

# **Do Robotic Calls From Credible Sources Influence Voter Turnout or Vote Choice? Evidence From a Randomized Field Experiment**

DARON R. SHAW

*University of Texas at Austin, Austin, Texas, USA*

DONALD P. GREEN

*Columbia University, New York, New York, USA*

JAMES G. GIMPEL

*University of Maryland at College Park, College Park, Maryland, USA*

ALAN S. GERBER *Yale University, New*

*Haven, Connecticut, USA*

*The effectiveness of prerecorded phone calls was assessed in the context of a Texas Republican primary election that featured a contest for state Supreme Court. Automated calls endorsing one of the judicial candidates were recorded by the sitting Republican governor and directed at more than a quarter million people identified as likely voters and probable supporters of the governor. Two experimental designs were used to evaluate the calls' effectiveness. The first design randomly assigned households to treatment and control conditions in order to gauge the calls' effects on individuals' voter turnout, as measured by public records. The second design randomly assigned precincts to treatment and control conditions in order to assess whether the calls increased the precinct-level vote margin of the endorsed candidate. Results suggest that the automated calls had weak and statistically insignificant effects on turnout and vote margins.*

*KEYWORDS field experiment, persuasion, source credibility, voter mobilization, voter turnout*

Address correspondence to Daron R. Shaw, University of Texas at Austin, 1 University Station A1800, Government Department, Austin, TX 78712, USA. E-mail: [dshaw@austin.utexas.edu](mailto:dshaw@austin.utexas.edu)

## INTRODUCTION

Although restricted by law in approximately half of U.S. states, automated phone calls remain an important communication channel for political campaigns. They are touted by those who use them as an important part of a winning campaign. This perspective is widely shared by industry professionals, who counsel that “robocalling is a very powerful and inexpensive medium candidates use to target specific voters” (Simonson 2009). Clearly, if these low-cost “robocalls” have even a modest positive effect on voters, they represent an attractive option for campaigns—especially down ballot, where money is tight and candidates and ballot measures struggle for voters’ attention.

Automated calls thus represent an attractive opportunity for social scientists. Robocalls are uniform “treatments” that can be deployed on a grand scale in order to study the conditions under which campaign messages mobilize or persuade voters. Of particular interest are automated messages that, according to theories of persuasion, seem likely to influence voters: messages that convey sincere endorsements from sources who are perceived to be credible by receivers (Hovland, Janis, and Kelley 1953; Milburn 1991).

The present study engages these practical and theoretical questions by leveraging an unusual research opportunity: we were permitted to conduct an experiment during the closing days of the March 2006 Texas Republican Primary election. This experiment featured the random assignment of an automated call from Republican Governor Rick Perry on behalf of a candidate for the state Supreme Court. In the first analysis, we assess the impact of the call on turnout among households identified as likely Republican supporters. In the second, we assess the impact of the call on votes cast for this Supreme Court candidate at the precinct level. Although the margin in the race was razor-thin, the mobilization and persuasion effects of automated calls appear to be too weak to have determined the outcome.

### Persuasion, Source Credibility, and Automated Messages

Studies of the persuasive effects of election communications date back to the earliest days of survey research (Berelson, Lazarsfeld, and McPhee 1954). Despite decades of investigation, researchers have had difficulty establishing how (and how much) the actions of candidates influence the decisions of voters. It is not so much that citizens fail to respond to the stimulus of the campaign. At the presidential level, conventions, debates, and even major policy initiatives sometimes move voters. It is that voters move—more often than not—in utterly predictable ways given their partisan dispositions and short-term conditions (see Campbell et al. 1960; Zaller 1992; Gelman and King 1993). Campaigns, from this vantage point, merely serve to activate

voters; true persuasion is minimal due to the perceptual biases of partisans and the lack of engagement of independents (Delli Carpini and Keeter 1997).

While the skepticism of many election scholars toward campaign effects is well known, equally well known is the subsequent proliferation of challenges to the “minimal effects” school of thought. Most of these challenges have resulted from improved data and methods, and most have focused on television. For example, experimental research has suggested that television ads both mobilize and persuade voters (Ansolabehere and Iyengar 1995). Other studies have used information on actual television advertising outlays to demonstrate a small but statistically significant effect on turnout (Goldstein and Freedman 2002) and candidate support (Johnston, Hagen, and Jamieson 2004; Shaw 2006). And, most recently, large-scale field experiments have demonstrated that television ads in gubernatorial elections produce a substantial, but short-lived, boost in candidate standing (Gerber et al. 2011).

Research on the persuasive effects of direct forms of campaign communication underscores the challenge of persuading voters and the importance of source credibility for persuasion. In general, experimental studies of the persuasive effects from direct mail (Gerber 2004; Cardy 2005), live phone calls (Adams and Smith 1980), or canvassing (Arceneaux 2005) have found them to be fairly weak. Amid these null findings, the importance of source cues has been discovered almost inadvertently. Nickerson (2007b) found that a progressive political organization canvassing on behalf of state legislative candidates in Oregon had little effect on the vote preferences of Democrats and a negative effect on political independents. Arceneaux and Kolodny (2009) also found a backlash in their experimental study of a liberal group’s canvassing efforts, which apparently persuaded Republicans to vote against the endorsed Democratic candidates.

Perhaps the most relevant literature for our study is that surrounding endorsements. In their analysis of newspaper and celebrity endorsements in the 2008 presidential election, Jamieson and Hardy (2009) contend that endorsements are important because they have heuristic value as predictors of a candidate’s ideology. Such ideological cues can work in two ways. First, they can pull the ideological perception of the endorsed candidate in line with the ideological perception of the endorser. Second, they can push the perception of the unendorsed farther from the perceived ideology of the endorser. Endorsements provide influential cues about the ideological identities of the candidates, perhaps enabling voters with relatively low levels of information to vote as though well informed (Popkin 1994).

Observational studies of endorsements typically corroborate the Jamieson and Hardy argument. Most of this research has focused on the role of endorsements during the nominating process. Cohen et al. (2003) examined “all publicly reported endorsements in a broad range of publications” and concluded that trial heat poll results and endorsements are “almost equally important predictors of primary outcomes.” Similarly, Steger (2007)

found that such polls and endorsements significantly predict votes in both parties during the primaries, while Gimpel (1998) and Rappoport, Stone, and Abramowitz (1991) demonstrated that endorsements by groups and unions drum up enthusiasm among their members.

No field experiments to date, however, have tested the persuasive effects of a credible party leader directing messages to fellow partisans. In contrast to the door-to-door campaigns studied by Nickerson (2007b) and Arceneaux and Kolodny (2009), the messages deployed here are conveyed by a well-known party leader rather than local activists. As we will explain below, the source of this message is particularly credible in this context, having originally appointed the candidate to his office and vouching for his conservative credentials to fellow Republicans (Hass 1981; Sternthal, Phillips, and Dholakia 1978). Perry's message and audience meet Iyengar and Valentino's (2000: 110) criteria for credible, effective communication: "A credible advertisement is one that 'resonates' with voters—that is, is congruent with their preexisting beliefs about the competing candidates' strengths or weaknesses ... [C]ampaign communication is most effective when it plays upon ... voters' prevailing expectations or predispositions."

Although source credibility is important in principle, in practice it may be undermined by voters' indifference to impersonal forms of communication. E-mail outreach appears to have no effect on turnout, whether from nonpartisan groups (Nickerson 2007a) or a political party (Stollwerk 2006), although text message reminders appear to have an effect when directed to people who opt in to this kind of communication (Dale and Strauss 2009). With regard to automated calls in particular, Ramirez (2005) found that nonpartisan messages directed at Latino voters by Latino media personalities and elected officials had no effect, and the same null findings hold for several other studies in which calls were made by non-celebrities (Green and Gerber 2008). One explanation for the failure of automated messages to raise turnout is that such calls are perceived negatively by respondents, who ignore or discount them (see Pew Research Center for the People & the Press 2008). It should be stressed, however, that while the literature suggests that automated calls rarely affect turnout, research has yet to examine whether such calls affect candidate preference. Even if the endorsement does not cause respondents to vote, it may increase the salience and appeal of the endorsed candidate.

### The Texas Automated Call Experiment

Our study takes place within the context of the 2006 Republican Primary contest in Texas. One of the candidates agreed to build an experiment into his automated call program, giving us an unprecedented opportunity to design and execute a statistically powerful and valid study of political phone call effectiveness. We ended up with a two-tiered experiment, giving us

insight into how automated phone calls affect both mobilization and persuasion.

To appreciate the strategic context of the race and to judge the external validity of the project, some background is in order. On March 7, 2006, Don Willett defeated Steve Smith in the Republican primary election for Place #2 on the Texas Supreme Court. Willett tallied 280,356 votes (50.54%) to Smith's 274,302 (49.45%), winning the nomination to the office by a mere 6,054 votes. The race occurred in the midst of an extensive but pedestrian set of GOP election contests. All statewide executive offices were on the ballot, including governor, lieutenant governor, attorney general, comptroller, land commissioner, agriculture commissioner, and railroad commissioner. Of these, the governor, lieutenant governor, and railroad commission elections were contested on the Republican side. U.S. Senator Kay Bailey Hutchison was also on the ballot but faced no opponent. Of the 32 U.S. House districts in Texas, six involved contested Republican primaries (districts 3, 14, 15, 17, 22, and 30). There were also two statewide judicial elections on the GOP ballot: an uncontested Supreme Court race and the contest between Willett and Smith.

Both the turnout numbers and campaign finance figures support the notion that Republican nomination elections were uninspiring. Out of 10,658,363 registered voters in Texas in March of 2006,<sup>1</sup> 655,919 cast votes in the Republican primary for governor and 627,163 voted for Hutchison in the uncontested Senate primary. Only 554,658 (or 5.2%) voted in the Smith-Willett Supreme Court contest. By way of comparison, 570,493 (or 5.6%) voted in the 2002 Republican primary for place 2 on the Supreme Court.

The campaign spending figures tell a similar story. Ultimately, Willett raised \$1,792,523 in contributions for the entire cycle. The best available data indicate that he had raised only \$890,000 for the March primary election. Smith, on the other hand, raised and spent only \$15,000. These figures were hardly out of the mainstream. The biggest spender for the March primary was Attorney General Greg Abbott, who spent about \$3.25 million on his election bid." Other than Abbott, Republican statewide and congressional candidates averaged approximately \$17,500 in primary election spending. In short, one could hardly make the case that either Willett or Smith was "shouting in a crowded cafeteria."

This interpretation, of course, could be erroneous if there was a spirited race on the Democratic side of the aisle. But while Texas has an open primary, there wasn't much on the Democratic ballot to capture the public's attention. There were contests for the gubernatorial and U.S. Senate nominations, but the Democrats were enormous underdogs to Perry and Hutchison, respectively, in the general election and little heat was generated. Just over 500,000 voters participated in the Democratic contest. Beyond these races, only the lieutenant governor and attorney general nominations were

contested. The turnout figures are unsurprising when one examines the spending figures; no statewide or congressional Democratic candidates spent more than \$500,000 for the primary, and the average figure was approximately \$12,500. All in all, then, the 2006 Texas primaries were low-spending, low-salience affairs—and probably offer a decent representation of these sorts of elections writ large.

It was against this backdrop—predictable top-of-the-ticket contests followed by a single close, down-ballot race—that we conducted our experiment. In an effort to beat back what was thought to be a serious challenge, Governor Perry recorded an automated phone call on Willett’s behalf. This call was directed at more than 200,000 households that, according to micro-targeting data purchased by the campaign, were both likely primary voters and strong Perry supporters. The text of the automated message is as follows:

This is Governor Rick Perry. Tuesday is primary day and I want to encourage you to vote for my friend, Don Willett, for Supreme Court. I appointed Don Willett to the Supreme Court because he’s a proven conservative who understands that the role of a judge is to interpret the law, not create it from the bench. Keep Don Willett working for the people of Texas on the Supreme Court by voting in the Republican primary on Tuesday.

The call stressed three themes thought to be persuasive to Republican primary voters: (1) Willett’s conservative credentials, (2) Governor Perry’s long-standing support for him, and (3) Willett’s record as an incumbent.

The calls were conducted between 10 AM and 3 PM on the day before the election.<sup>3</sup> The timing was designed to place a large number of recorded messages on respondents’ answering machines, which would be played once respondents returned from work. Up to three attempts were made for each phone number. In the end, 85% of the targeted voters were contacted in some way, either through an answering machine (54%) or a live answer (31%).<sup>4</sup> These rates are typical for political calling campaigns.

The central question is whether the robotic call succeeded in its effort to increase support for Willett, either by increasing voter turnout or by persuading Smith voters to vote for Willett instead. As suggested earlier, the literature on robotic phone calls has to date considered only its effects on voter turnout, and the experiment reported here builds on that literature by considering the further question of whether robotic calls affect voter preference. The working hypothesis is that calls conveying a message from a prominent and credible source, such as the sitting Republican governor, should have a persuasive effect. Governor Perry, according to polls conducted during the week before the primary election, enjoyed overwhelming approval ratings among Republicans and ran uncontested in the primary.

## RESEARCH DESIGN

In order to assess the effects of the robotic calls on vote outcomes, we conducted two experimental analyses. The first randomly assigned precincts containing between 100 and 700 putatively Republican and pro-Perry households to treatment and control groups. We say “putatively Republican” because the Texas voter file does not record voters’ party preferences. Instead, Republican households were identified based on whether a vote had been cast in a prior Republican primary election, and Perry supporters were identified by a microtargeting firm.<sup>6</sup> The turnout and persuasion hypotheses are assessed by comparing average outcomes in the treatment and control precincts. Turnout can be measured at the level of the individual voter. The 815 treatment precincts contained a total of 147,660 voters, as compared to 782 control precincts, which contained 144,942 voters. Vote choice is observed only at the precinct level, which means that the votes of Perry supporters are intermingled among other Republican primary voters. Random assignment, however, means that the baseline level of support for the judicial candidates among households not targeted for robotic calls is expected to be the same in treatment and control precincts. Thus, the difference in vote margins in treatment and control districts provides an unbiased indication of the persuasive effect of the recorded calls.<sup>7</sup>

The second analysis occurred in precincts containing fewer than 100 or more than 700 Republican households. Within these precincts, households, rather than precincts, were randomly assigned to treatment or control groups. Robotic calls were directed only to phone numbers belonging to individuals in the treatment group. A total of 105,270 voters lived in households whose phone number was assigned to the treatment group. The control group comprised 65,657 voters. Because the probability of being assigned to the treatment group increases with the number of voters with the same phone number, our analysis controls for the number of voters in the household.<sup>8</sup> All analyses below also adjust standard errors for the fact that clusters of individuals were assigned to treatment by household or precinct (Green and Vavreck 2008). As we will see, the results are unaffected by the inclusion of covariates that measure voting in prior elections. As shown in the Appendix, for both experiments, the treatment and control groups are closely balanced with respect to past turnout, with no significant differences across six prior elections.

## RESULTS

The two analyses each provide a reading of the mobilizing effects of robotic calls. The precinct-level analysis, reported in Table 1, suggests that the calls had weak effects. When we take into account the fact that this experiment

**TABLE 1** Effects of Robotic Calls on Voter Turnout Based on Precinct Level Random Assignment

	Control group	Treatment group	Difference
Percentage Voting	34.6%	35.3%	0.68%
Robust Cluster Standard Error			(0.57)
(N)	(144,942)	(147,660)	

contains 292,602 individuals in 1,597 clusters, we obtain an estimated treatment effect of 0.68 percentage points with a standard error of 0.57 percentage points. The estimate falls short of statistical significance at the conventional 5% level. When we further control for individuals' voter turnout in six primary and general elections from 2000 through 2004, the estimated effect is 0.42 percentage points with a standard error of 0.49 percentage points.

A more precise estimate of the treatment effect may be obtained by shifting the analysis to the individual/household level. Table 2 reports that for both one-person and two-person households, the estimated treatment effect is 0.2 percentage points. Taking into account that some individuals are clustered in households with more than one voter, regression analysis indicates that for all 166,278 individuals, the average treatment effect is estimated to be 0.20 percentage points with a standard error of 0.27 percentage points. Again, this estimate is statistically indistinguishable from zero at the 5% significance level. When we further control for individuals' voter turnout in primary and general elections from 2000 through 2004, the estimated effect is 0.22 percentage points with a standard error of 0.22 percentage points.

The turnout results hint that the robotic calls *might* have nudged turnout upward to a *slight* extent. We now turn to the precinct level vote outcomes to see whether this effect is detectable in terms of the precinct-level votes cast for the two candidates. Precinct-level results were obtained for 577

**TABLE 2** Effects of Robotic Calls on Voter Turnout Based on Household Level Random Assignment, by Household Size

	Control group	Treatment group	Difference
One-Voter Households			
Percentage Voting	27.6%	27.8%	0.17%
Standard Error			(0.30)
(N)	(45,767)	(45,761)	
Two-Voter Households			
Percentage Voting	37.9%	38.1%	0.24%
Robust Cluster Standard Error			(0.53)
(N)	(19,233)	(55,517)	

treatment precincts and 583 control precincts. The loss of some observations is due to the fact that certain counties make their precinct-level results difficult to obtain; because attrition is unrelated to treatment assignment, it does not cause bias, although it does increase our standard errors somewhat. Table 3 reports that the total number of votes cast in the Republican primary averaged 214.3 in the treatment group and 211.1 in the control group. The overall turnout rate climbed from 8.44% of registered votes in control precincts to 8.55% in treatment precincts. Looking just at the votes cast for Willett and Smith in the judicial contest of interest, the average vote totals are 169.0 and 165.9 in treatment and control precincts. Although these effects are in the positive direction, none of these turnout differences is statistically distinguishable from zero.

Turning to the key outcome of interest, the difference between the number of votes cast for Willett and for Smith, we find evidence of a small but statistically insignificant effect. In treatment precincts, Willett's vote margin average is 17.6, as compared to 16.0 in control markets. This difference of 1.67 votes has a standard error of 2.54 votes, and so we cannot distinguish it statistically from zero. Nevertheless, it is interesting to consider what the estimates would mean if they were taken at face value. Robotic calls seem to have generated 31 additional judicial votes per precinct and increased Willett's net votes by 1.7 votes. This pattern suggests that only a portion of the votes mobilized by Governor Perry were actually cast for Willett.

Although we find the estimated persuasive effect of robotic calls to be statistically indistinguishable from zero, it is instructive to consider what this estimate implies in terms of the cost efficiency with which votes are generated by this medium. The average precinct in our sample contained 183 target voters. If the average household contains 1.5 voters, the average precinct would receive 122 calls. At 5 cents per call (and ignoring setup fees,

**TABLE 3** Estimated Effects of Robotic Calls on Precinct-Level Outcomes

Outcome measure	Control group	Treatment group	difference	standard error
GOP Primary Votes	211.1	214.3	3.16	(8.59)
Cast				
GOP Turnout (Votes Cast/# Registered)	8.44%	8.54%	0.10%	(0.35)
Votes Cast for Willett or Smith	165.9	169.0	3.16	(6.58)
Willett Votes	90.9	93.3	2.41	(3.62)
Net Willett - Smith Difference	16.0	17.6	1.67	(2.54)
<i>N</i>	583	577		

*Note.* Standard errors in parentheses. None of the differences between treatment and control means approaches statistical significance, as all *t*-ratios are less than 1.0.

recording expenses, and data-management costs) and assuming a completion rate of 84%, robotic calls would cost  $\$0.05 \times 0.84 \times 122 = \$5.12$  per precinct. Thus, to improve Willett's net votes by an average of 1.67, the campaign would spend approximately \$5.12, which amounts to \$3.07 per vote. This level of cost-efficiency seems outstanding by the standards of most campaign tactics (Green and Gerber 2008), but there's a catch: the estimated effect of 1.67 could easily be attributed to chance. Reduce this estimate by one standard error, and the cost per vote skyrockets.

## DISCUSSION

This experiment was designed to speak to three research objectives. First, this project examines both the turnout and vote choice effects of automated calls: no other project that we know of measures these two distinct types of influence. Second, this project analyzes automated phone calls in the context of a major statewide political campaign: the overwhelming majority of comparable studies focus on nonpartisan mobilization without any direct connection to a particular candidate. In short, expanding the scope of effects under scrutiny and considering a more competitive (albeit one occurring within the context of a low-salience primary) election allows us additional leverage over the question of how political phone calls impact voters.

The third component of our project assesses what might be taken to be the "maximally effective" call given the speaker's high source credibility. The spokesperson—a sitting Republican governor with a very positive approval rating among Republicans—was expected to be credible in the eyes of Republican respondents. The message—establishing that Willett is a conservative, was appointed by Perry, and believes in judicial restraint—was expected to be persuasive in a Republican primary contest. The calls were produced by a professional firm, and the audio quality of the call was excellent.

Despite the quality of the treatment and size of our experiments, we were unable to detect effects that were substantively large or statistically distinguishable from zero. Our null findings are subject to two interpretations: Either basic suppositions about source credibility need to be reconsidered (Pornpitakpan 2004) or the communication channel, automated calling, robs the message of its force (Green and Gerber 2008).

There are, of course, limits to what we can claim based on a single set of experiments. The setting for our study is a primary election. Although one might think that Republican partisans would be especially receptive to a Republican governor's message, it may be that such messages would be more effective in a general election context, where a party leader could call on fellow partisans on behalf of a partisan cause. It is also possible that an attack message—as opposed to an advocacy message—might produce greater effects.

The question of how far these results generalize to other settings, messages, and voters is best answered empirically. Given the low cost of conducting an automated calling campaign (in states where such campaigns are permitted), extending this line of research in new directions is fairly straightforward, and the current research provides a template for studying both voter turnout and vote choice as outcomes.

## NOTES

1. These are “non-suspense” registered voters. Including suspense voters yields 12,722,671 registrants. Data are from the Texas Secretary of State.

2. Contribution figures, displayed by time, are available through the [followthemoney.org](http://followthemoney.org) Web site.

3. Registrants who had already voted (Texas has a 2-week early voting period in advance of their primary and general elections) were excluded from the selection universe.

4. A small number (less than 1%) of the control group was contacted inadvertently, due to duplicate names in the campaign’s database. Given the high contact rate and negligible rate of inadvertent contacts, we report only intent-to-treat estimates below. Complier average causal effects could be estimated by dividing the intent-to-treat estimates by 0.84. This adjustment has no effect on the statistical significance of the estimates.

5. Unfortunately, there were no public polls that asked about the Supreme Court primary election, so we do not have a precise baseline estimate of support for either Willett or Smith. Furthermore, this was Willett’s first run for statewide office, so we cannot gauge his underlying support from previous elections. We do know, however, that Smith finished second in a 2004 Republican primary election for state Supreme Court, garnering 259,877 votes and 46.8%. Moreover, members of the Willett campaign have told us that they recall internal polling that showed the 2006 race to be “very close/within margin of error” prior to the experiment.

6. Microtargeting involves conducting an extensive, large  $N$  survey (5,000 respondents is typical) of registered voters and then using the results to build models of turnout and vote choice, which can be used to target an appropriate segment of the broader universe of registered voters. It is, advocates argue, a more cost-effective alternative to statewide voter identification calls.

7. In principle, those who receive the calls could spread the word to their neighbors, in which case the causal effect of the calls could involve those outside the experimental target group. Given other experiments that show the lack of household-to-household spillover in treatment effects, we regard this possibility as unlikely.

8. The probability of being assigned to the treatment group was 50% for one-person households and 75% for two-person households.

## REFERENCES

- Adams, W. C. and D. J. Smith. (1980). Effects of canvassing on turnout and preferences: A field experiment. *Public Opinion Quarterly*, 44(3), 389–395.
- Ansolabehere, S. and S. Iyengar. (1995). *Going negative: How attack advertising shrinks and polarizes the electorate*. New York, NY: The Free Press.
- Arceneaux, K. (2005). Using cluster randomized field experiments to study voting behavior. *Annals of the American Academy of Political and Social Science*, 601, 169–179.
- Arceneaux, K. and R. Kolodny. (2009). Educating the least informed: Group endorsements in a grassroots campaign. *American Journal of Political Science*, 54(A), 755–770.

- Berelson, B., P. Lazarsfeld, and W. McPhee. (1954). *Voting: A study of opinion formation in a presidential campaign*. Chicago, IL: The University of Chicago Press.
- Campbell, A., P. Converse, W. Miller, and D. Stokes. (1960). *The American voter*. New York, NY: John Wiley & Sons.
- Cardy, E. 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(1), 28–40.
- Cohen, M., D. Karol, H. Noel, and J. Zaller. (2003). Polls or pols? the real driving force behind presidential nominations. *Brookings Review*, 3(Summer), 36–39.
- Dale, A. and A. Strauss. (2009). Don't forget to vote: Text message reminders as a mobilization tool. *American Journal of Political Science*, 53(4), 787–804.
- Delli Carpini, M. and S. Keeter. (1997). *What Americans know about politics and why it matters*. New Haven, CT: Yale University Press.
- Gelman, A. and G. King. (1993). Why are American presidential election polls so variable when votes are so predictable? *British Journal of Political Science*, 23, 409–451.
- Gerber, A. S. (2004). Does campaign spending work? *American Behavioral Scientist*, 47(5), 541–574.
- Gerber, A. S., J. G. Gimpel, D. P. Green, and D. R. Shaw. (2011). How large and long-lasting are the persuasive effects of televised campaign ads? Results from a randomized field experiment. *American Political Science Review*, 105(1), 135–150.
- Gimpel, J. G. (1998). Packing heat at the polls: Gun ownership, interest group endorsements, and voting behavior in gubernatorial elections. *Social Sciences Quarterly*, 79(3), 634–648.
- Goldstein, K. and P. Freedman. (2002). Campaign advertising and voter turnout: New evidence for a stimulation effect. *Journal of Politics*, 64(3), 721–140.
- Green, D. P. and A. S. Gerber. (2008). *Get out the vote: How to increase voter turnout* (2nd ed.). Washington, DC: Brookings Institution Press.
- Green, D. P. and L. Vavreck. (2008). Analysis of cluster-randomized experiments: A comparison of alternative estimation approaches. *Political Analysis*, 16(2), 138–152.
- Hass, R. G. (1981). Effects of source characteristics on cognitive response and persuasion. In R. E. Petty, T. M. Ostrom, and T. C. Brock (Eds.), *Cognitive responses in persuasion* (pp. 44–72). Hillsdale, NJ: Lawrence Erlbaum.
- Hovland, C. I., I. Janis, and H. Kelley. (1953). *Communication and persuasion*. New Haven, CT: Yale University Press.
- Iyengar, S. and N. A. Valentino. (2000). Who says what? Source credibility as a mediator of campaign advertising. In A. Lupia, M. D. McCubbins, and S. L. Popkin (Eds.), *Elements of reason: Cognition, choice, and the hounds of rationality* (pp. 108–129). New York, NY: Cambridge University Press.
- Jamieson, K. H. and B. W. Hardy. (2009). Media, endorsements and the 2008 primaries. In S. S. Smith and M.J. Springer (Eds.), *Reforming the presidential nomination process*, Washington, DC: Brookings Press.
- Johnston, R., M. G. Hagen, and K. H. Jamieson. (2004). *The 2000 presidential election and the foundations of party politics*. Cambridge, UK: Cambridge University Press.
- Milburn, M. (1991). *Persuasion and politics: The social psychology of public opinion*. Pacific Grove, CA: Brooks/Cole Publishing.

- Nickerson, D. (2007a). Does email boost turnout? *Quarterly Journal of Political Science*, 2(4), 369–379.
- Nickerson, D. (2007b). *Don't talk to strangers: Experimental evidence of the need for targeting*. Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL.
- Pew Research Center for the People & the Press. (2008). *Robocalls now top type of campaign outreach*. April 3.
- Samuel Popkin. (1994). *The reasoning voter* (2nd ed.). Chicago, IL: University of Chicago Press.
- Pornpitakpan, C. (2004). The persuasiveness of source credibility: A critical review of five decades' evidence. *Journal of Applied Social Psychology*, 34(2), 243–281.
- Ramirez, R. (2005). Giving voice to Latino voters: A field experiment on the effectiveness of a national nonpartisan mobilization effort. *Annals of the American Academy of Social and Political Sciences*, 601, 66–84.
- Rappoport, R. B., W.J. Stone, and A. I. Abramowitz. (1991). Do endorsements matter? Group influence in the 1984 democratic caucuses, *American Political Science Review*, 85, 193–203.
- Shaw, D. R. (2006). *The race to 270*. Chicago, IL: University of Chicago Press.
- Simonson, L. (2009). *Robo calls for campaigns can work*. Retrieved November 29, 2009, from <http://ezinearticles.com/?Robo-Calls-For-Campaigns-Can-Work&id=3137846>
- Steger, W. P. (2007). Who wins nominations and why? An updated forecast of the presidential primary vote. *Political Research Quarterly*, 60, 91–99.
- Sternthal, B. L., L. W. Phillips, and R. Dholakia. (1978). The persuasive effect of source credibility: A situational analysis. *Public Opinion Quarterly*, 42, 285–314.
- Stollwerk, A. F. (2006). *Does e-mail affect voter turnout? An experimental study of the New York City 2005 election*. Unpublished manuscript. New Haven, CT: Yale University.
- Zaller, J. (1992). *The nature and origins of mass opinion*. New York, NY: Cambridge University Press.

#### AUTHOR NOTES

Daron R. Shaw received his BA and PhD from the University of California, Los Angeles, before joining the faculty at the University of Texas at Austin in the fall of 1994. His most recent book is *Unconventional Wisdom* (Oxford University Press), which examines across time survey data to inform the popular conversation about voting and elections in the United States. In 2006, he published *The Race to 270* (University of Chicago Press) which analyzes the effects of TV advertising and candidate visits on the 2000 and 2004 presidential elections. In addition, Shaw has published articles in a wide array of political science journals. Before accepting a position at the University of Texas, Professor Shaw worked as a survey research analyst for Market Opinion Research (1989) and Market Strategies, Inc. (1992). He has also

conducted and analyzed polls and voting results for several campaigns, including presidential campaigns in 1992, 2000, and 2004. Professor Shaw is currently a research fellow at the Hoover Institution, a member of the board of overseers for the National Election Study, a member of the Fox News Decision Team, a member of the advisory board for the Annette Strauss Institute, and a presidential appointee to the National Historical Publications and Records Commission. He has also served as a consultant for the Tomas Rivera Policy Institute and the Texas Poll. Professor Shaw is a distinguished teaching professor at the University of Texas and teaches American government, campaigns and elections, political parties, public opinion and voting behavior, and applied survey research.

Donald P. Green is professor of political science at Columbia University, having moved there in 2011 after 22 years at Yale University. The author of four books and more than one hundred essays, his research interests span a wide array of topics: voting behavior, partisanship, campaign finance, rationality, research methodology, and hate crime. Professor Green's recent books include *Partisan Hearts and Minds: Political Parties and the Social Identities of Voters* (Yale University Press 2002) and *Get Out the Vote: How to Increase Voter Turnout* (Brookings Institution Press 2008).

James G. Gimpel is a professor of government at the University of Maryland, College Park. He has been on the faculty there since January 1992. From 1989 to 1991, he worked on Capitol Hill in the U.S. Senate, then did a brief stint at the Department of Housing and Urban Development in the Office of Policy Development and Research. His recent research and teaching interests include American politics, political geography, voting and elections, state politics, U.S. immigration policy, and public opinion. Professor Gimpel is the editor of *American Politics Research*, a scholarly journal specializing in the empirical study of American political behavior and institutions.

Alan S. Gerber graduated from Yale University (summa cum laude, Phi Beta Kappa) and holds a PhD in economics from MIT. He is professor of political science and director of the Center for the Study of American Politics at Yale University, where he teaches courses on experimental methods, statistics, and American politics. His current research focuses on the application of experimental methods to the study of campaign communications, and he has designed and performed experimental evaluations of many campaigns and fundraising programs, both partisan and nonpartisan in nature. His experimental research has appeared in numerous academic journals including the leading journals in political science: *American Political Science Review*, *American Journal of Political Science*, and *The Journal of Politics*, as well as the *Proceedings of the National Academy of Sciences*. He currently serves as an editor of the *Quarterly Journal of Political Science*. He has received various academic honors and awards, including the Heinz Eulau Award for the best article in the American

Political Science Review (2002), and was recently selected to be a fellow in residence at the Center for Advanced Studies in the Behavioral Sciences (2004-2005).

**APPENDIX** Covariate Balance Between Treatment and Control Groups

Precinct Level Random Assignment		
	Control group	Treatment group
% Voting 2004 General	69.9%	69.4%
% Voting 2002 General	64.9%	65.2%
% Voting 2000 General	66.3%	65.9%
% Voting 2004 Primary	39.5%	40.2%
% Voting 2002 Primary	40.5%	41.0%
% Voting 2000 Primary	53.9%	54.1%
(N)	(144,942)	(147,660)
Household Level Random Assignment		
One-Voter Households	Control group	Treatment group
% Voting 2004 General	62.2%	62.0%
% Voting 2002 General	56.9%	56.6%
% Voting 2000 General	59.3%	58.9%
% Voting 2004 Primary	34.9%	34.4%
% Voting 2002 Primary	34.3%	34.0%
% Voting 2000 Primary	46.9%	45.7%
(N)	(45,767)	(45,761)
Two-Voter Households	Control group	Treatment group
% Voting 2004 General	68.3%	69.3%
% Voting 2002 General	64.4%	65.0%
% Voting 2000 General	64.8%	65.8%
% Voting 2004 Primary	44.3%	45.0%
% Voting 2002 Primary	44.9%	45.2%
% Voting 2000 Primary	56.1%	57.3%
(N)	(19,233)	(55,517)