

Getting Out the Vote in Local Elections: Results from Six Door-to-Door Canvassing Experiments

Donald P. Green
Alan S. Gerber
David W. Nickerson
Yale University

Prior to the November 6, 2001 elections, randomized voter mobilization experiments were conducted in Bridgeport, Columbus, Detroit, Minneapolis, Raleigh, and St. Paul. Names appearing on official lists of registered voters were randomly assigned to treatment and control groups. A few days before Election Day, the treatment group received a face-to-face contact from a coalition of nonpartisan student and community organizations, encouraging them to vote. After the election, voter turnout records were used to compare turnout rates among people assigned to treatment and control groups. Consistent with the recent experimental results reported by Gerber and Green (2000b), the findings here indicate that face-to-face voter mobilization was effective in stimulating voter turnout across a wide spectrum of local elections.

Among the many distinctive attributes of American federalism is the frequency with which elections are held. Due to a profusion of state, municipal, and primary elections, the American voter has more opportunities to cast ballots than citizens of any other country. However, few Americans avail themselves of these abundant opportunities. Voter turnout slumps from presidential election years to even-numbered midterm elections. And in off-years, during which many local and some state elections are held, turnout levels fall even lower (Morlan 1984). Despite the immediate relevance of local issues to voters' lives, the typical U.S. municipal election draws between one-fifth and one-half of the *registered* electorate.

As Harold Gosnell (1927) noted in his path-breaking study of voter mobilization in Chicago during the 1924 and 1925 elections, the quiescence of local elections makes them ideal laboratories for studying methods for increasing voter turnout. Amid limited campaigning and few newsworthy political events, the effects of interventions designed to increase turnout are more readily detected. In addition, low voter turnout rates reduce statistical uncertainty, which is maximal when half of the sample casts ballots. Despite these advantages, local elections tend to attract little attention from students of politics, except insofar

as they involve heated racial politics or other circumstances that make them atypical.

In recent years, the study of electoral turnout has increasingly focused on the subject of voter mobilization. Building on the early works of Gosnell (1927) and Eldersveld (1956), the recent scholarship of Rosenstone and Hansen (1993), Verba, Schlozman, and Brady (1995), and Putnam (2000) has emphasized the responsiveness of voters to their social and political environments. A citizen's level of electoral participation and civic engagement more generally is said to respond to blandishments from family members, political parties, and social networks. By implication, a dearth of mobilization activities may account for the low voter turnout rates typical of local elections.

The present study is patterned after the recent field experimental work of Gerber and Green (2000b). Examining the effects of nonpartisan get-out-the-vote (GOTV) drives on voter turnout in the 1998 midterm elections in New Haven, Gerber and Green found that face-to-face canvassing raised turnout rates from approximately 44% in the control group to 53% among those canvassed. This randomized experiment, which involved tens of thousands of registered voters, provides the clearest indication to date of the effectiveness of face-to-face mobilization, although the authors point out that one must be cautious about drawing generalizations based on a single study.

This essay provides six replications of the Gerber and Green (2000b) experiments, spanning a range of competitive and uncompetitive local elections. During the months leading up to the November 6, 2001 election, we collaborated with a variety of nonpartisan organizations to examine the effectiveness of door-to-door canvassing. Names appearing on lists of registered voters were randomly assigned to treatment and control groups. Treatment groups were visited during the days leading up to the election. Control groups were not contacted. After the election, we obtained voter turnout records from each county and calculated the turnout rates in each control and treatment group. It should be stressed that in contrast to most survey-based analyses of voter mobilization, our study does not rely on voters' self-reported turnout or self-reported contact with GOTV campaigns.

This essay begins with a brief overview of our statistical model and estimation procedure. Next, we describe the experiments conducted at each of the face-to-face canvassing sites. We then analyze the effectiveness of the get-out-the-vote campaign in each site and for the sample as a whole. The results indicate that canvassing significantly increases voter turnout across a range of political and social environments. These mobilization effects are significant, both substantively and statistically, and similar in magnitude to other recent experiments (Gerber and Green 2000b; Michelson 2003).

Research Design and Statistical Model

Unlike observational studies of voter mobilization, which examine the correlation between voting and contact with campaigns, experimental studies ran-

domly manipulate whether voters are approached by campaigns. Experimental control eliminates two problems associated with observational data. First, if campaigns target voters who are especially likely to go to the polls, the observed correlation between contact and voter turnout may be spurious. We might observe a strong correlation even if GOTV campaigns were ineffective. Second, if respondents' recollections of whether they were contacted is vague or distorted, the correlation between self-reported contact and turnout will misrepresent the true causal influence of contact.

The principal complication that arises in experimental studies of voter mobilization is that some citizens assigned to the treatment group cannot be reached. We must therefore distinguish between the intent-to-treat effect and the effects of actual contact. The intent-to-treat effect is simply the observed difference in voter turnout between those assigned to the treatment and control groups. If everyone in the treatment group is actually contacted, the intent-to-treat effect is identical to the actual treatment effect. In practice, however, contact rates are lower than 100% because target voters are often unavailable when canvassers visit their residences.

Consider the linear probability model,

$$Y = a + bX + u, \quad (1)$$

where Y is a dichotomous $\{0,1\}$ variable indicating whether a citizen cast a vote, and $X \in \{0,1\}$ represents whether he or she was actually contacted by a canvassing campaign. The coefficient b is the treatment effect, the boost in turnout caused by contact with the mobilization campaign. Contact is itself a function of whether a person was assigned to the treatment or control condition of the experiment. Let the variable Z , also a dichotomous $\{0,1\}$ variable, represent the random assignment to one of these experimental groups, such that

$$X = cZ + e. \quad (2)$$

To estimate the actual treatment effect (b) given a contact rate (c), we must adjust the intent-to-treat effect (t) as follows:

$$t/c \Rightarrow b. \quad (3)$$

In other words, to estimate the actual treatment effect, we take the intent-to-treat estimate and divide it by the observed contact rate. This estimator is equivalent to performing a two-stage least-squares regression of vote (Y) on actual contact (X) using randomization (Z) as an instrumental variable (Angrist, Imbens, and Rubin 1996; Gerber and Green 2000b). So long as we have information about the rate at which subjects assigned to the treatment group are actually contacted by the canvassers, we can accurately estimate the effects of contact.

A similar approach may be used for nonlinear probability models. One complication in studying local elections is that voter turnout tends to be very low, particularly among certain subgroups such as registered voters who did not vote in a previous national election. OLS may produce inadmissible predictions that

voting will occur with less than zero probability. Rivers and Vuong (1988) propose a two-stage conditional probit estimator that parallels the instrumental variables estimator described above.¹ The probit transformation ensures that predicted vote probabilities are confined to the range between 0 and 1.

Door-to-Door Canvassing Sites

Using official lists of voters gathered immediately after the close of registration, we compiled a database of registered voters' names and addresses. Names of individuals residing at the same address were grouped into households, which were in turn grouped geographically into walk lists. We restricted our attention to households with fewer than five registered voters, and in two sites, Raleigh and St. Paul, the population was restricted to households with at least two voters. One registered voter from each household was selected for study, and these voters were randomly assigned to treatment and control groups. The walk lists given to canvassers contained the names and addresses of people in the treatment group, and they were instructed to approach only these residences.

Although the canvassing sites cannot be construed as a random sample of municipal elections occurring nationwide, our study is strengthened by the fact that the get-out-the-vote campaigns took place in very different political and demographic settings. Some elections were tightly contested; others were devoid of meaningful competition. Some sites have large populations of racial and ethnic minorities; others are predominantly white. Our aims in drawing from such a diverse collection of sites are twofold: to better gauge the average treatment effect of canvassing and to examine whether the treatment effects vary systematically with electoral competitiveness or other characteristics of the sites.

Site 1: Bridgeport, Connecticut. Bridgeport is a racially diverse, low-income urban area that votes overwhelmingly Democratic. The November 6 election featured a local school board election and local city council races. Due to the city's lopsided Democratic majority, all but one of these races were uncompetitive, and the remaining election occurred in a district that was outside the area we canvassed. Turnout, as expected, was low.

ACORN, a community organization championing the interests of low- and moderate-income families, conducted a door-to-door campaign in hopes of generating sufficient support among voters for a "living wage" ordinance (raising the minimum wage to \$11.08 per hour) that had been introduced in the city council earlier in the year. Beginning on October 20th and each weekend thereafter, ACORN volunteers followed walk lists urging every treatment household to vote in the upcoming election.

ACORN did not field many volunteers, but those who participated were experienced and well trained. These volunteers, a diverse group of African Americans

¹ Estimates produced by the Rivers and Vuong method turn out to be almost identical to those obtained using maximum likelihood.

and Latinos, were largely female high-school graduates in their 30s and 40s. Some but not all of the canvassers spoke Spanish.

Site 2: Columbus, Ohio. The neighborhoods canvassed in Columbus were near The Ohio State University campus and tended to be heavily populated by students. Since the Franklin County Clerk only recently began collecting data on birth dates, the age of the voters contacted cannot be determined from voter registration records. Based on the observations of those who conducted the canvassing, it appears that the majority of those canvassed were 25 years old or younger. The only election on the slate was for City Council. Despite the at-large format of the city council election, the races were not competitive, and turnout was expected to be low.

Canvassing occurred during the weekend prior to the election. The canvassers were recruited from the local Public Interest Research Group (PIRG) chapter and volunteers from around campus. Unlike the canvassers in Bridgeport, the Columbus canvassers had little political experience. Only one of the volunteers had ever gone door-to-door for any purpose.

Site 3: Detroit, Michigan. With a closely contested mayoral race, the Detroit elections were among the most interesting in 2001. As in Bridgeport, canvassing was conducted under the auspices of ACORN. However, the crew of canvassers, who were predominantly young, African American, and female, had no previous political experience. After receiving a half-hour training session, the canvassers took to the streets during the weekend prior to Election Day, canvassing all day Saturday, Sunday, and Monday.

Site 4: Minneapolis, Minnesota. Minneapolis elected a wide array of officials in 2001: mayor, city council, school board, and the boards governing city parks, library, and taxation. Turnout was expected to be moderate by local election standards. The neighborhood canvassed was racially mixed and working class. Canvassing was conducted on the Saturday before Election Day by the Twin Cities PIRG chapter. Most canvassers were white college students with no previous experience with door-to-door activity.

Site 5: St. Paul, Minnesota. This election focused solely on the mayor's office, and the race was expected to be very close. As it happened, only 400 votes eventually separated the winner and the loser—a mere 0.6%. Two neighborhoods were canvassed, a poor racially mixed neighborhood and a predominantly white working-class neighborhood.

The local PIRG chapter again organized the canvassing effort with a slightly broader range of volunteers drawn from both colleges and community organizations. The canvassers were inexperienced but received a brief training session before venturing into the field during the Sunday before Election Day.

Site 6: Raleigh, North Carolina. In Raleigh, our canvassing experiment focused on a municipal run-off election. Rather than holding conventional municipal primary elections, Raleigh conducts an all-comers election the first Tuesday of October. In the event that no candidate receives an outright majority, the two top

candidates advance to an Election Day run-off. Both the mayoral and the city council elections featured closely contested races.

Canvassing focused on two neighborhoods, one racially mixed and the other predominantly white. Raleigh was the only site where canvassers were paid an hourly wage (\$10 an hour). Half of the canvassers assembled by the local PIRG group were students (mostly North Carolina State University), and half were members of the community. One-third of the canvassers were African American. Only 20% of the canvassers had any experience in canvassing, and the overwhelming majority responded to a prec canvassing survey by indicating that their principal motive for canvassing was to earn money. On the Saturday before Election Day, the canvassers received twenty minutes of instruction before heading into the field.²

An overview of the six sites is presented in Table 1. Looking solely at the regions within each site that were targeted for canvassing, one sees that the variation across sites is considerable. Data from the 2000 Census indicate that the region canvassed in Detroit is 94% black, whereas Columbus is 82% white. Hispanics account for nearly half of the population in the canvassed regions of Bridgeport. St. Paul has a large Asian population. Home-ownership rates vary widely as well. The large student population in Columbus makes for a low median age and 9% rate of home ownership. St. Paul and Raleigh have higher median ages and home ownership rates of nearly 50%. While not a random sample of cities or elections, the sites in this study span a wide array of urban profiles.

Canvassing Scripts

Although the characteristics of the sites and canvassers varied, they tended to follow similar procedures when going door-to-door. Each canvasser was equipped a clipboard, a map, and a target list of names and addresses. The scripts used in Columbus were broadly representative of the kind of scripts used in other sites:

Hi, how are you? I'm _____ with Ohio Youth Vote. We're reminding people that there's an election this Tuesday. Are you [Name]? [If yes:] Hi! I'd just like to give you this little reminder about voting this Tuesday. [Hand reminder sheet and check "reached" on disposition sheet.] [If no:] Oh, may I please speak with [Name]? And by the way, are you registered to vote? [When Name appears, repeat script with person listed on sheet. Check "reached" if they are there, and record the number of voting age people you spoke with in "other:"] Have a nice day!

²In contrast to the other sites, where canvassing occurred without incident, the canvassing effort in Raleigh encountered problems. Some white residents refused to open their doors to black canvassers. Two black canvassers were accosted by white residents and expelled from the neighborhood. A coincidental canvassing effort by white supremacists seeking to deport Arabs raised residents' general level of hostility to canvassers; indeed, the local police stopped and questioned some of the white canvassers in the PIRG campaign, thinking that they were white supremacists. Whether these events altered the effectiveness of the canvassing effort is a matter of speculation.

TABLE 1
Characteristics of Canvassing Sites, Focusing only on Regions of Each Site that Were Actually Canvassed

	Bridgeport	Columbus	Detroit	Minneapolis	Raleigh	St. Paul
Total City Population	139,529	771,470	951,270	382,618	276,093	287,151
Population in Canvassed Areas	19,115	8,222	17,412	12,177	43,030	17,904
White	37%	82%	4%	36%	70%	54%
Black	28%	5%	94%	26%	22%	20%
Asian	4%	9%	0%	6%	2%	17%
Hispanic	47%	3%	1%	23%	6%	7%
Median Age	26	24	35	26	37	29
Owners	20%	9%	28%	20%	49%	48%
Type of Election	School Board	City Council	Mayoral	Mayoral	Mayoral/City Council	Mayoral
Competitiveness	Low	Low	High	Medium	High	High
Voter Turnout Rate among Subjects in the Control Group	9.9%	8.2%	43.3%	25.0%	29.4%	37.6%
N of Subjects in the Control Group	911	1,322	2,482	1,418	2,975	1,104
N of Subjects in the Treatment Group	895	1,156	2,472	1,409	1,685	1,104

Source for Demographic Profile: 2000 Census.

Canvassers were thus responsible for conveying a brief reminder about the upcoming election, in some cases distributing a flyer, and recording the disposition of each visit.³

Data and Design Issues

The procedures by which subjects were assigned at random to treatment and control groups varied slightly across sites. Subjects in Detroit, Minneapolis, and St. Paul were stratified into walk lists before random assignment, while Bridgeport and Columbus were not, but in each of these sites subjects were assigned the same probability of receiving a treatment. In Raleigh, the proportion of subjects assigned to the treatment condition varied across walk lists.⁴ This across-list variation means that the multivariate analyses that follow include dummy variables for each walk list. These walk-list dummy variables are useful for the other sites as well, since they absorb some intra-site variability in turnout rates.⁵

After the election, we obtained voting histories and registration lists from local registrars. These lists were merged with names in the treatment and control groups in order to calculate voter turnout rates. Since both the registration lists and the voter turnout lists came from the same sources, we experienced no difficulties merging the two lists into a unified database. We also obtained information about whether each citizen participated in the 2000 presidential election. This information enables us to check whether random assignment to treatment and control groups was indeed uncorrelated with past voting behavior. It also provides a useful covariate in a multivariate analysis, as past behavior helps reduce the disturbance variance in models predicting voting in 2001.

Results

Randomization Check

Randomization procedures are designed to create treatment and control groups with equivalent pretreatment vote propensities. In order to check that random

³The treatment thus comprises both a personal appeal and distribution of a leaflet. Other experimental evidence seems to show that leaflets alone have minimal effects on turnout (Gerber and Green 2000a). Not reported here are embedded experiments in which the content of the leaflet was varied randomly, sometimes urging subjects to vote and in other cases presenting them with a voter guide culled from a local newspaper. Varying the content of the flyer had small and statistically insignificant effects.

⁴This variation was introduced to enable us to study the interaction between the treatment given to any single individual and the quantity of treatment given to a particular block. In the end, this study proved too small to detect this interaction with any statistical power.

⁵Random assignment within walk lists give us the luxury of being able to discard walk lists (including both treatment and control subjects) when we suspected that canvassers had falsified their records of whom they contacted. In Raleigh, we discarded one walk list because a canvasser implausibly claimed to have contacted every other house in a rigid sequence. In Detroit, we discarded three lists from one canvasser whose records involved an implausible sequence of contacts and noncontacts. These lists were discarded before we gathered data on voter turnout in 2001.

assignment performed this function, we calculated voter turnout rates for treatment and control groups in the 2000 elections, a year before canvassing occurred. In five of the six sites, pretreatment differences are negligible. In three cases, the treatment group voted at slightly lower rates than the control group, and in two cases, slightly higher. In St. Paul, the treatment group voted at rates that were noticeably higher, with a two-tailed p value of .052. We checked the integrity of the randomization procedures used in St. Paul and found them to be sound. Nevertheless, it will be necessary to analyze the St. Paul results in two ways, one based on a simple comparison of treatment and control and the other using past voting behavior as a covariate.⁶

Intent-to-Treat Effects

The intent-to-treat effects of face-to-face canvassing can be calculated by examining the turnout rates among those assigned to the treatment and control groups, making no allowance for the fact that only some of those assigned to the treatment groups were actually contacted. Column two of Table 2 presents these turnout rates for treatment and control groups in each city. In every site, the treatment group turned out at a higher rate than the control group. For example, in Detroit, where over 40% of registered voters cast ballots, turnout in the treatment group was 2.4 percentage points higher than in the control group. In Bridgeport, where turnout in the control group was an abysmal 9.9%, turnout in the treatment group was 4.0 percentage points higher. The outlier in this set of experiments was Raleigh, where turnout was negligibly higher in the treatment group.

Taking all of the experiments into account, face-to-face canvassing has an intent-to-treat effect of 2.1 percentage points. This estimate is statistically significant at the .01 level using a one-tailed test. These intent-to-treat estimates give some indication of how much get-out-the-vote drives like the ones studied here tend to raise aggregate levels of turnout. As we note below, more intensive GOTV campaigns, which contact much higher proportions of the subjects in the treatment group (e.g., Michelson 2003), can be expected to have much larger intent-to-treat effects.

The Effects of Actual Contact

In order to estimate the mobilizing effect of canvassing among those who are contacted, one must make a statistical adjustment for the fact that many people in the treatment group were never contacted. As shown in Table 2, the limiting factor in these GOTV campaigns is the fact that they contact less than half of

⁶ An additional test of randomization examined the joint significance of age, race, gender, party, and past voting history as predictors of assignment to treatment and control groups. Dummy variables were used to mark missing values in order to avoid deleting observations. Both the null and alternative models included dummy variables for walk lists. As expected, the test statistic is non-significant, $F(12, 18729) = .66, p = .80$.

TABLE 2
Treatment Effects, by Site

	Pre-Treatment Difference in Voting Rates, 2000 Election	Post-Treatment Difference in Voting Rates, 2001 Election	Percentage of Treatment Group Actually Contacted	Effects of Actual Treatment on Voting in 2001
Bridgeport (n = 1,806)	-.3 (2.3)	4.0*** (1.5)	28.1	14.4*** (5.3)
Columbus (n = 2,478)	-.4 (2.0)	1.4 (1.1)	14.3	9.7 (7.9)
Detroit (n = 4,954)	.2 (1.4)	2.4** (1.4)	30.9	7.8** (4.5)
Minneapolis (n = 2,827)	-.4 (1.8)	1.9 (1.6)	18.5	10.1 (8.7)
Raleigh (n = 4,660)	.4 (1.4)	.1 (1.4)	44.6	.2 (3.2)
St. Paul (n = 2,208)	3.2* (1.6)	4.6*** (2.1)	32.1	14.4*** (6.4)
All Sites (n = 18,933)	.4 (.7)	2.1*** (.6)	29.6	7.1*** (2.2)

* $p < .10$, two-tailed test. ** $p < .05$, one-tailed test. *** $p < .01$, one-tailed test. Standard errors in parentheses.

Notes: Differences between treatment and control groups were calculated from OLS regressions of voting in 2000 or 2001 on a dummy variable for experimental treatment, with dummy variables for each walk list as covariates. Contact rates (column 3) were calculated from an OLS regression of actual contact on a dummy variable for experimental treatment, with dummy variables for each walk list as covariates. Actual treatment effects (column 4) were estimated from a 2SLS regression of voting in 2001 on contact, with the experimental treatment as an excluded instrumental variable. Both stages of the 2SLS regression included covariates for each walk list.

their walk lists; indeed, in two of the sites, fewer than one citizen in five was actually contacted. It should be emphasized that for purposes of this calculation, contact is defined quite conservatively to include GOTV conversations with intended subjects or their housemates.⁷ Excluded from the definition of contact are instances where canvassers found no one at home, could not locate the address, discovered that they had the wrong address, or were told to go away before making their GOTV appeal.

The rightmost column of Table 2 reports the actual contact effects. The influence of actual contact in Bridgeport, for example, is estimated to be a 14.4

⁷Of the 8,721 subjects assigned to the treatment group, 18% were contacted directly, and another 11% were contacted indirectly insofar as canvassers spoke with another voting-age member of the household. If one assumes that only direct conversations with canvassers influence turnout, the effects of actual treatment will be larger than what we report in Tables 2 and 3 because the estimated intent-to-treat effects are divided by .18 instead of .29.

percentage-point jump in the probability of voting. Four of the six estimates exceed the estimate of 8.7 percentage points reported by Gerber and Green (2000b), although the standard errors associated with the estimates for Columbus and Minneapolis are quite high. Combining all of the sites (but controlling for walk lists and therefore for inter- and intrasite variation), we find an average treatment effect of 7.1 percentage points. This estimate is statistically significant at the .01 level using a one-tailed test. This estimate also falls within one standard error of the Gerber and Green findings (2000b, 659).

In sum, the experimental results reaffirm the effectiveness of face-to-face canvassing as a means of mobilizing voters. Across a wide range of electoral settings, ranging from the sleepy local election in Bridgeport to the closely contested mayoral race in St. Paul, canvassing had a profound effect on voter participation. This effect turns up in places as different economically and demographically as Columbus and Detroit. Although Raleigh appears to be an outlier *ex post*, we cannot reject the null hypothesis of homogeneous treatment effects across the six sites [$F(5,18737) = 1.06, p = .38$]. Moreover, the findings square with the results of other experiments, such as Gerber and Green's (2000b) study of voter mobilization in New Haven's 1998 elections and Michelson's (2003) study of mobilization prior to a local election in a rural town with a large Latino population.

Multivariate and Nonlinear Models

These findings are underscored by a two-stage probit analysis, which is presented in Table 3. We report two versions of this analysis. The first model includes actual contact as a regressor and intended contact as an instrument. Dummy variables marking each walk list in each site are included as covariates at both stages

TABLE 3
Two-Stage Probit Coefficients, with and without Covariates

	Probit Estimates	Standard Errors
Model Without Covariates		
Canvassing	.211**	.069
Model Including Covariates		
Canvassing	.217**	.076
Voting in 2000	1.571**	.030
1 Registered Voter in the Household	.143*	.053
2 Registered Voters in the Household	.192*	.049
3 Registered Voters in the Household	.018	.053

Note: ** $p < .01$, one-tailed test. * $p < .01$, two-tailed test. Both specifications include dummy variables (not shown) marking each walk list in each site. The dummy variables for the number of registered voters in each household treats four voter households as the base category. Estimation method is 2-stage conditional maximum likelihood, see Rivers and Vuong (1988).

of the regression. In the second model, we also control for a set of covariates that predict voter turnout in 2001: voting in 2000 and dummy variables marking whether a household contained one, two, three, or four registered voters.

The two models provide nearly identical estimates of the effectiveness of canvassing. The two probit coefficients are .211 and .217. These estimates imply that a person who would otherwise have a 50% chance of voting would vote with approximately a 58.5% probability after being canvassed face-to-face. Ordinarily, the inclusion of covariates reduces the standard errors associated with an experimental treatment by reducing the disturbance variance. Here, the standard error increases slightly, reflecting an unexpected correlation between voting in 2000 and the treatment for one of the sites (St. Paul). Nevertheless, the probit coefficients in both specifications suggest that contact with canvassers raises turnout by a statistically significant margin ($p < .01$, one-tailed test).

Conclusion

Building upon previous results, these experimental findings demonstrate that mobilization campaigns have the potential to increase turnout substantially in local elections. Each successful contact with a registered citizen raises that individual's probability of voting by approximately 7 percentage points, which is considerable given the fact that local elections often attract only 25% of the electorate to the polls. This figure, moreover, is a conservative estimate. When calculating the effects of actual treatment, we regarded any conversation with a member of the household as a "contact." Only about half of these conversations occurred directly with the subject in the treatment group; the remainder involved urging a housemate to vote and requesting that this message be passed along to the intended subject. Had we restricted the definition of contact to direct conversations with the subject, the apparent effects of canvassing would have been much greater.

The success with which these door-to-door campaigns mobilized voters is especially impressive given the meager budgets on which these campaigns operated. Our experimental results suggest that 12 successful face-to-face contacts translated into one additional vote. Consider what this finding implies for a large scale GOTV campaign. Suppose one were to hire campaign workers at a rate of \$10 per hour. According to our records for Bridgeport and Columbus, where canvassers traveled in pairs but approached different doors, canvassers contacted eight voters per hour. In Raleigh and St. Paul, the rate was five contacts per hour, but this figure reflects the fact that in these sites canvassers not only traveled in pairs but also went in pairs up to every door. Had the teams of canvassers split up, the contacts per hour would presumably have doubled. If we imagine that the average canvasser makes eight contacts per hour, the cost per vote would be \$15. This figure is quite similar to those reported in previous experimental studies using face-to-face canvassing and notably smaller than comparable cost-per-vote

figures associated with commercial phone banks or direct mail (Gerber and Green 2000b, 2001).⁸

One of the paradoxes of local elections is that individual votes have a greater likelihood of affecting the outcome, yet fewer eligible voters participate. The same logic applies to arguments based on the indirect effects that voters can have on elections by mobilizing their friends and neighbors (Shachar and Nalebuff 1999). With such small numbers of voters casting ballots, mobilization campaigns would seem to be a promising strategy for influencing an election. And yet, the overall level of GOTV activity tends to be low in local elections. In lopsided contests, campaigns have little incentive to do this type of work; in competitive contests, campaigns seem content to focus their energies on persuading voters who regularly vote in local elections. This pattern tends to leave undisturbed the massive age and socioeconomic disparities between voters and nonvoters that have long been the focus of scholarship on local voter turnout (Hamilton 1971; Oliver 1999). The present study suggests that nonpartisan groups, as well as partisan groups that choose to use nonpartisan appeals, have the potential to alter this pattern through face-to-face contact with potential voters. Even in settings where the election outcome seems to be a foregone conclusion, this type of personal contact has a marked effect on voter participation.

Manuscript submitted 16 July 2002

Final manuscript received 9 October 2002

References

- Angrist, Joshua D., Guido W. Imbens, and Donald P. Rubin. 1996. "Identification of Causal Effects Using Instrumental Variables." *Journal of the American Statistical Association* 91 (June): 444–55.
- Eldersveld, Samuel J. 1956. "Experimental Propaganda Techniques and Voting Behavior." *American Political Science Review* 50(1): 154–65.
- Gerber, Alan S., and Donald P. Green. 2000a. "The Effect of a Nonpartisan Get-Out-The-Vote Drive: An Experimental Study of Leafleting." *Journal of Politics* 62(3): 846–57.
- Gerber, Alan S., and Donald P. Green. 2000b. "The Effects of Canvassing, Direct Mail, and Telephone Contact on Voter Turnout: A Field Experiment." *American Political Science Review* 94(3): 653–63.
- Gerber, Alan S., and Donald P. Green. 2001. "Do Phone Calls Increase Turnout?" *Public Opinion Quarterly* 65 (Spring): 75–85.
- Gosnell, Harold F. 1927. *Getting-Out-The-Vote: An Experiment in the Stimulation of Voting*. Chicago: University of Chicago Press.
- Hamilton, Howard D. 1971. "The Municipal Voter: Voting and Nonvoting in City Elections." *American Political Science Review* 65(4): 1135–40.

⁸Naturally, a complete accounting of costs must take into account the fixed costs of sustaining organizations that can recruit and inspire canvassers. The canvassing campaigns studied here were put together with only a few weeks' preparation. Better organization and training could improve the hourly productivity of canvassers.

- 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.
- Morlan, Robert L. 1984. "Municipal vs. National Election Voter Turnout: Europe and the United States." *Political Science Quarterly* 99(3): 457–70.
- Oliver, J. Eric. 1999. "The Effects of Metropolitan Economic Segregation on Local Civic Participation." *American Journal of Political Science* 43(1): 186–212.
- Putnam, Robert C. 2000. *Bowling Alone: The Collapse and Renewal of American Community*. New York: Simon and Schuster.
- Rivers, Douglas, and Quang H. Vuong. 1988. "Limited Information Estimators and Exogeneity Tests for Simultaneous Probit Models." *Journal of Econometrics* 39(3): 347–66.
- Rosenstone, Steven J., and John Mark Hansen. 1993. *Mobilization, Participation, and Democracy in America*. New York: Macmillan Publishing Company.
- Shachar Roni, and Barry Nalebuff. 1999. "Follow the Leader: Theory and Evidence on Political Participation." *American Economic Review* 89(3): 525–47.
- Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge: Harvard University Press.

Donald P. Green is A. Whitney Griswold Professor of political science and Director, Institution for Social and Policy Studies, Yale University, New Haven, CT 06520-8209 (donald.green@yale.edu). Alan S. Gerber is professor of political science, Yale University, New Haven, CT 06520-8209 (alan.gerber@yale.edu). David W. Nickerson is a Ph.D. candidate in political science, Yale University, New Haven, CT 06520-8209 (david.nickerson@yale.edu).