	“Trickle Down” “Phyllis Hunter”
E	1356	Gore negative Bush negative	“Trickle Down” “Gore-gantuan”
Control II	832	None	None

Demographic Characteristics of Sample

Group	No HS Educ.	Black	Female	Inc. > 75K	Democrat	Age > 60
A	6%	7%	45%	23%	41%	17%
B	6%	8%	49%	24%	42%	14%
C	6%	8%	47%	27%	42%	14%
Ad not shown	5%	7%	46%	25%	41%	19%
D	10%	10%	52%	24%	43%	26%
E	8%	9%	55%	23%	42%	16%
Control Wave II	11%	9%	51%	19%	40%	23%

Appendix *continued*

Measures Used in the Analysis

Measure	Description
log(Democratic Spending)	Advertising spending by Gore and DNC in state after Democratic convention in dollars (from Kirk Brown of Hickman and Brown)
Registration rules	Number of days before the election a state requires voters to register in order to vote. Ranges from 0 (same day registration) to 31 days.
Registered to Vote	1 ("At current address" or "At a different address"); 0 (Else)
Gubernatorial race	1 (Gubernatorial race in state), 0 (no Gubernatorial race)
Senate race	1 (Senate race in state), 0 (no Senate race)
Black	1 (black respondent); 0 (otherwise)
Education	1 ("Less than High School") to 5 ("Post graduate degree [Masters, professional, doctorate]")
Attention to Campaign (orig.)	1 ("None") to 5 ("A great deal")
Recoded into:	Low (1 or 2) otherwise 0; High (4 or 5) otherwise 0
Attention to Campaign (low)	Omitted category = 3
Attention to Campaign (high)	
Age	Age of respondent (must be over 18 and eligible to vote to be a participant)
Age ²	Square of Age
Democrat	Indicator variable for identifying as Democrat (1); 0 otherwise
Republican	Indicator variable for identifying as Republican (1); 0 otherwise
HH Income	HH income in 1999: 1 ("Less than \$5K") to 17 ("More than \$125K")

Acknowledgements

The authors would like to thank several people who provided invaluable advice throughout the course of the project. These people include Charles Riemann, Daniel Slotwiner, Alan Gerber, Don Green, Simon Jackman, Doug Rivers, Robert Shapiro, Paul Sniderman, John Geer, Richard Lau, Jonathan Krasno, Daniel Galvin, and Andrew Roach. We would also like to thank the three anonymous reviewers along with the guidance of Professor William G. Jacoby. Joshua D. Clinton thanks the financial support of the National Science Foundation through its graduate fellowship program. John S. Lapinski thanks the Institution for Social and Policy Studies at Yale University, and specifically, its experimental research grant program for the generous financial support that made this project possible.

Manuscript submitted May 23, 2002

Final manuscript received October 4, 2002

References

- Ansolabehere, Stephen, and Shanto Iyengar. 1995. *Going Negative: How Political Advertising Shrinks and Polarizes the Electorate*. New York: Free Press.
- Ansolabehere, Stephen D., Shanto Iyengar, and Adam Simon. 1999. “Replicating Experiments Using Aggregate and Survey Data: The Case of Negative Advertising and Turnout.” *American Political Science Review* 93(4): 901–10.
- Ansolabehere, Stephen, Shanto Iyengar, Adam Simon, and Nicholas Valentino. 1994. “Does Attack Advertising Demobilize the Electorate?” *American Political Science Review* 88(4): 829–38.
- Bartels, Larry M. 1993. “Messages Received: The Political Impact of Media Exposure.” *American Political Science Review* 87(2): 267–85.
- Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Brown, Kirk. 2001. Personal interview.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald Stokes. 1960. *The American Voter*. New York: Wiley.
- Campbell, James E. 2000. *The American Campaign: U.S. Presidential Campaigns and the National Vote*. College Station: Texas A&M Press.
- Carsey, Thomas M. 2000. *Campaign Dynamics: The Race for Governor*. Ann Arbor: University of Michigan Press.
- Castellanos, Alex. 2001. Personal interview.
- Dillman, Don A. 2000. *Mail and Internet Surveys: The Tailored Design Method*. New York: John Wiley and Sons.
- Dubin, Jeffrey A., and Douglas Rivers. 1989/1990. “Selection Bias in Linear Regression, Logit and Probit Models.” *Sociological Methods and Research* 18(2/3): 360–90.
- Finkel, Steven, and John Geer. 1998. “A Spot Check: Casting Doubt on the Demobilizing Effect of Attack Advertising.” *American Journal of Political Science* 42(2): 573–95.
- Freedman, Paul, and Ken Goldstein. 1999. “Measuring Media Exposure and the Effects of Negative Campaign Ads.” *American Journal of Political Science* 43(4): 1189–1208.
- Gerber, Alan S., and Donald P. Green. 2000. “The Effects of Personal Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment.” *American Political Science Review* 94(3): 653–64.
- Hamilton, David L., and Mark P. Zanna. 1972. “Differential Weighting of Favorable and Unfavorable Attributes in Impressions of Personality.” *Journal of Experimental Research in Personality* 6(1): 204–12.
- Iyengar, Shanto, and Donald R. Kinder. 1987. *News That Matters*. Chicago: University of Chicago Press.
- Iyengar, Shanto, and Adam Simon. 1993. “News Coverage of the Gulf Crisis and Public Opinion: A Study of Agenda-Setting, Framing and Priming.” *Communication Research* 20: 365–38.
- Jamieson, Kathleen Hall. 2001. *Everything You Think You Know About Politics . . . And Why You’re Wrong*. New York: Basic Books.
- Johnson, Richard, Andre Blais, Henry E. Brady, and Jean Crete. 1992. *Letting the People Decide: Dynamics of a Canadian Election*. Stanford: Stanford 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(4): 877–90.
- Kanouse, David E., and L. Reid Hanson. 1972. “Negativity in Evaluation.” *Attribution: Perceiving the Causes of Behavior*. Edward E. Jones et al., eds. Morristown, NJ: General Learning Press.
- King, Gary, James Honaker, Anne Joseph, and Kenneth Scheve. 2001. “Analyzing Incomplete Political Science Data: An Alternative Algorithm for Multiple Imputation.” *American Political Science Review* 95(1): 49–70.
- Knapp, Bill. 2001. Personal interview.

- Krosnick, Jon A., and Donald R. Kinder. 1990. "Altering the Foundations of Support for the President through Priming." *American Political Science Review* 84(2): 497–512.
- Lau, Richard R. 1982. "Negativity in Political Perception." *Political Behavior* 4(4): 353–78.
- Lau, Richard R. 1985. "Two Explanations for Negativity Effects in Political Behavior." *American Journal of Political Science* 29(1): 119–38.
- Lau, Richard R., and Gerald Pomper. 2002. "Effectiveness of Negative Campaigning in U.S. Senate Elections." *American Journal of Political Science* 46(1): 47–66.
- 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): 851–76.
- Lawton, Dale L., and Paul Freedman. 2001. "Beyond Negativity: Advertising Effects in the 2000 Virginia Senate Race." Delivered at the Annual Meeting of the Midwest Political Science Association.
- Lazarsfeld, Paul F., Bernard Berelson, and Hazel Gaudet. 1948. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign*. New York: Columbia University Press.
- Mendelberg, Tali. 2001. *The Race Card: Campaign Strategy, Implicit Messages, and the Norm of Equality*. Princeton: Princeton University Press.
- Milburn, Michael A. 1991. *Persuasion and Politics: The Social Psychology of Public Opinion*. Pacific Grove, CA: Brooks/Cole Publishing Company.
- Pan, Zhongdang, and Gerald M. Kosicki. 1997. "Priming and Media Impact on the Evaluations of the President's Performance." *Communication Research* 24(1): 3–30.
- Rosenstone, Steven, and John Mark Hansen. 1993. *Mobilization, Participation, and Democracy*. New York: Macmillan.
- Schriefer, Russ. 2001. Personal interview.
- Shaw, Daron R., Rodolfo O. de la Garza, and Jongho Lee. 2000. "Examining Latino Turnout in 1996: A Three-State, Validated Survey Approach." *American Journal of Political Science* 44(2): 338–46.
- Smith, T.M.F. 1983. "On the Validity of Inferences from Non-random Sample," *Journal of the Royal Statistical Society. Series A (General)* 146(4): 394–403.
- Sigelman, Lee, Philip W. Roeder, Malcolm E. Jewell, and Micheal A. Baer. 1985. "Voting and Non-voting: A Multi-Election Perspective." *American Journal of Political Science* 29(4): 749–65.
- Sorauf, Frank J. 1988. *Money in American Elections*. Glenview, IL: Scott, Foresman/Little.
- Wattenberg, Martin P., and Craig Leonard Brians. 1999. "Negative Campaign Advertising: Demobilizer or Mobilizer." *American Political Science Review* 93 (December): 877–91.
- West, Darrell M. 1997. *Air Wars: Television Advertising in Election Campaigns, 1952–1996*. Washington: Congressional Quarterly Inc.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven: Yale University Press.
- Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Joshua Clinton is assistant professor of political science, Princeton University, Princeton, NJ (clinton@princeton.edu). John S. Lapinski is assistant professor of political science and Resident Fellow at the Institution for Social and Policy Studies, Yale University, New Haven CT (john.lapinski@yale.edu).