

I'm Asking for Your Support: The Effects of Personally Delivered Campaign Messages on Voting Decisions and Opinion Formation

Kevin Arceneaux*

Temple University, Department of Political Science and Institute for Public Affairs, Faculty Affiliate, 453 Gladfelter Hall, 1115 West Berks Street, Philadelphia, PA 19122, USA; kevin.arceneaux@temple.edu

ABSTRACT

In addition to mobilizing supporters to vote, partisan campaigns use get-out-the-vote tactics as a means to boost support for their candidate. Although observational studies have attempted to estimate the effects of grassroots campaigning on political attitudes, they are unable to establish causality convincingly. Because campaigns strategically target potential supporters, comparing the attitudes of those whom campaigns contact to those they do not may only reveal spurious and biased relationships. In this paper I use a randomized field experiment to isolate the influence of personally delivered campaign messages on candidate support and attitude formation. I find that both door-to-door canvassing and commercial phone bank calls can have strong effects on voting preferences, but these tactics appear to have only weak effects on the actual beliefs that subjects possess about candidates and the degree to which those beliefs are weighted in their candidate preference. Although previous field experiments show that phone calls are less cost-effective at boosting turnout than door-to-door canvassing, they may be equally effective at increasing candidate support.

*I owe thanks to Don Green, Greg Huber, David Nickerson, the editors, anonymous reviewers, and seminar participants at Temple University, University of Connecticut, University of Memphis, University of South Carolina, Vanderbilt University, and Yale University for their helpful comments. The project would have been impossible to complete without generous funding from the Institution for Social and Policy Studies at Yale University and the cooperation and tireless efforts of Neri Holguin.

Supplementary electronic data for this article is available at <http://dx.doi.org/10.1561/100.00006003_supp.zip>

MS submitted 20 October 2005; final version received 4 October 2006

ISSN 1554-0626; DOI 10.561/100.00006003

© 2007 K. Arceneaux

Unsurprisingly, election campaigns concentrate an immense amount of effort on getting their supporters to vote. To attain this goal, campaigns use what is perhaps the most straightforward get-out-the-vote tactic: dispatch campaign workers, or the candidates themselves, to contact supporters in person and ask them to vote. This age-old form of labor-intensive electioneering continues to be used despite the ability of ever-present mass media to reach more people with greater efficiency. In the 2004 U.S. presidential elections, for instance, the two major-party presidential campaigns knocked on over 17 million doors and completed 50.7 million phone calls (Bergan et al. 2005).¹ The continued use of grassroots mobilization does not appear to be a mistake, either. An impressive amount of evidence collected across varied settings shows that grassroots campaigns, both partisan and non-partisan, can boost turnout through door-to-door canvassing (e.g. Arceneaux 2005; Gerber and Green 2000, 2001; Green and Gerber 2004; Green, Gerber, and Nickerson 2003; Michelson 2003; Vavreck, Spiliotes, and Fowler 2002) and volunteer phone banks (Nickerson 2005, 2006; Nickerson, Friedrich, and King 2006; Ramirez 2005; Rosenstone and Hansen 1993).

Although partisan campaigns use grassroots mobilization as a way to spur their supporters to the polls (and boost overall turnout as a consequence), they also have an interest in persuading others to support their candidate.² They expend a great deal of effort devising tactics and messages in an attempt to persuade voters that their candidate deserves trust and power. The persuasive aspect of grassroots campaigning has certainly not been overlooked by political scientists. Using observational survey data and, occasionally, public records, numerous studies report a positive correlation between contact with a partisan campaign and support for its candidate (Hillygus 2005; Rosenstone and Hansen 1993; Vavreck, Spiliotes, and Fowler 2002; Wielhouwer and Lockerbie 1994; but see Huckfeldt and Sprague 1992; Kramer 1970). On the surface, these findings suggest that partisan grassroots campaigns persuade as well as mobilize. Yet because these studies rely on observational data it is difficult to rule out plausible alternative explanations. For instance, partisan campaigns target citizens who are predisposed to support their candidate (Wielhouwer 2003), which raises the possibility that the positive relationship between partisan campaign contact and vote choice is a spurious one. Consequently, it is difficult to infer the causal effects of grassroots campaign contact from observational data without fully accounting for the selection process.³

Field experiments offer a more convincing approach to estimating the causal effects of personally delivered campaign messages. By randomly assigning campaign contact, this method allows researchers to compare treatment groups that are similar on both

¹ These figures do not include the millions of door knocks and phone calls conducted by issue advocacy groups (see Bergan et al. 2005 for an estimate).

² I use the term *partisan campaign* to refer to campaigns that advocate a particular candidate or issue.

³ To be fair, researchers in these studies attempt to address this issue by including a spate of covariates that predict both campaign contact and voting preferences. Unfortunately, it is unlikely that the demographic and attitudinal characteristics typically measured in these studies are capable of fully accounting for the selection process, leaving open the possibility of omitted variable bias (cf. Arceneaux, Gerber, and Green 2006).

observed and unobserved attributes. Observational methods, in contrast, can only balance on observables. Because randomization also reduces unobserved differences, post-study differences in treatment groups can be more confidently attributed to the campaign intervention.⁴ Although field experiments have been put to good use studying the effects of grassroots mobilization, they have mostly focused on get-out-the-vote efforts (Cardy 2005; Gerber and Green 2000, 2001; Green, Gerber, and Nickerson 2003; Michelson 2003; McNulty 2005; Nickerson, Friedrichs, and King 2006; Nickerson 2006). Only a few analyze the effects of partisan electioneering on candidate support; however, for various reasons they are unable to test fully the effects of personal contact.⁵ Gerber (2004) finds that negative direct mail messages increase the vote share of challengers (although not incumbents), but he does not study the effects of personally delivered messages. I studied the effects of door-to-door canvassing on precinct-level election outcomes in a separate study, but the experiment lacked sufficient power to discern significant effects (Arceneaux 2005). Using survey data, Nickerson (2005) finds no evidence that phone calls made by Democratic Party volunteers influenced vote preference in the 2002 Michigan gubernatorial race, but he cautions that the null results may have been caused by the campaign's failure to target effectively citizens open to voting Democratic.⁶

This study uses a field experiment to assess specifically whether personally delivered campaign messages can influence voting behavior. Following previous get-out-the-vote field experiments (e.g. Gerber and Green 2000), I analyze the effects of both face-to-face contact and phone calls placed by a paid phone bank. Observational studies, in contrast, lump together all forms of campaign contact and are unable to disentangle the effects of door-to-door canvassing from phone banks. Furthermore, field experimental evidence shows that paid phone calls tend to have small and often negligible effects on voter turnout (Green and Gerber 2004), but these studies do not address the impact of campaign phone calls on vote choice. Finally, I go beyond studying vote choice, and consider the effects of partisan campaign contact on individuals' beliefs about the candidate as well as the general opinion formation process. In doing so, I hope to contribute to the marriage of the grassroots mobilization literature with the broader study of campaign effects.

⁴ Random assignment reduces differences in treatment groups *within sampling variability*. It is still possible, if unlikely, to observe a spurious difference between treatment groups by chance. Standard frequentist methods are able to estimate the probability of such an occurrence (i.e. Pr[Type I error]) and, as discussed in the conclusion, it is the main reason why replication is an important ingredient to the external validity of experimental methods.

⁵ In addition to the studies mentioned here, there are a handful of studies purporting to be field experiments that study campaign contact and vote choice (Adams and Smith 1980; Blydenburgh 1971; Bochel and Denver 1971; Bositis, Baer, and Miller 1985), but they suffer debilitating methodological flaws. They either do not randomly assign contact or ignore random assignment when analyzing the data by comparing contacted individuals in the treatment group to the control group, which raises the specter of selection bias inherent in observational studies.

⁶ Michigan does not record party registration in the voter files, making it difficult for political parties to target its supporters. As a result, the Democratic Party ended up contacting just as many Republicans as Democrats, according to the survey data (Nickerson 2005).

GENERATING HYPOTHESES

The overriding goal of any partisan grassroots outreach campaign is to mobilize supporters and persuade others to support their candidate. In pursuit of this goal, campaigns craft and disseminate messages that attempt to influence citizens' beliefs about their candidate's issue positions and characteristics. While judging the persuasive effect of campaign contact is relatively straightforward – it either affects voting preferences and behavior or it does not – campaign messages may affect beliefs about the candidates in multiple ways. Message effects have been written about extensively in the broader campaign literature, and it is not possible to detail completely the scholarship and all of its nuances here. Instead, I concentrate on three high-profile hypotheses that have received a great deal of attention in previous work and consider how they apply to personally delivered campaign messages.

The Hypodermic Stimulus

The initial scholarship on campaign effects conducted in the 1940s and 1950s failed to find evidence that exposure to campaign appeals in the mass media causes voters to alter their voting preferences or beliefs about candidates' issue positions and personality characteristics (e.g. Berelson, Lazarsfeld, and McPhee 1954; Klapper 1960). Although the debate over the effects of campaign messages remains unsettled, subsequent research continues to find inconsistent evidence on whether political advertisements do much to change people's beliefs and opinions about candidates (Goldstein and Ridout 2004). Yet even if political advertising has minimal effects, it is conceivable that the impersonal way in which mass-based appeals are delivered reduces their ability to affect beliefs and voting preferences. Perhaps campaign messages are able to influence directly people's beliefs about candidates when the messages are delivered in person. Before individuals can accept a message they must first receive it (McGuire 1968; Zaller 1992). It is a simple matter to tune out a political commercial, but quite another to ignore a campaign worker in a face-to-face encounter in one's doorway. Furthermore, the personal nature of social contact may make it a more salient and potent form of communication. Indeed, the early proponents of the minimal effects model carved out a special role for the influence of social contact on voting decisions (Berelson, Lazarsfeld, and McPhee 1954).

Priming

Other scholars argue that campaign and media messages do not directly alter individuals' beliefs about the candidates inasmuch as they cause some considerations to be more accessible than others in the minds of voters (e.g. Iyengar and Kinder 1987; Krosnick and Kinder 1990).⁷ As a result, individuals place more weight on the considerations

⁷ Miller and Krosnick (2000) propose a different causal mechanism for media news priming. They contend that news stories prime attitudes not by increasing the accessibility of particular considerations, but by causing people to infer that the issues addressed in news stories are more important and thereby deserve greater weight in their attitudes. They do not address whether this logic also applies to priming caused by campaign messages.

raised by campaign messages in their evaluation of the candidates. For instance, revelations of Gary Hart's extra-marital affair prior to the 1988 U.S. presidential primary primed voters to weight attitudes about personal morality more heavily in their assessment of Hart (Stoker 1993). Moreover, campaigns may develop messages with more of an eye toward priming than belief change. Jacobs and Shapiro (1994) cleverly employ private polling data collected by the 1960 Kennedy presidential campaign to demonstrate that the campaign strategically emphasized popular issues in the hope that voters would attach more weight to those issues when evaluating his candidacy.

Opinionation

Irrespective of whether campaign messages affect beliefs about the candidates or prime particular considerations, some scholars contend that exposure to a message causes individuals, at the very least, to form *an* opinion, any opinion, on the subject addressed in the message (Ansolabehere, Behr, and Iyengar 1991; Kim, Scheufele, and Shanahan 2005; Zaller 1992, 205). For instance, some individuals may not believe a campaign message that contends Candidate A is tough on crime, yet in doing so they do form an opinion about whom they believe is tougher on crime (e.g. Candidate B or perhaps none). Consequently, when asked to venture an opinion about which candidate offers the most draconian crime policies, individuals who received the campaign message may not be more likely to opine that it is Candidate A, nor may they place more weight on their beliefs about crime policy in their evaluation of Candidate A, but they should be more likely to simply offer an opinion on who is tough on crime.

STUDY DESIGN

Background

In the spring of 2004 I worked with the campaign for Deanna Archuleta-Loeser, who was running in an open-seat closed Democratic primary for a county commissioner seat in Bernalillo County, New Mexico. The district included Albuquerque and is heavily Democratic. The winner of the primary would not face Republican opposition in the fall general election. Archuleta-Loeser had never held elective office before and faced two opponents in the primary. One of her opponents, Judy Woodward, ran an organized campaign and enjoyed name recognition, having held a prominent local elected office.⁸ In the end, Archuleta-Loeser (henceforth, the candidate) won a plurality of votes (48.5% to Woodward's 36%).

The low-salience nature of the electoral environment offers several advantages. First, the lack of intense news coverage reduces the amount of background noise. Many citizens may have learned about these candidates mostly through contact with the campaign.

⁸ According to newspaper accounts, the third candidate, John Eder, was also well known but did not run a serious campaign.

Second, because the candidates were competing in a closed primary, partisanship is removed as a factor in the voting decision. As a result, if any campaign message effects are found, it is possible to rule out the alternative explanation that subjects were merely responding to partisan cues.⁹ These electoral features provide a sensible place to conduct an initial field experiment of this topic. If specific campaign messages do not influence vote preference and opinion formation in this setting, it is unlikely that they would have much effect in noisier settings in which any given message is likely to be drowned out by the greater amount and variety of competing messages. Finally, this field experiment generalizes to a broader swath of the U.S. electoral landscape than studies of highly visible elections, such as presidential or senatorial contests. Across all levels of governments, there are over 500,000 elected offices (Bureau of the Census 1995). The vast majority of the elections associated with these offices are quiescent affairs, and many lack partisan cues because they are decisive primary contests like this one (e.g. one-party congressional districts) or because they are non-partisan.

Protocol

The field experiment was designed to study the effect of door-to-door canvassing and phone calls on voting behavior and opinion formation. The campaign identified a list of 18,512 registered Democrats living in all 101 precincts in the County Commissioner district who they believed to be likely primary voters. The campaign slated 60 precincts, which included 17,175 of the individuals on the list, to be eligible for the door-to-door canvassing experiment. The target universe was stretched too thin over the remaining precincts to justify canvassing. Of the 60 precincts, 11 were assigned to be canvassed by the candidate herself (3,227 individuals), 21 were assigned to be canvassed by paid campaign workers (5,623 individuals), and the remaining 28 were assigned to the control group (8,325 individuals).¹⁰ For the phone treatment, the list of 3,844 individuals with phone numbers in the target universe was randomly divided between treatment and control groups with 3,189 assigned to receive a telephone call from a paid phone bank and 655 assigned to the control group.¹¹ Nearly a third of the individuals ($n = 1,337$) lived in precincts that were outside the canvassing experiment, while the remaining two-thirds ($n = 2,507$) were randomly distributed across treatment and control groups in the canvassing experiment. Because of random assignment the overlap in experimental treatment groups does not pose a threat to inference. See Table 1 for the breakdown of assignment by conditions.

⁹ Of course, this does not control for other types of cues, such as ethnicity, age, or gender. Yet given the central role that partisanship plays in voting decisions (Campbell et al. 1960; Green, Palmquist, and Schickler 2002), it is important to rule out partisan cue-taking, as opposed to persuasion, as the causal mechanism. Also, it is worth noting that the candidate's major competitor was also a woman.

¹⁰ Prior to the experiment, the campaign conducted an initial walk of some precincts. To ensure that the pre-experimental contact does not confound the experimental stimuli, I stratified the random assignment by previous contact.

¹¹ Care was taken so that individuals in the same household were not split across groups. If an individual was selected to be in the treatment group, all individuals in that person's household were also included.

Table 1. Treatment assignment summary.

Door-knock assignment	Number of Precincts	Phone assignment			Total
		Phone call	Control group	Not assigned	
Candidate visit	11	396	83	2,748	3,227
Field visit	21	694	124	4,805	5,623
Control group	28	981	229	7,115	8,325
Not assigned	41	1,118	219	–	1,337
Total	101	3,189	655	14,668	18,512

Door-to-door canvassers were directed to their assigned precincts armed with the list of registered Democrats' addresses. Canvassers were trained to read a persuasive script that emphasized the candidate's support for "planned growth" to limit urban sprawl (a popular issue in the view of the campaign) and asked citizens to vote for her. Canvassers also left a glossy campaign brochure that included information on the candidate's issues positions and, in many places (including the large title), claimed that the candidate is a "strong leader." The brochures were left with individuals to whom the canvassers spoke and at houses where the subjects were not available. The candidate followed a similar script as the field canvassers, but she was free to modify and personalize her message if she deemed necessary. The canvassers worked all of the precincts assigned to the treatment group. The paid phone bank called subjects in the phone treatment group and read an interactive script that also emphasized the candidate's support for planned growth, but there was no mention of the candidate's leadership qualities.¹² Following the primary election, a commercial phone bank conducted a survey of the 3,881 registered Democrats in the treatment and control groups for whom the campaign had phone numbers, completing 988 interviews (25.5% survey response rate) with 685 respondents in the canvassing experiment, 976 in the phone experiment, and 673 in both. (After the experiment began, the campaign discovered an additional 37 phone numbers for individuals in its target universe, and they were included in the survey sample.)

Table 2 displays summary statistics of demographic characteristics and vote history within the treatment groups for the target populations and survey samples of both experiments. Beginning with the canvassing experiment, there are virtually no pairwise imbalances in the observed covariates for the target population and the joint test of significance fails to reject the null hypothesis that the treatment groups are drawn from the same population as the

¹² These activities were the bulk of the campaign. Activities outside of the experimental protocol, such as direct mail, were randomly spread across the treatment and control groups. Consequently, if non-experimental activities did have effects, random assignment causes them to be the same in each group, and they cancel out when estimating treatment effects.

Table 2. Summary statistics for canvassing and phone experiment.

Variable	Canvassing experiment						Phone experiment			
	Target population			Sample respondents			Target population		Sample respondents	
	Control	Candidate	Field	Control	Candidate	Field	Control	Treatment	Control	Treatment
Percent Vote 1992	29.2	30.4	29.2	61.0	62.2	62.9	52.7	57.1	61.2	59.9
Primary	(1.5)	(1.8)	(1.3)	(2.5)	(4.5)	(4.6)	(2.0)	(0.9)	(3.4)	(1.8)
Percent Vote 1992	52.0	53.6	52.3	80.9	76.0	79.5	68.9	72.0	79.1	75.9
General	(2.3)	(2.3)	(1.7)	(1.6)	(3.4)	(3.4)	(1.8)	(0.8)	(2.9)	(1.5)
Percent Vote 1994	30.5	30.6	30.2	66.1	60.4	74.5	58.8	61.1	69.2	66.7
Primary	(1.6)	(1.6)	(1.4)	(2.1)	(4.3)	(4.9)	(1.9)	(0.9)	(3.3)	(1.7)
Percent Vote 1994	46.4	47.2	46.2	80.8	71.7	82.7	69.9	72.8	81.6	77.0
General	(2.2)	(2.3)	(1.7)	(1.7)	(4.2)	(3.4)	(1.8)	(0.8)	(2.7)	(1.5)
Percent Vote 1996	24.0	22.6	23.3	65.5	49.9	63.0	52.5	54.9	63.2	58.1
Primary	(1.3)	(1.4)	(1.1)	(2.7)	(3.9)	(4.3)	(2.0)	(0.9)	(3.4)	(1.8)
Percent Vote 1996	55.6	57.7	55.9	83.7	82.3	86.8	74.2	78.2	83.1	82.1
General	(2.3)	(2.4)	(1.5)	(2.1)	(1.9)	(3.2)	(1.7)	(0.7)	(2.7)	(1.4)
Percent Vote 1998	29.9	29.5	29.0	72.5	71.2	73.8	62.9	66.8	72.6	71.5
Primary	(1.6)	(1.7)	(1.5)	(2.4)	(3.6)	(3.3)	(1.9)	(0.8)	(3.2)	(1.6)
Percent Vote 1998	57.5	58.4	57.0	89.3	85.3	89.8	77.4	81.8	86.6	85.9
General	(2.3)	(2.0)	(1.8)	(1.6)	(2.2)	(2.2)	(1.6)	(0.7)	(2.4)	(1.2)
Percent Vote 2000	27.2	28.4	26.7	88.0	88.9	90.2	79.4	83.1	89.1	87.0
Primary	(1.4)	(1.3)	(1.1)	(2.0)	(1.9)	(2.1)	(1.6)	(0.7)	(2.2)	(1.2)
Percent Vote 2000	70.7	72.7	72.0	91.0	90.7	92.9	82.0	85.0	91.5	89.2
General	(1.7)	(1.7)	(1.2)	(1.4)	(1.7)	(1.4)	(1.5)	(0.6)	(2.0)	(1.1)
Percent Vote 2002	30.6	32.2	31.3	91.7	90.6	91.8	82.1	86.1	92.5	89.7
Primary	(1.4)	(1.3)	(1.2)	(1.4)	(2.2)	(1.8)	(1.5)	(0.6)	(1.9)	(1.1)
Percent Vote 2002	62.0	64.3	62.2	92.2	90.8	92.5	82.1	86.0	93.0	89.8
General	(1.8)	(1.8)	(1.4)	(1.2)	(2.2)	(1.6)	(1.5)	(0.6)	(1.8)	(1.1)
Percent female	58.3	59.3	57.5	59.1	54.2	50.9	54.2	54.9	53.2	55.4
	(0.4)	(0.5)	(0.8)	(3.5)	(4.6)	(4.4)	(1.9)	(0.9)	(3.5)	(1.8)

Percent newly registered	7.9 (0.7)	8.1 (0.9)	7.7 (0.6)	6.4 (1.1)	7.6 (1.8)	6.4 (1.5)	15.1 (1.4)	11.8 (0.6)	5.0 (1.5)	8.6 (1.0)
Mean age	50.2 (1.2)	50.6 (1.2)	50.1 (0.9)	64.2 (1.3)	61.2 (2.2)	64.6 (1.5)	59.4 (0.7)	60.3 (0.3)	63.9 (1.2)	64.4 (0.6)
Mean household size	1.9 (0.1)	2.0 (0.1)	1.9 (0.03)	1.8 (0.1)	1.7 (0.1)	1.8 (0.1)	1.6 (0.03)	2.0 (0.02)	1.6 (0.06)	1.8 (0.03)
Number of observations*	28	11	21	27	10	21	655	3,189	201	775
Treatment contact rate	0%	7.1%	10.6%	0%	33.1%	44.9%	0%	37.6%	0%	62.8%
Joint test of significance	$F(20, 39) = 0.815, p = 0.682$		$F(20, 39) = 0.304, p = 0.997$	$F(20, 37) = 1.019, p = 0.465$	$F(20, 37) = 0.903, p = 0.586$	$F(15, 3,828) = 0.7146, p = 0.772$	$F(15, 960) = 1.009, p = 0.443$			

* In the canvassing experiment the number of observations refers to the number of precincts. Fixed effects for strata were included in the joint test of significance for the canvassing experiment. Standard errors are in parentheses.

control group (candidate treatment, $p = 0.682$; field treatment, $p = 0.997$).¹³ In the survey sample there are some imbalances in vote history for the 1994 elections and the 1996 primary.¹⁴ Pairwise imbalances in random samples can happen by chance, and the joint test is statistically insignificant (candidate treatment, $p = 0.465$; field treatment, $p = 0.586$). Inferences are unaffected by including these unbalanced covariates in subsequent analyses.

There is more evidence of pairwise imbalance in the target population for the phone experiment, but this has to do with the design of the experiment and not a failure in randomization. Individuals who live in multi-person households were more likely to be placed in the treatment group. The appropriate randomization check is an incremental F -test, which tests whether the inclusion of covariates significantly predicts assignment beyond household size, and it does not ($p = 0.772$). With the exception of the newly registered indicator there is no evidence of pairwise imbalance in the survey sample, and the incremental F -test is statistically insignificant ($p = 0.443$). Since household size is correlated with assignment by design, it is imperative that it is included as a covariate in subsequent statistical analyses.

The available evidence, then, suggests that there are no serious flaws in the randomization (or odd random draws) in both the target populations and survey samples for both experiments, which bolsters their internal validity. Despite the relatively low response rate in the post-election survey, it is still possible to make valid inferences from these data, because the treatment and control groups in the survey sample still contain individuals who, prior to the study, were just as likely to vote, prefer a particular candidate, possess the same knowledge about campaign issues, and so on.¹⁵ Yet differences between the survey samples and their target populations shown in Table 2 strongly suggest limitations in the generalizability of inferences based on the survey sample. In particular, relative to the target population, survey-takers were far more likely to have been contacted by a canvasser or the phone bank.¹⁶ Survey-takers are ostensibly more likely to be home and answer calls from strangers, and they are also older and more likely to vote. Consequently, the findings below generalize only to the population of people who are likely to be exposed to grassroots campaign messages.

Estimating Quantities of Interest

The campaign, like all campaigns, was able to reach only a portion of those initially targeted. The failure to treat some subjects does not pose a threat to causal inference if the

¹³ Females are slightly less likely to be in the field group relative to the candidate ($t = -1.91$, $p = .06$, two-tailed).

¹⁴ Subjects in the field group were more likely to vote in the 1994 primary relative to those in the candidate group ($t = 2.158$, $p = 0.04$); those in the candidate group were less likely to vote in both the 1994 general and 1996 primary elections relative to the control group ($t = -2.027$, $p = 0.05$ for 1994; $t = -3.323$, $p = 0.001$ for 1996) and field group ($t = -2.039$, $p = 0.05$ for 1994; $t = -2.271$, $p = 0.03$ for 1996). All p -values are two-tailed.

¹⁵ The survey response rates in the treatment and control groups are not statistically different from one another in both the canvassing experiment ($F[2, 57] = 0.09$, $p = 0.918$) and phone experiment ($\chi^2 = 0.25$, $p = 0.618$), which, taken together with the results from the randomization checks, bolsters this claim.

¹⁶ Subjects are coded as contacted if they actually spoke with a campaign canvasser or phone bank caller. Individuals who were not home or did not answer the door or phone are coded as not contacted.

data are analyzed appropriately. The intent-to-treat (ITT) effect compares everyone in the treatment group irrespective of contact (i.e. those whom the campaign *intended* to treat) to everyone in the control group. Because these groups were formed randomly, there should be no differences in outcomes, and to the extent there are differences, it is possible to attribute them to the effect of the treatment. Substantively, the ITT estimates the number of individuals a campaign persuaded or mobilized per the number of individuals they attempted to contact. The average-treatment-on-treated (ATT) effect estimates the number of individuals a campaign persuaded or mobilized per the number of actual contacts, and measures the effect of being exposed to a campaign message. The ATT estimand is properly estimated by using random assignment as an instrument to compare those who were actually contacted to those who were not (see Angrist, Imbens, and Rubin 1996; Gerber and Green 2000).¹⁷

Because random assignment for the canvassing experiment was conducted at the precinct level, it is inappropriate to estimate the standard errors in the individual-level analysis using the number of subjects as the total n . The correlation in individual behavior within precincts lowers the effective n , which causes the standard errors to grow. Failure to account for precinct randomization underestimates the uncertainty around the treatment effect and the probability of committing a Type I error. Consequently, I estimate the standard errors for point estimates in the door-to-door canvassing experiments using a robust covariance matrix estimator that clusters on precinct (cf. Arceneaux 2005).¹⁸

FINDINGS

The post-election survey asked subjects which candidate they supported in the county commissioner race (see the Appendix for question wording). Respondents who selected the candidate from the response items are coded as supporters (vote preference = 1), and those who chose a different candidate or who said that they “did not vote in this race” are coded as non-supporters (vote preference = 0). Respondents who volunteered that they “did not know” or refused to answer the question are removed from the analyses. This decision allows the dependent variable to remain dichotomous and, thus, enables me to estimate ATT effects. As Supplemental Appendix A shows, the results do not change substantively if don’t know and refusal responses are coded separately and ITT effects are estimated via multinomial logit.

I estimate the effects of door-to-door canvassing and campaign phone calls on support for the candidate by regressing the vote preference indicator on the treatment indicators

¹⁷ It would be inappropriate to ignore random assignment when comparing those contacted by the campaign to those who were not. This approach treats the data as if they were observational and may introduce bias into causal estimates. People whom the campaign reached may simply be different from those who were not contacted in ways that are correlated with the dependent variable.

¹⁸ Also, the canvassing models are weighted by the proportion of the sampled precinct that completed the survey to take into account the variability of precinct size.

and covariates obtained from the voter file that measure sex, recent registration status, age, household size, and vote history in the previous primary. Canvassing experiment models also include fixed effects for strata (see footnote 10), and vote history for the 1994 primary, 1994 general, and 1996 primary, which showed evidence of imbalance in the survey sample. Because the canvassing and phone experiments were conducted independently of one another and only overlap partially, I estimate separate models for each. Yet even if the sample is restricted to just those respondents who were in both experiments, the substantive results do not change (see Supplemental Appendix B). The results are displayed in the first four columns of Table 3.

Both the door-to-door canvassing and commercial phone bank efforts increased support for the candidate among subjects in the survey sample. Respondents living in precincts canvassed by the candidate were 13.3 percentage points more likely to support her than those in the control group ($p = 0.016$, one-tailed t -test), and those living in precincts canvassed by campaign workers were 8.6 percentage points more likely to support the candidate ($p = 0.078$, one-tailed t -test). Similarly, subjects assigned to receive phone calls were 11.7 percentage points more likely to support the candidate than the control group ($p = 0.009$, one-tailed t -test). Among those who were actually reached by the campaign, the candidate boosted her support by 42.3 percentage points ($p = 0.015$, one-tailed t -test), the field workers by 18.3 percentage points ($p = 0.085$, one-tailed t -test), and the phone bank by 18.6 percentage points ($p = 0.009$, one-tailed t -test). Although the candidate's visits had by far the largest effect on her support, I cannot rule out the null hypothesis that it is no different from the field worker and phone effects.

In addition to boosting support for the candidate, campaigns use electioneering in the hope that it mobilizes their supporters (and not their non-supporters) on Election Day. Combining the survey responses with official voting records, I placed respondents into four categories: (1) those who did not vote, (2) supporters who voted, (3) non-supporters who voted, and (4) non-responders who voted. Using multinomial logit, I regressed this variable on the same set of treatment indicators and covariates employed in the vote preference analyses. Because I am using multinomial logit, it is only possible to estimate the ITT effects. Results are reported in the remaining columns of Table 3. There is modest evidence that the candidate and the phone calls mobilized supporters, but the size of the field canvassing standard error makes its ITT effect estimates too imprecise to rule out chance as an explanation. Respondents living in precincts canvassed by the candidate were 10.1 percentage points more likely to vote for her relative to the control group ($p = 0.08$, one-tailed t -test) and individuals in the phone group were 6.9 percentage points more likely to vote for her than the control group ($p = 0.103$, one-tailed t -test). The campaign did not mobilize individuals who expressly opposed the candidate.

Did the campaign also succeed in convincing subjects to accept its messages about the candidate? Recall that the door-to-door canvassers talked to subjects about the candidate's support for planned growth as well as their belief that she is a strong leader. The post-election survey asked respondents to identify the candidate in the race who they believed to be the strongest supporter of planned growth and best described by the phrase "strong leader" (see the Appendix for question wording). On each item, respondents who chose the candidate are coded as 1 and those who did otherwise are coded as 0. ITT and ATT

Table 3. The effect of personally delivered campaign messages on vote preference and mobilization.

	Vote preference				Vote preference and mobilization					
	ITT estimates		ATT estimates		Canvassing experiment			Phone experiment		
	Canvassing experiment	Phone experiment	Canvassing experiment	Phone experiment	For vs. abstain	Against vs. abstain	DK/RF vs. abstain	For vs. abstain	Against vs. abstain	DK/RF vs. abstain
Candidate treatment	0.133 (0.06)		0.423 (0.19)		0.413 (0.29)	0.027 (0.33)	0.356 (0.28)			
Field treatment	0.086 (0.06)		0.183 (0.13)		0.311 (0.30)	0.123 (0.32)	0.037 (0.24)			
Phone treatment		0.117 (0.05)		0.186 (0.08)				0.284 (0.23)	-0.270 (0.24)	0.371 (0.24)
Female	0.090 (0.04)	0.040 (0.04)	0.066 (0.04)	0.033 (0.04)	0.267 (0.21)	-0.219 (0.25)	-0.200 (0.29)	0.135 (0.18)	-0.229 (0.21)	-0.134 (0.19)
Newly registered	0.231 (0.16)	0.150 (0.16)	0.331 (0.15)	0.133 (0.16)	0.338 (0.80)	-0.530 (0.86)	1.294 (1.29)	0.074 (0.77)	-0.347 (0.76)	0.752 (0.88)
Age	-0.001 (0.00)	-0.001 (0.00)	-0.001 (0.00)	-0.002 (0.00)	0.000 (0.01)	0.006 (0.01)	0.023 (0.01)	0.002 (0.01)	0.012 (0.01)	0.026 (0.01)
Household size	0.057 (0.03)	0.035 (0.02)	0.056 (0.03)	0.042 (0.02)	0.109 (0.12)	-0.100 (0.14)	-0.061 (0.11)	0.003 (0.10)	-0.144 (0.14)	-0.200 (0.11)
Voted 2002	0.339 (0.14)	0.252 (0.13)	0.413 (0.11)	0.231 (0.14)	1.761 (0.65)	0.181 (0.65)	1.914 (1.08)	1.611 (0.68)	0.596 (0.65)	1.654 (0.82)
Voted 1994	-0.097 (0.06)		-0.103 (0.06)		0.443 (0.27)	0.882 (0.38)	1.200 (0.43)			
Primary 1994	0.085 (0.07)		0.073 (0.07)		-0.613 (0.38)	-1.090 (0.48)	-1.377 (0.38)			
General 1996	0.050 (0.05)		0.033 (0.05)		0.676 (0.24)	0.779 (0.28)	0.457 (0.29)			
Primary 1996	0.046 (0.21)	0.293 (0.17)	0.018 (0.20)	0.343 (0.17)	-2.096 (0.97)	-0.418 (0.90)	-3.232 (1.28)	-1.515 (0.86)	-0.998 (0.89)	-3.021 (0.98)
<i>n</i>	437	603	437	603	685			976		
Adjusted R^2 or pseudo R^2	0.04	0.02	0.02	0.02	0.06			0.03		
F or χ^2	3.610	2.723	5.372	2.664	106.344			84.640		
p -value	0.000	0.013	0.000	0.015	0.000			0.000		

Note: Clustered standard errors in parentheses. Vote preference ITT effects estimated with OLS. ATT effects estimated with two-stage least squares, using random assignment as an instrument for campaign contact. Vote preference/mobilization models estimated using multinomial logit. Fixed effects for strata not shown in canvassing models.

effects are shown in Table 4. Although it appears that the candidate increased the probability of treatment group respondents choosing her as the champion of planned growth and a strong leader, the field canvass and phone calls had smaller, statistically insignificant effects.¹⁹ Individuals who were contacted by the candidate were especially likely to come away with a positive view of her leadership skills – 41 percentage points more likely than the control group ($p = 0.019$, one-tailed t -test). In light of the weak effects by the other modes of campaigning it is likely that, by meeting the candidate in person, subjects developed a positive overall impression of her rather than a straightforward acceptance of her campaign messages.²⁰

Aside from (or even in lieu of) influencing individuals' beliefs about the candidates, scholars have posited that campaign messages prime voters to weight some beliefs more than others in their voting decisions. Following the approach that previous studies use (e.g. Miller and Krosnick 2000), I investigate this hypothesis by regressing vote preferences on the interaction between treatment indicators and respondents' beliefs about the candidate. The ITT and ATT estimates are reported in Table 5.²¹ The priming hypothesis is supported if the interaction term and the combined effect (interaction term + direct effect) are positive. These criteria are met only in the phone experiment models. Relative to the control group, respondents in the phone treatment group gave more than twice as much weight to their beliefs about the candidate's support for planned growth. Specifically, control group respondents who believed the candidate to be the strongest champion of planned growth were 14.6% more likely to support her relative to those who did not deem her the strongest proponent of the issue ($p = 0.045$, one-tailed t -test), whereas respondents in the phone treatment group were 31.4% (37.5% if they were actually contacted) more likely to support her if they perceived her to be a supporter of planned growth ($p < 0.0001$, one-tailed t -test).²²

Also, beliefs about the candidate's leadership abilities are more weakly related to vote preference in the phone treatment group than in the control group. Because the phone callers did not address the leadership qualities of the candidate, these statistically significant negative findings are not unexpected. Curiously, though, there is no evidence of priming in the canvassing experiment. Although the explanation cannot be ruled out without collecting additional evidence, it seems unlikely that phone calls would be more effective than door-to-door canvassing at priming. The possibility remains that either

¹⁹ At best, the field canvass treatment effect is of borderline statistical significance ($p = 0.111$, one-tailed t -test).

²⁰ I thank an anonymous reviewer for pointing out this parsimonious alternative explanation. It would be possible to test this in a replication of the experiment by asking survey respondents to rate the candidate on popular issue positions that he or she did not explicitly take in the campaign. If the message is the mechanism, as opposed to the interpersonal contact, candidate contact should affect only those issue items addressed by him or her.

²¹ It was not possible to estimate multinomial logit models that included respondents who did not answer the vote preference question, because the interaction terms caused there to be insufficient observations in some cells of the analysis.

²² The standard errors for the combined effects in the ITT and ATT models are 0.036 and 0.068, respectively.

Table 4. The effect of personally delivered campaign messages on the acceptance of campaign messages.

	ITT estimates						ATT estimates					
	Planned growth			Strong leader			Planned growth			Strong leader		
	Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment	
Candidate treatment	0.060 (0.036)			0.135 (0.049)			0.184 (0.121)			0.410 (0.192)		
Field treatment	0.040 (0.036)			0.055 (0.044)			0.091 (0.077)			0.126 (0.100)		
Phone treatment		0.029 (0.032)			-0.034 (0.037)			0.046 (0.051)				-0.053 (0.058)
Female	0.014 (0.032)	0.006 (0.027)		0.109 (0.029)	0.078 (0.029)		0.008 (0.032)	0.006 (0.027)		0.096 (0.030)	0.079 (0.029)	
Newly registered	-0.022 (0.105)	-0.015 (0.102)		0.019 (0.095)	-0.068 (0.118)		0.015 (0.118)	-0.021 (0.103)		0.091 (0.100)	-0.061 (0.118)	
Age	-0.003 (0.001)	-0.002 (0.001)		-0.003 (0.001)	-0.003 (0.001)		-0.003 (0.001)	-0.003 (0.001)		-0.003 (0.001)	-0.002 (0.001)	
Household size	0.083 (0.021)	0.064 (0.016)		0.065 (0.025)	0.054 (0.017)		0.082 (0.021)	0.066 (0.016)		0.064 (0.025)	0.052 (0.017)	
Voted 2002	0.112 (0.084)	0.095 (0.087)		0.026 (0.070)	-0.012 (0.104)		0.148 (0.094)	0.087 (0.088)		0.100 (0.081)	-0.003 (0.104)	
Primary 1994	0.010 (0.043)			-0.040 (0.038)			0.002 (0.044)			-0.057 (0.041)		
Primary 1994	0.012 (0.064)			0.089 (0.057)			0.014 (0.063)			0.095 (0.058)		
General 1996	0.070 (0.031)			0.049 (0.040)			0.064 (0.030)			0.036 (0.043)		
Primary Constant	0.092 (0.131)	0.162 (0.115)		0.146 (0.131)	0.370 (0.131)		0.081 (0.138)	0.176 (0.115)		0.123 (0.136)	0.354 (0.129)	
<i>n</i>	685	976		685	976		685	976		685	976	
Adjusted <i>R</i> ²	0.04	0.03		0.04	0.02		0.04	0.03		0.01	0.01	
<i>F</i>	5.349	5.370		7.106	4.770		5.633	5.374		8.866	4.753	
<i>p</i> -value	0.000	0.000		0.000	0.000		0.000	0.000		0.000	0.000	

Note: Clustered standard errors in parentheses. ITT effects estimated with OLS. ATT effects estimated with two-stage least squares, using random assignment as an instrument for campaign contact. Fixed effects for strata not shown in canvassing models.

Table 5. Testing for priming effects.

	ITT estimates		ATT estimates	
	Canvassing experiment	Phone experiment	Canvassing experiment	Phone experiment
Candidate treatment	0.000 (0.088)		-0.002 (0.301)	
Field treatment	0.029 (0.060)		0.071 (0.159)	
Phone treatment		0.127 (0.052)		0.210 (0.087)
Planned growth	0.236 (0.062)	0.146 (0.086)	0.238 (0.063)	0.145 (0.086)
Strong leader	0.503 (0.062)	0.664 (0.066)	0.503 (0.062)	0.661 (0.066)
Candidate \times growth	-0.053 (0.079)		-0.180 (0.309)	
Field \times growth	0.057 (0.083)		0.063 (0.156)	
Phone \times growth		0.168 (0.093)		0.230 (0.137)
Candidate \times leader	0.069 (0.126)		0.230 (0.483)	
Field \times leader	-0.032 (0.083)		-0.046 (0.162)	
Phone \times leader		-0.233 (0.077)		-0.349 (0.116)
Female	0.050 (0.032)	0.007 (0.030)	0.047 (0.032)	0.004 (0.030)
Newly registered	0.131 (0.097)	0.116 (0.084)	0.163 (0.097)	0.154 (0.090)
Age	0.001 (0.001)	0.000 (0.001)	0.001 (0.001)	0.000 (0.001)
Household size	-0.004 (0.016)	-0.005 (0.015)	-0.008 (0.016)	-0.001 (0.015)
Voted 2002 Primary	0.242 (0.095)	0.203 (0.068)	0.263 (0.086)	0.223 (0.075)
Voted 1994 Primary	-0.088 (0.051)		-0.094 (0.049)	
Voted 1994 General	0.029 (0.057)		0.026 (0.055)	
Voted 1996 Primary	0.029 (0.041)		0.032 (0.042)	

Table 5. Continued.

	ITT estimates		ATT estimates	
	Canvassing experiment	Phone experiment	Canvassing experiment	Phone experiment
Constant	-0.034 (0.138)	0.013 (0.100)	-0.042 (0.141)	0.020 (0.107)
<i>n</i>	437	603	437	603
Adjusted <i>R</i> ²	0.47	0.46	0.47	0.45
<i>F</i>	54.601	82.462	58.679	79.948
<i>p</i> -value	0.000	0.000	0.000	0.000

Note: Dependent variable 5 vote preference. Clustered standard errors in parentheses. ITT effects estimated with OLS. ATT effects estimated with two-stage least squares, using random assignment as an instrument for campaign contact. Fixed effects for strata not shown in canvassing models.

the lack of evidence in the canvassing experiment or the presence of evidence in the phone experiment is due to sampling variability.

The models reported in Table 6 suggest that, even if the campaign did little to change respondents' beliefs about the candidate or the weights attached to those beliefs, it may have encouraged them to form opinions about the issues addressed by the campaign messages. In these analyses the dependent variable is coded as 1 if respondents expressed an opinion on either the planned growth or strong leader survey items and 0 if they said "don't know." Respondents who simply did not answer the questions are removed from the analyses, but the results do not change if they are included (see Supplemental Appendix A). Respondents in the field canvassing group were approximately 8 percentage points (16–17 percentage points if contacted) more likely than the control group to express an opinion on both items ($p = 0.059$ and 0.047 for planned growth ITT and ATT; $p = 0.039$ and 0.034 for strong leader ITT and ATT, one-tailed t -tests). The candidate had an effect only on opinions about leadership qualities ($p = 0.013$ for ITT and 0.029 for ATT, one-tailed t -tests), and again because the phone bank did not address the leadership issue, it had an effect only on opinions about the candidates' position on planned growth ($p = 0.051$, one-tailed t -tests). Therefore, with one inexplicable exception, the evidence points to consistent opinionation effects across campaigning modes.

Of course, it remains unspecified whether the specific campaign messages caused treatment group respondents to form opinions on these issues or if campaign contact merely increased interest in the campaign and, as a result, caused treatment group subjects to seek information on their own. No statistically significant relationship exists between treatment assignment or campaign contact and expressing an interest in getting involved in the campaign (see Supplemental Appendix C), which suggests that the campaign

Table 6. Effect of personally delivered campaign messages on opinionation.

	ITT estimates						ATT estimates					
	Planned growth			Strong leader			Planned growth			Strong leader		
	Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment		Canvassing experiment	Phone experiment	
Candidate treatment	-0.006 (0.043)			0.136 (0.060)			-0.011 (0.119)			0.383 (0.199)		
Field treatment	0.077 (0.049)			0.080 (0.045)			0.164 (0.096)			0.176 (0.094)		
Phone treatment		0.073 (0.044)			-0.048 (0.041)			0.113 (0.069)			-0.074 (0.063)	
Female	-0.026 (0.039)	-0.017 (0.037)		0.064 (0.040)	0.024 (0.035)		-0.022 (0.040)	-0.019 (0.037)		0.051 (0.041)	0.026 (0.035)	
Newly registered	-0.154 (0.160)	-0.127 (0.160)		-0.242 (0.114)	-0.305 (0.099)		-0.125 (0.160)	-0.136 (0.161)		-0.165 (0.116)	-0.298 (0.100)	
Age	-0.003 (0.002)	-0.003 (0.001)		-0.003 (0.001)	-0.001 (0.001)		-0.003 (0.002)	-0.003 (0.001)		-0.003 (0.002)	-0.001 (0.001)	
Household size	0.061 (0.026)	0.063 (0.020)		0.038 (0.023)	0.032 (0.019)		0.058 (0.026)	0.067 (0.020)		0.037 (0.024)	0.029 (0.019)	
Voted 2002 Primary	0.098 (0.118)	0.054 (0.146)		-0.192 (0.052)	-0.218 (0.078)		0.116 (0.114)	0.041 (0.146)		-0.123 (0.064)	-0.209 (0.078)	
Voted 1994 Primary	0.120 (0.064)			0.061 (0.064)			0.121 (0.064)			0.045 (0.067)		
Voted 1994 General	-0.109 (0.078)			0.035 (0.069)			-0.109 (0.077)			0.046 (0.075)		
Voted 1996 Primary	0.013 (0.051)			0.082 (0.053)			0.011 (0.050)			0.067 (0.056)		
Constant	0.542 (0.178)	0.541 (0.172)		0.793 (0.134)	0.957 (0.120)		0.538 (0.176)	0.567 (0.172)		0.766 (0.147)	0.939 (0.118)	
<i>n</i>	527	736		511	714		527	736		511	714	
Adjusted <i>R</i> ²	0.02	0.03		0.03	0.01		0.03	0.02		0.02	0.00	
<i>F</i>	2.664	4.623		5.871	3.041		2.643	4.625		6.540	3.041	
<i>p</i> -value	0.005	0.000		0.000	0.006		0.005	0.000		0.000	0.006	

Note: Clustered standard errors in parentheses. ITT effects estimated with OLS. ATT effects estimated with two-stage least squares, using random assignment as an instrument for campaign contact. Fixed effects for strata not shown in canvassing models.

messages themselves are at work. Nevertheless, it is a plausible alternative explanation, and it will take additional experimentation to nail down the specific mechanisms at work.

DISCUSSION

In recent years over a dozen field experiments have demonstrated that grassroots campaigning can substantially increase turnout. The evidence reported here extends this research by demonstrating that personal forms of campaigning can affect people's vote preferences in addition to their decision to turn out and vote. Both door-to-door canvassing and phone calls caused survey respondents to support the candidate by sizable margins. Previous field experiments have found little evidence that partisan get-out-the-vote messages delivered by commercial phone banks do much to boost turnout (Cardy 2005; Panagopoulos 2006, but see Nickerson 2007). Yet these data suggest that commercial phone banks may be able to influence voting preferences if not the decision to vote. These findings are also in line with field experiments of non-partisan campaigns that show professional phone banks are able to affect voting behavior when they use interactive scripts to make a personal connection with contacts (Green and Gerber 2004; Nickerson 2007).

On the basis of numerous field experiments, Green and Gerber (2004: 39, 78) conclude that door-to-door GOTV canvassing produces one vote for every \$18.67 spent on labor (assuming that canvassers make eight contacts per hour and are paid \$16 an hour) and quality commercial phone banks produce one vote for every \$45 (assuming \$1.50 per completed call), but they are careful to point out that these are conservative estimates that do not account for other campaign goals, such as persuasion. The experiment reported here allows us to do such an accounting. The field canvassing results provide a conservative estimate of the effect that canvassing has on candidate support with the campaign workers attracting one supporter for every 12 attempted contacts (ITT effect = 8.6%, Table 3). If canvassers were paid \$16 an hour and they attempted to contact 20 houses per hour (a more conservative number than Green and Gerber's), the campaign would, on average, attract one supporter for every \$9.60 spent on canvassing. The phone bank generated one supporter for every nine attempts (ITT effect = 11.7%, Table 3). So, assuming \$1.50 per completed call and a 50% contact rate, a commercial phone bank would attract one supporter for every \$6.75. Using the multinomial logit ITT effect estimates in Table 3 (7.6% for field canvassing and 6.9% for phone calls), it is also possible to estimate the cost-effectiveness of mobilizing supporters to vote (note that both of these estimates were accompanied by considerable uncertainty). With the same assumptions about costs, a canvassing operation would mobilize one supporter for every \$10.40 in labor cost and a phone bank would mobilize one supporter for every \$10.50. In short, it appears that personally delivered campaign appeals can be a highly cost-effective way to attract and mobilize supporters in low salience elections.

Going beyond previous research, I also investigate the effects of personally delivered campaign messages on beliefs about the candidate. Campaign professionals expend a great deal of effort on crafting messages aimed at positively influencing how voters perceive their candidate. Thus far, studies of message effects are mostly isolated to research on mass media. Yet like most initial studies of a topic, the experimental findings in this regard raise

more questions than they answer. Although respondents who were personally contacted by the campaign were more likely to have *an* opinion about the candidates' leadership qualities and position on urban planning, there is no clear evidence of either a hypodermic stimulus or priming. It does not appear that the canvassing or phone calls did much to cause individuals to develop the opinion that the campaign's candidate, in particular, supported planned growth or possessed strong leadership abilities, unless they were directly contacted by the candidate herself. Although it is plausible that the candidate was just particularly persuasive, it is also possible that the personal touch of her visit left respondents with generally favorable impressions of her. Consequently, even though the campaign was successful at attracting support for their candidate through personal contact, doing so did not require altering respondents' beliefs about the candidate or the weight accorded to those beliefs in their voting decisions. It will take additional study to establish whether this is a general pattern and, if so, sort out the specific explanations for it.

Indeed, as is typical of almost all scientific endeavors, much room is left for future research. The experiment reported here demonstrates that grassroots campaigns can influence voters' attitudes about the candidate, yet it does not address the mechanisms that underlie these effects. Is it the campaign messages themselves? Is it an interest in the campaign fostered by the messages that leads individuals to learn more about the candidates? Or is it merely an instance in which personal communication fosters trust? In addition to questions about causal mechanisms, it is also important to replicate this study in more salient electoral contexts. At the very least, the experiment demonstrates that campaigns can have a sizable effect on citizens' voting decisions in quiescent elections. Are these effects diminished in more visible elections and, if so, by how much. Answering this last question is important, because, for one, it has implications for the cost-effectiveness of using grassroots campaigns to boost support, and more generally, replication is necessary to establish the generalizability of these findings to other electoral contexts.

APPENDIX: SURVEY QUESTIONNAIRE

Q1. In talking to people about elections, we often find that a lot of people were not able to vote because they were sick, or they just didn't have time. How about you – did you vote in the Democratic primary elections [today/last Tuesday]?

1. Yes, did vote
2. No, didn't vote
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

[If Q1 answer is YES then ASK Q2a]

Q2a. In the Democratic primary election that just passed, the candidates for County Commissioner District 3 were (**ROTATE**) John Eder [EEH der], Judy Woodward, and Deanna Archuleta Loeser [DEE-na Are-Chew-LEH-ta LOH-sir]. Who did you vote for?

1. John Eder
2. Judy Woodward
3. Deanna Archuleta Loeser

4. Other (DO NOT READ)
5. Did not vote in this race (DO NOT READ)
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

[If Q1 answer is NO or Refused or DK then ASK Q2b]

Q2b. In the Democratic primary election that just passed, the candidates for County Commissioner District 3 were (**ROTATE**) John Eder [EEH der], Judy Woodward, and Deanna Archuleta Loeser [DEEena Are-Chew-LEH-ta LOH-sir]. Which candidate for County Commissioner would you have voted for?

1. John Eder
2. Judy Woodward
3. Deanna Archuleta Loeser
4. Other (DO NOT READ)
5. Did not vote in this race (DO NOT READ)
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

Q3. Of the three candidates for County Commissioner District 3, (**ROTATE**) John Eder [EEH der], Judy Woodward, and Deanna Archuleta Loeser [DEEena Are-Chew-LEH-ta LOH-sir], which do you think is the strongest supporter of “planned growth” to limit development on the fringes of Bernalillo [BURN-nal-e-YO] county?

1. John Eder
2. Judy Woodward
3. Deanna Archuleta Loeser
4. All of them or some combination (DO NOT READ)
5. None of them (DO NOT READ)
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

Q4. Would you say that the phrase “strong leader” best describes (**ROTATE**) John Eder [EEH der], Judy Woodward, or Deanna Archuleta Loeser [DEEena Are-Chew-LEH-ta LOH-sir]?

1. John Eder
2. Judy Woodward
3. Deanna Archuleta Loeser
4. All of them or some combination (DO NOT READ)
5. None of them (DO NOT READ)
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

Q5. And finally, my last question is: Looking ahead to the upcoming election in November, how likely are you to volunteer your time, contribute money, or attend rallies on behalf of a political candidate? Would you say you are . . .

1. Very likely
2. Somewhat likely

3. Not too likely
4. Not at all likely
8. Don't Know (DO NOT READ)
9. Refused (DO NOT READ)

REFERENCES

- Adams, William C., and Dennis J. Smith. 1980. "Effects of Telephone Canvassing on Turnout and Preferences: A Field Experiment." *Public Opinion Quarterly* 44: 389–95.
- Angrist, Joshua D., Guido W. Imbens, and Donald B. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (434): 444–55.
- Ansolahehere, Stephen, Roy Behr, and Shanto Iyengar. 1991. "Mass Media and Elections: An Overview." *American Political Quarterly* 19 (1): 109–39.
- Arceneaux, Kevin. 2005. "Using Cluster Randomized Field Experiments to Study Voting Behavior." *Annals of the American Academy of Political and Social Science* 601 (September): 169–79.
- Arceneaux, Kevin, Alan S. Gerber, and Donald P. Green. 2006. "Comparing Experimental and Matching Methods Using a Large-Scale Field Experiment on Voter Mobilization." *Political Analysis* 14 (1): 37–62.
- Bergan, Daniel E., Alan S. Gerber, Donald P. Green, and Costas Panagopoulos. 2005. "Grassroots Mobilization and Voter Turnout in 2004." *Public Opinion Quarterly* 69 (5): 760–77.
- 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.
- Blydenburgh, John C. 1971. "A Controlled Experiment to Measure the Effects of Personal Contact Campaigning." *Midwest Journal of Political Science* 15 (2): 365–81.
- Bochel, J. M., and D. T. Denver. 1971. "Canvassing, Turnout, and Party Support: An Experiment." *British Journal of Political Science* 1 (3): 257–69.
- Bositis, David A., Denise L. Baer, and Roy E. Miller. 1985. "Cognitive Information Levels, Voter Preferences, and Local Partisan Political Activity: A Field Experimental Study on the Effects of Timing and Order of Message Presentation." *Political Behavior* 7 (3): 266–84.
- Bureau of the Census. 1995. *1992 Census of Governments*, Volume 1, Number 2. Washington, DC: Government Printing Office.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. Chicago: University of Chicago Press.
- Cardy, Emily A. 2005. "An Experimental Field Study of the GOTV and Persuasion Effects of Partisan Direct Mail and Phone Calls." *Annals of the American Academy of Political and Social Science*, 601 (September): 28–40.
- Gerber, Alan. 2004. "Does Campaign Spending Work? Field Experiments Provide Evidence and Suggest New Theory." *American Behavioral Scientist* 47 (5): 541–74.
- 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.
- Gerber, Alan S., and Donald P. Green. 2001. "Do Phone Calls Increase Voter Turnout? A Field Experiment." *Public Opinion Quarterly* 65 (1): 75–85.
- Goldstein, Ken, and Travis Ridout. 2004. "Measuring the Effects of Televised Political Advertising in the United States." *Annual Review of Political Science* 7: 205–26.
- Green, Donald P., and Alan S. Gerber. 2004. *Get Out the Vote!* Washington, D.C.: Brookings.
- Green, Donald P., Alan S. Gerber, and David W. Nickerson. 2003. "Getting Out the Vote in Local Elections: Results from Six Door-to-Door Canvassing Experiments." *Journal of Politics* 65 (4): 1083–96.
- Green, Donald, Brad Palmquist, and Eric Schickler. 2002. *Partisan Hearts and Minds*. New Haven: Yale University Press.
- Hillygus, D. Sunshine. 2005. "Campaign Effects and the Dynamics of Turnout Intention in Election 2000." *Journal of Politics*, 67(1) 50–68.

- Huckfeldt, Robert, and John Sprague. 1992. "Political Parties and Electoral Mobilization: Political Structure, Social Structure, and the Party Canvass." *American Political Science Review* 86 (1): 70–86.
- Iyengar, Shanto, and Donald R. Kinder. 1987. *News That Matters: Television and American Opinion*. Chicago: University of Chicago Press.
- Jacobs, Lawrence R., and Robert Y. Shapiro. 1994. "Issues, Candidate Image, and Priming: The Use of Private Polls in Kennedy's 1960 Presidential Campaign." *American Political Science Review* 88 (3): 527–40.
- Kim, Sei-Hill, Dietram A. Scheufele, and James Shanahan. 2005. "Who Cares About the Issues? Issue Voting and the Role of News Media During the 2000 U.S. Presidential Election." *Journal of Communication* 55 (1): 103–21.
- Klapper, Joseph T. 1960. *The Effects of Mass Communication*. New York: Free Press.
- Kramer, Gerald H. 1970. "The Effects of Precinct-Level Canvassing on Voting Behavior." *Public Opinion Quarterly* 34: 560–72.
- 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.
- Michelson, Melissa R. 2003. "Getting Out the Latino Vote: How Door-to-Door Canvassing Influences Voter Turnout in Rural Central California." *Political Behavior* 25 (3): 247–63.
- McGuire, William J. 1968. "Personality and Susceptibility to Social Influence." In *The Handbook of Personality Theory and Research*, eds Edgar F. Borgatta and William W. Lambert. Chicago: Rand McNally.
- McNulty, John E. 2005. "Phone-Based GOTV – What's on the Line? Field Experiments with Varied Partisan Components, 2002–2003." *Annals of the American Academy of Political and Social Science* 601 (September): 41–65.
- Miller, Joanne, and Jon Krosnick. 2000. "News Media Impact on the Ingredients of Presidential Evaluations: Politically Knowledgeable Citizens Are Guided by a Trusted Source." *American Journal of Political Science* 44 (2) 295–309.
- Nickerson, David. 2005. "Partisan Mobilization Using Volunteer Phone Banks and Door Hangers." *Annals of the American Academy of Political and Social Science* 601 (September): 10–27.
- Nickerson, David W. 2006. "Volunteer Phone Calls Can Increase Turnout: Evidence From Eight Field Experiments." *American Politics Research* 34 (3): 271–92.
- Nickerson, David. 2007. "Quality is Job One: Comparing Professional and Volunteer Phone Banks." *American Journal of Political Science* 51 (2): forthcoming.
- Nickerson, David, Ryan Friedrich, and David King. 2006. "Partisan Mobilization Campaigns in the Field: Results from a Statewide Turnout Experiment in Michigan." *Political Research Quarterly* 59 (1): 85–98.
- Panagopoulos, Costas. 2006. "Are Partisan Voters More Responsive to Partisan or Nonpartisan GOTV Messages?" Unpublished manuscript. New Haven, CT: Institution for Social and Policy Studies, Yale University.
- Ramirez, Ricardo. 2005. "Giving Voice to Latino Voters: A Field Experiment on the Effectiveness of a National Nonpartisan Mobilization Effort." *Annals of the American Academy of Political and Social Science* 601 (September): 66–84.
- Rosenstone, Steven J., and John Mark Hansen. 1993. *Mobilization, Participation, and Democracy in America*. New York: Macmillan.
- Stoker, Laura. 1993. "Judging Presidential Character: The Demise of Gary Hart." *Political Behavior* 15 (2): 193–223.
- Vavreck, Lynn, Constantine J. Spilotes, and Linda L. Fowler. 2002. "The Effects of Retail Politics in the New Hampshire Primary." *American Journal of Political Science* 46 (3): 595–610.
- Wielhouwer, Peter W. 2003. "In Search of Lincoln's Perfect List: Targeting in Grassroots Campaigns." *American Politics Research* 31: 632–69.
- Zaller, John. 1992. *The Nature and Origin of Mass Public Opinion*. Cambridge: Cambridge University Press.

