



Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment
Author(s): Alan S. Gerber, Donald P. Green, Ron Shachar
Source: *American Journal of Political Science*, Vol. 47, No. 3 (Jul., 2003), pp. 540-550
Published by: [Midwest Political Science Association](#)
Stable URL: <http://www.jstor.org/stable/3186114>
Accessed: 23/06/2011 22:24

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=mpsaa>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Midwest Political Science Association is collaborating with JSTOR to digitize, preserve and extend access to *American Journal of Political Science*.

<http://www.jstor.org>

Voting May Be Habit-Forming: Evidence from a Randomized Field Experiment

Alan S. Gerber Yale University
Donald P. Green Yale University
Ron Shachar Tel Aviv University

Habit is a frequently mentioned but understudied cause of political action. This article provides the first direct test of the hypothesis that casting a ballot in one election increases one's propensity to go to the polls in the future. A field experiment involving 25,200 registered voters was conducted prior to the November general election of 1998. Subjects were randomly assigned to treatment conditions in which they were urged to vote through direct mail or face-to-face canvassing. Compared to a control group that received no contact, the treatment groups were significantly more likely to vote in 1998. The treatment groups were also significantly more likely to vote in local elections held in November of 1999. After deriving a statistical estimator to isolate the effect of habit, we find that, ceteris paribus, voting in one election substantially increases the likelihood of voting in the future. Indeed, the influence of past voting exceeds the effects of age and education reported in previous studies.

For the better part of a century, political scientists have charted individual and group differences in political participation. Scholars such as Harold Gosnell (1927) observed early on that voter turnout rates differed markedly among groups defined by ethnicity, class, gender, and region. In the 1950s, these aggregate patterns were corroborated by panel surveys showing that individuals with strong party attachments and high levels of education were much more likely to vote in successive elections (Campbell et al. 1960). Pooling cross-sectional surveys over time, analysts since the 1960s have found large and enduring turnout differences among age cohorts (Miller and Shanks 1996). In sum, persistence in voting behavior ranks among the most robust empirical generalizations in political science.

What accounts for this persistence? This question is seldom the subject of theoretical reflection or empirical investigation, despite the importance of temporal linkages for a wide array of social science hypotheses. One answer, implicit in most social-psychological research, is that the psychological impetus to vote endures over

time. Certain people have longstanding feelings of civic obligation, interest in political affairs, and a sense of themselves as voters. These attitudes, or enduring response tendencies, continually express themselves over a series of elections (Campbell et al. 1960; Milbrath 1965; Verba and Nie 1972). A second answer, which grows out of research on contextual effects, holds that certain voters are continually mobilized by campaigns and people in their social environment (Huckfeldt and Sprague 1992; Lake and Huckfeldt 1998). By either of these accounts, voting tendencies persist because people typically make similar choices under similar circumstances.

An alternative explanation holds that the act of voting is self-reinforcing. When people abstain from voting, their subsequent proclivity for voting declines; when they vote, they become more likely to vote again. Voting and abstention, in other words, are habit forming. Attitudes and the environment help explain whether voting habits take root, but one's pattern of behavior itself has an independent effect on subsequent conduct.¹

Alan S. Gerber, Professor, Department of Political Science, Yale University, 77 Prospect Street, New Haven, CT 06520 (alan.gerber@yale.edu). Donald P. Green, Professor, Department of Political Science, Yale University, 77 Prospect Street, New Haven, CT 06520, (donald.green@yale.edu). Ron Shachar, Faculty of Management, Tel Aviv University, Tel Aviv, Israel 69978 (rroonn@post.tau.ac.il).

¹This hypothesis has a long intellectual pedigree. Aristotle, for example, argued that ethical behavior shapes one's ethical sensibilities and subsequent ethical choices: "We become just by doing just acts, temperate by doing temperate acts, brave by doing brave acts" (quoted in Wilson 1995, 108).

American Journal of Political Science, Vol. 47, No. 3, July 2003, Pp. 540–550

©2003 by the Midwest Political Science Association

ISSN 0092-5853

Although complementary, these hypotheses have quite distinct theoretical and methodological implications. If voting is not only a recurrent manifestation of enduring psychological or situational factors but also reflects prior voting, variations in the political environment will produce effects that extend beyond their immediate consequences. A decline in mobilization activity would lead to lower turnout in both the election to come and future elections as well. Thus, statistical models that focus solely on the instantaneous effects of mobilization (e.g., Rosenstone and Hansen 1993; Gerber and Green 2000) will underestimate its long-term effects. Similarly, if voting is self-reinforcing, cross-sectional analyses that explain voting solely by reference to various demographic or social-psychological traits (Wolfinger and Rosenstone 1980) would describe who tends to develop voting habits yet overlook the manner in which these habits perpetuate themselves or change over time. The notion that voting is habit-forming introduces a dynamic component to what has been a largely static understanding of political participation.

The research literature on voting habits remains undeveloped, impeded by the limitations of nonexperimental data. In their analysis of National Election Study survey data, Brody and Sniderman (1977, 347–8) report that past voting behavior predicts current turnout, controlling for a host of individual-level traits, such as age, race, income, education, sex, and psychological involvement in politics. Although this list of control variables is long, this kind of regression analysis leaves open the possibility that the apparent influence of past voting stems from “unobserved heterogeneity” among individuals rather than habit. Unmeasured (or inadequately measured) psychological orientations or socioeconomic conditions that caused past voting behavior might also cause current turnout (Nownes 1992, 210; Plutzer 2002, 44). What is dubbed habit may simply reflect an inability to account for all of the persistent causes of voting.

Concerned that voter turnout in one election may predict turnout in the next simply because factors absent from our model affect turnout in both elections, Green and Shachar (2000) propose an instrumental variables estimator to correct for persistent unmeasured causes of voting. Their analyses of the 1972-4-6 and 1992-4-6 American National Election Study Panel Surveys indicate that voter turnout indeed has a profound effect on participation in subsequent elections. Those voting in a given election were in some instances 50 percentage points more likely to vote in the next election. Although these results bolster Brody and Sniderman’s contention that “Voting is for many a habit” (1977, 349), they ultimately rest on

the untestable statistical assumptions associated with the instrumental variables model. Much the same type of concern applies to Plutzer’s (2002) latent growth curve analysis of turnout using the Youth-Parent Socialization panel study, which also turns up evidence of habit formation. Before one can accept habit as fact, one must corroborate these nonexperimental results with experimental data.

No experiments to date have set out to investigate habit, but evidence of habit formation may be gleaned from two small experiments that track the enduring effects of interviewing respondents prior to an election. Kraut and McConahay (1973) and Yalch (1976) randomly assigned lists of registered voters to treatment and control conditions. Subjects in the treatment condition were contacted as part of an opinion survey; different contacts or none at all occurred in the control condition. Kraut and McConahay (1973, 42) conducted face-to-face interviews with subjects in the treatment condition two weeks prior to a 1970 Democratic primary election in May and found that the treatment group voted at significantly higher rates in May and again the following August. Yalch (1976) conducted personal interviews with respondents in the treatment condition prior to the June 1973 special local election and tabulated turnout rates in that election, a July run-off election, and the March 1974 primary election. The effects of the treatment were powerful not only for the June election, but for the July run-off as well. Turnout in the statewide primary election of 1974, however, saw the treatment group return to rates close to that of the district as a whole.²

The possibility that prior voting might shape the propensity to participate in the current election is bolstered by findings in psychology showing that prior behavior can influence future behavior, even if the subject’s prior decision does not alter the apparent costs and benefits associated with the future decision. The classic study by Freedman and Fraser (1966) finds that subjects who are asked to participate in a short survey are more likely to agree to a subsequent larger request. This result demonstrates the effectiveness of the “foot-in-the-door” technique, whereby participation makes future participation more likely, even when initial compliance does not alter the terms of the subsequent request. More recent work has focused on gauging the strength and generality of this effect (see Beaman et al. 1983 for an early

²Yalch’s tabulation of the district’s turnout rates (1976, 335) makes precise comparisons problematic. Yalch did not track turnout rates among a control group of people who were registered at the start of his experiment.

meta-analysis of 120 experimental groups), and on detailing the psychological mechanisms supporting it.³

Taken together, the psychology findings and the handful of studies that speak to the issue of habit formation in political contexts lend credence to the notion that voting is self-reinforcing. The case is far from settled, however. The most relevant previous studies are the two experimental studies measuring interview effects. Both involve relatively small numbers of subjects, and neither establishes the effects of habit with much statistical precision. Although regression analyses using nonexperimental data produce consonant findings, the inherent uncertainties surrounding model specification call out for an experiment designed to test the habit hypothesis.

This article reports the results of a large-scale field experiment conducted prior to the November general election of 1998. More than 25,000 subjects were randomly assigned to treatment conditions in which they were urged to vote through direct mail or face-to-face canvassing. Using public records, we tracked their voting behavior in the 1998 and 1999 November elections. We derive a statistical estimator that enables us to gauge the effects of voting in 1998 on electoral participation in 1999. This model is estimated and the results shown to be robust under a wide variety of specifications, even those that control for voting behavior prior to 1998. The estimates indicate that casting a ballot in one election profoundly increases one's propensity to vote in the next election. Our findings underscore the potential importance of habit in political participation and other forms of social action.

Detecting Habit Formation in Experimental Data

In the context of electoral participation, the concept of habit implies that if two people whose psychological propensities to vote are identical should happen to make different choices about whether to go to the polls on election day, these behaviors will alter their likelihoods of voting in the next election. In other words, holding preexisting individual and environmental attributes constant, merely going to the polls increases one's chance of returning. As Green and Shachar (2000) point out in their dis-

³Recent examples of this work include Gorassini and Olson (1995), who investigate whether the foot-in-the-door effect is explained by self-perception change, and Cialdini, Trost, and Newsom (1995) who show that the foot-in-the-door technique is more effective on those scoring higher on a measure of preference for consistency.

cussion of "consuetude" effects, the *ceteris paribus* clause is crucial. The claim is not simply that individual differences in voting propensity persist over time. That much is obvious from simple cross-tabulations of voting behavior among respondents in panel studies. Rather, the hypothesis is that the propensity to vote changes when one votes.

To formalize this hypothesis and show how an experimental stimulation of turnout can be used to produce a statistical estimate that isolates habit's effects, one must construct a model that allows for both unobserved heterogeneity among individuals and the potential force of habit. Unobserved heterogeneity can be modeled by supposing that voters can be divided into two groups according to the probability they will vote (in the absence of an experimental stimulus). "High-turnout" voters vote with the probability P_H , and "low-turnout" voters vote with the probability P_L , with $P_H > P_L$.⁴ Assume for simplicity that the proportion of high-turnout voters in the treatment group, which was formed by random assignment, is the same as the proportion of high-turnout voters in the control group. The probability that subjects in the treatment group and control group, respectively, vote in the 1998 elections is denoted by:

$$P_{98}^T = \alpha(P_H + Z) + (1 - \alpha)(P_L + Z) = P_{98}^{BASE} + Z, \quad (1)$$

where α is the proportion of high-turnout voters, and Z is the effect of the experimental treatment on the probability of turnout.⁵

$$P_{98}^C = \alpha P_H + (1 - \alpha) P_L = P_{98}^{BASE}, \quad (2)$$

The probability that voters in the treatment group and control group vote in 1999 equals:

$$P_{99}^T = P_{98}^{BASE} + \delta P_{98}^T + \mu_{99} \quad (3)$$

$$P_{99}^C = P_{98}^{BASE} + \delta P_{98}^C + \mu_{99}, \quad (4)$$

where δ is the "habit" effect (an increase in the probability of voting attributable to having voted in the previous election), and μ captures the change in voting probabilities across elections. Equations (3) and (4) assume that canvassing or sending mail over a year before an election has no direct effect on voting behavior. We have not tested this assumption, though in principle it could be subjected to

⁴The arguments can be extended to any number of voter types.

⁵The assumption that the experiment produces a similar increase in turnout rates for each type of voter can be relaxed without any effect on the analysis. In the two type case, Z is replaced by $\alpha Z_H + (1 - \alpha)Z_L$, where Z_i represents the average percentage increase in turnout produced by the experiment for voters of type i .

experimental verification.⁶ Subtracting the 1999 turnout in the control group from the turnout in the treatment group yields:

$$P_{99}^T - P_{99}^C = \delta(P_{98}^T - P_{98}^C). \quad (5)$$

Rewriting this equation yields an expression for δ , the effect of habit on voting:

$$\frac{P_{99}^T - P_{99}^C}{P_{98}^T - P_{98}^C} = \delta. \quad (6)$$

Notice that, since $(P_{98}^T - P_{98}^C)$ equals Z , the denominator is the 1998 “treatment effect.” The effect of habit is measured by calculating the ratio of the “long-term” increase in treatment group turnout to the “short-term” increase in treatment group turnout. Replacing these true probabilities with probabilities estimated from experimental data gives us an estimate of δ .⁷ If there were no treatment effect, the estimate of the effect of habit would be undefined. The practical implication of this observation is that, if the effects of our 1998 intervention were small, the standard error of the estimates of the habit effect will tend to be very large. Thus, only experimental interventions that significantly increase (or decrease) turnout in a given election can offer insights into voting’s long-term consequences, and weak interventions must be studied with very large samples.

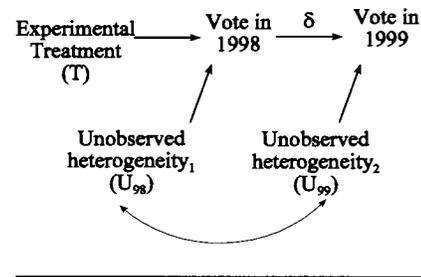
⁶We believe this assumption is justified based on common sense and the observed practices of campaign professionals (who concentrate mobilization efforts on the days and hours prior to voting). The only study that directly addresses the effect of mobilization months in advance of an election is observational research by Niven, who finds significant decay in the effects of early mobilization (Niven 2002). While the Niven study is an important addition to the literature, a skeptic might require some additional evidence. The time frame of Niven’s study does not extend beyond contacts about 6 months prior to the election. Also, the study suffers from limitations inherent in observational research. It is unclear whether the key comparison, that of potential voters canvassed at different points in time with those not canvassed, isolates the effect of canvassing (or the timing of canvassing) from the effect of unmeasured differences across the sample.

⁷Another implication of this derivation is that one cannot infer the effects of habit (δ) through a regression of voting behavior in the 1999 election on voting behavior in the previous election. It can be shown that as long as both types of voters are represented in the population, the effects of habit will be confounded with differences among voters. A bivariate regression in essence compares the turnout rates in 1999 among those who did or did not vote in the previous election:

$$\hat{\delta} = (P_{99}^C | \text{Vote}_{98}) - (P_{99}^C | \text{Abstain}_{98})$$

It can be shown that this estimator gives biased estimates whenever δ lies between 0 and 1, i.e., when there is unobserved heterogeneity across voters. OLS regression may suggest a habit effect even when δ is really zero.

FIGURE 1 Model of Habit and Unobserved Heterogeneity in Voting Behavior



This model and its statistical implications are illustrated in Figure 1. The diagram posits a possible causal relationship (δ) between voting in 1998 and voting in 1999. Unobserved causes of the vote in 1998 (U_{98}) are possibly correlated with unobserved causes of voting in 1999 (U_{99}). The randomized treatment (T) affects voting in 1998 but is uncorrelated with both U_{98} and U_{99} .⁸ Thus, we have all the necessary ingredients for an instrumental variables regression, or what Green and Gerber (2002) term a “downstream” experimental analysis. If the relationship between voting at each point in time is attributable to persistent unobserved factors, rather than habit, an instrumental variables regression should show δ to be near zero. On the other hand, if an instrumental variables regression should produce the same estimate of δ as an ordinary least-squares regression, we would infer that U_{98} and U_{99} are weakly correlated. Our model, in sum, allows the data to describe the relative mixture of habit and unobserved heterogeneity.

Experimental Design

The experiment was conducted in New Haven, CT, a city with a population of approximately 100,000 residents. In September 1998 we obtained a complete list of registered voters. Using this list we created a data set of all households with one or two registered voters, from which we excluded all names with post office box addresses. After these exclusions, we were left with 25,200 individuals whose participation or abstention in the 1998 and 1999 election

⁸The assumption that the treatment is uncorrelated with U_{99} means that interventions prior to the 1998 election have no direct effect on voting in 1999. Another way to state this assumption is to say that the effectiveness of mobilization campaigns decays over time, such that GOTV messages cease to be effective when delivered more than one year prior to an election.

TABLE 1 Assignment to Experimental Conditions

	No Personal Canvassing	Personal Canvassing
Number of Mailings Sent		
None (Control Group)	10,073	2,492
One	3,309	778
Two	3,515	826
Three	3,353	854
Total	20,250	4,950

Cell entries are the number of subjects in each condition.

could be determined from public records. It turns out that 14.2% of the 28,380 subjects whose names appeared on the voter rolls in 1998 had been dropped from the rolls in 1999. However, this type of attrition has no effect on the results reported below because the attrition rates are almost identical for treatment and control groups.⁹

The 1998 and 1999 elections in New Haven were only moderately competitive. In 1998, there were no municipal contests, and the major state and federal elections were not close. The congressional incumbent received over 70% of the vote, and the incumbent governor received nearly twice as many votes as his opponent. In 1999, the city held elections for mayor and Board of Alderman. New Haven is heavily Democratic, and the Democratic incumbent mayor had no major party opponent but did face a long-shot challenge from a candidate he had defeated soundly in the primary election, who ran as an independent. The mayor won an easy general election victory. Three out of four candidates for the Board of Alderman elections faced no major party opponent, and in only 1 of the 30 races was the victory margin within 10%. The level of campaign activity in 1998 and 1999 was modest, and conversations with local party officials and residents revealed that there was no significant door-to-door canvassing effort either year.

⁹Comparing Table 1 above to Table 2 of Gerber and Green (2000) reveals attrition rates of 14.1% among those not assigned to the personal-canvassing condition and 14.6% among those assigned to personal canvassing. Differences in attrition rates across the mail treatment groups are negligible as well. There appears to be no systematic tendency for turnout-enhancing interventions to reduce attrition in our database. We are informed by the city's registrar of voters that nonvoters are not purged from the registration rolls. Instead, a change-of-address postcard is mailed to all registered voters, regardless of their voting participation. If it is returned as undeliverable on successive attempts, those registrants are placed on an inactive list. The experiments did not include voters on the inactive lists.

Our experiments were designed to measure the effect of personal canvassing and direct mail appeals on voter turnout. Through a series of independent random assignments, the sample was divided into control and experiment groups. Table 1 shows the sample sizes for the treatment and control groups for the 2×4 design of personal canvassing and direct mail treatments. The treatment and control groups for the two experiments overlap. The control group consists of 10,073 respondents who were contacted neither by mail or in person. We assigned 4,950 subjects to the personal-canvassing treatment and 12,635 to receive direct mail.¹⁰ Although habit formation manifests itself in subtle changes in voting rates, the large sample sizes used in our experiment enable us to detect these small shifts.

Personal-canvassing procedure. Working in collaboration with the League of Women Voters, we developed a face-to-face canvassing campaign. During each Saturday and Sunday for four weeks prior to the 1998 election, we sent canvassers out to contact randomly selected registered voters. The canvassers were paid \$20 per hour and were primarily graduate students. The experiment city has a substantial minority population, as well as a significant non-English speaking population. More than half of our canvassers were African American or fluent in Spanish, and when possible canvassers were matched to the racial and ethnic composition of the neighborhoods they worked.

For safety reasons, all canvassers worked in pairs, and canvassing ceased at 5:00, when the sun began to set. This procedure constrained both the pool of available canvassing labor and our ability to contact people who were out during the day. Unlike conventional canvassing efforts, ours targeted specific households rather than entire streets, which meant that more time was devoted to locating addresses and walking from one address to the next. Due to these limitations, canvassers were able to contact 1,462 of the 4,950 (29.5%) people in the personal-canvassing treatment group. Examination of the data by voting ward showed a fairly consistent contact rate across the 29 regions of the city.

In order to test the relative effectiveness of different political messages, we divided the treatment group into three subgroups. The messages were designed to reflect both themes used in actual mobilization appeals and social science explanations of voting. Upon contacting one

¹⁰The 1998 experiment also involved an intervention in which a random subset of subjects were called by a commercial phone bank urging them to vote. This treatment did not stimulate voter turnout in 1998, and so we ignore it for purposes of this analysis. Imai's (2002) recent critique of the phone experiment conducted in 1998 thus has no bearing on the results presented here.

of the names on the treatment list, canvassers read the following introduction:

Hi. My name is _____. I'm part of Vote New Haven'98, a nonpartisan group working together with the League of Women Voters to encourage people to vote. I just wanted to remind you that the elections are being held this year on November 3rd.

After appealing either to the citizen's sense of civic duty, neighborhood solidarity, or desire to cast a pivotal vote,¹¹ we sought to estimate the effects of pressing citizens to promise to go to the polls. Ordinarily, canvassers closed their appeal by saying, "We hope you'll come out and vote." For a random subset of the treatment group, the closing statement was augmented with the question, "Can I count on you to vote on November 3rd?" Since this variation in wording proved to be nonsignificant, we combine all of those canvassed into a single treatment group for purposes of analysis.

Direct mail. The direct mail experiment was intended to measure the turnout effect of both the number of mailings a voter received and the message conveyed by the mailings. To gauge the effect of varying the number of mailings, we divided the direct mail treatment group into three subgroups and sent these groups 1, 2, or 3 mailings, respectively. As shown in Table 1, each of these three subgroups contained approximately 4,900 persons. The mailings were sent out at three points in time: 15 days before the election, 13 days before the election, and 8 days before the election. The subgroup that was sent three pieces of mail was included on all three mailing dates; the subgroup receiving two mailings was sent mail on the two mailing dates closest to the election, and the remaining subgroup was sent mail eight days before election day. Within each of these groups, we created three additional

¹¹In the *civic duty* condition, the script went on to say "We want to encourage everyone to do their civic duty and exercise their right to vote. Democracy depends on the participation of our country's citizens." In the *close election* condition, the script instead contended that "Each year some election is decided by only a handful of votes. Who serves in important national, state, and local offices depends on the outcome of the election, and your vote can make a difference on election day." In the *neighborhood solidarity* condition, canvassers stated that "Politicians sometimes ignore a neighborhood's problems if the people in that neighborhood don't vote. When politicians see a lot of people turning out to vote, they know they should pay attention to issues important to people who live around here." As is pointed out in Green and Gerber(2000), turnout did not vary to a statistically significant degree depending on the type of message used. Thus, we omit this aspect of the experiment from the presentation. It should be noted, however, that the results we report do not change if interactions are included for different messages.

subgroups, one set of households for each of the three political messages we tested. (Note that subjects who received a certain appeal through personal canvassing received the same type of appeal by mail.) To avoid sending anyone the same mail piece twice, nine different postcards were required, three for each form of appeal. The postcards were prepared by a professional political consulting firm specializing in political direct mail. All nine treatments were three-color, 8.5 × 11 postcards.

Results

Both personal canvassing and direct mail had statistically significant effects on voter turnout in 1998. Consider, first, the effects of personal canvassing on voter participation in 1998 depicted in Table 2. Turnout among those in the control group was 48.1%, as compared to 51.1% among those assigned to the personal-canvassing treatment. This three percentage-point gap understates the effects of personal canvassing, since only 29.5% of those in the treatment group were actually contacted. When recalculated to take the contact rate into account, the apparent effect of personal canvassing is $(51.1 - 48.1)/.295 = 10.2$ percentage points.

The effects of mail on voter turnout are easier to calculate, since contact rates are not an issue. As one would expect, the effects of mail grew stronger as the number of mailings increased. While one or two postcards had faint effects, a regimen of three mailings appeared to boost turnout by 1.5 percentage points in 1998. A simple linear regression of turnout on the number of mailings rejects the null hypothesis of no effect at $p < .05$, using a one-sided test.

The canvassing and direct mail interventions raised voter turnout in 1998. The question is whether those additional voters were more likely to return to the polls in the mayoral election of 1999. The 1999 election can be characterized as an uneventful reelection of a Democratic incumbent in a city where Democrats hold a large majority of party registrants. The mayor rolled to an easy victory in the Democratic primary, whereupon his primary opponent declared his candidacy as an Independent. Although the mayor was white and his opponent black, this election featured little by way of racial acrimony or mobilization. In our sample of one and two voter households, just 39.4% of those registered went to the polls.

The second column of Table 2 shows that turnout rates in 1999 tended to be higher in the treatment groups. While 39.2% of the control group in the personal-canvassing experiment cast ballots in 1999, 40.3% of the treatment group did so. In the mail experiment, the

TABLE 2 Voter Turnout in 1998 and 1999, by Treatment Prior to the 1998 Election

	Percentage Voting in 1998	Percentage Voting in 1999	Number of Observations
<i>Personal Canvassing Experiment</i>			
Subjects in the control group	48.1%	39.2%	20,250
Subjects in the treatment group	51.1	40.3	4,950
<i>Direct Mail Experiment</i>			
Subjects in the control group	48.5	39.2	12,565
Subjects sent one piece of mail	47.7	38.3	4,087
Subjects sent two pieces of mail	49.0	39.3	4,341
Subjects sent three pieces of mail	50.0	41.1	4,207

TABLE 3 Turnout Rates in 1998 and 1999, Controlling for Turnout in 1996

	Voter Turnout in 1998	Voter Turnout in 1999	N
<i>Abstained or Not on Voter Rolls in 1996</i>			
Control Group	23.5%	19.6%	4,568
Treatment Group: Mail	23.5	19.9	4,658
Treatment Group: Personal Canvassing	27.5	21.4	1,084
Treatment Group: Mail & Personal Canvassing	26.0	20.7	1,135
<i>Voted in 1996</i>			
Control Group	67.6	54.7	5,505
Treatment Group: Mail	69.8	56.1	5,519
Treatment Group: Personal Canvassing	70.7	55.9	1,408
Treatment Group: Mail & Personal Canvassing	71.2	55.9	1,323

control group voted at a rate of 39.2% in 1999, as compared to 41.1% among those who received three postcards a year earlier. On the whole, the contrast between treatment and control was less marked in 1999 than in 1998, but our experimental interventions seem to have left an enduring imprint.

When looking at persistent differences in turnout rates among experimental groups, it is natural to wonder whether some kind of accident of randomization occurred, whereby more longstanding voters happened to fall into certain treatment groups. It is instructive to note that a similar pattern of results emerges when we control for voting in the 1996 presidential election, prior to our experimental intervention.¹² As Table 3 shows, mail and

personal contact stimulated voting in both 1998 and 1999, among those who did not cast ballots in 1996. Those contacted both by mail and in person voted at a rate of 26.0% in the 1998 election, as compared to 23.5% among the control group. In 1999, the corresponding gap was 20.7% versus 19.6%—smaller, but still in the expected direction. This pattern also holds for those who voted in 1996. Regardless of whether subjects were contacted by mail or in person or both, they were more likely to vote in both 1998 and 1999 than comparable subjects in the control group.

To this point, we have limited our inspection of the data to what might be termed a “reduced form” analysis, looking at the relationship between the dependent variable and the experimental treatment groups. Examining the data in this way supports the notion that our treatments had enduring influence, but our specific aim is to examine the effects of voting in 1998 on voting in 1999. This can be done in a preliminary way by lumping all

¹²Also, when we use voting in 1996 as the dependent variable, the treatments are nonsignificant (χ^2 , 6 d.f. = 10.0, $p > .10$), as one would expect based on random assignment.

of the treatment groups together and calculating δ . Those who were either canvassed or sent mail turned out at rates of 49.4% and 39.8% in 1998 and 1999. Those who were not contacted in any way turned out at rates of 47.6% and 38.8%. Using Equation 6, we estimate δ to be $1.0/1.8 = .55$. This estimate means that, *ceteris paribus*, voting in 1998 raised the probability of voting in 1999 by 55 percentage points.

In order to estimate the effects of habit using all of the distinct treatments, we turn to regression analysis. As noted in the previous section, we cannot obtain a reliable estimate of the habit effect by means of an OLS regression of voting in 1999 on voting in 1998, because voting in 1998 is potentially correlated with unmodeled causes of voting in 1999. Therefore, we turn to two-stage least-squares (2SLS) regression, using the experimental treatment groups as the instrumental variables. The first stage model regresses voter turnout in 1998 on the mail and personal-canvassing treatments;¹³ the second-stage model regresses voter turnout in 1999 on the predicted values from the first-stage equation.

Table 4 reports 2SLS regression estimates of the experimental effects. The coefficient of .467 indicates that voting in 1998 raised the probability of voting in 1999 by 46.7 percentage points. Other things being equal, registered voters who did not vote in 1998 had a 16.6% chance of voting in 1999, as compared to 63.3% among those who voted in 1998. By any standard, this is a very large effect. Its magnitude is all the more remarkable given that instrumental variables estimation eliminate the influence of other factors that might have encouraged voting in both elections. This dramatic shift in probabilities appears to be the result of *voting per se*.

This finding is little affected by alterations to the statistical model. We reestimated this regression controlling for voting ward, past voting behavior, age, party registration, and the number of persons in the household. If by luck of the draw, random assignment had generated an unusually participatory treatment group, these controls should reduce the size of our estimate. Moreover, by reducing the amount of unexplained variation in voting, these covariates slightly improve the precision with which we estimate the effects of the experimental interventions. As shown in Table 4, adding these control variables to the model slightly increases the effects of voting in 1998, while lowering its standard error. Given an estimate of

¹³The mail treatment is coded as the number of mailings that a person received. The personal-canvassing treatment was coded as a dummy variable. Vote in 1998 was regressed on the treatment variables, and then vote in 1999 was regressed on the predicted values of vote in 1998, with appropriate corrections to the standard errors.

TABLE 4 Estimated Effects of Voter Turnout in 1998 on Voter Turnout in 1999

	No Covariates	Covariates Included*
<i>Linear Regression (2SLS)</i>		
Estimate of δ	.467	.504
Standard error**	(.202)	(.184)
<i>2-Stage Probit Estimates</i>		
Estimate of δ	1.308	1.477
Standard error**	(.664)	(.756)

*Covariates: whether the voter cast a ballot in the 1996 general election whether the voter abstained in the 1996 general election (those not registered in 1996 constitute the base category), the number of registered voters in the household, the ward in which the subject lives, age, age squared, and whether the subject is registered with one of the two major parties. The first stage equation of the linear model regresses vote in 1998 on these covariates and the two treatment variables (number of mailings and whether the voter was in the personal canvassing group). The second-stage equation regresses vote in 1999 on the predicted values of this equation and the covariates. The conditional two-stage probit model uses vote in 1998 as a predictor, with controls for these covariates and the regression error in the first-stage regression (see Rivers and Vuong 1988, 353).

**Standard errors for the probit equations were obtained by jackknifing.

.504 with a standard error of .184, we easily reject the null hypothesis that voting in 1998 did nothing to stimulate future voting ($p < .05$, one-tailed test).

The same finding turns up when we switch from linear regression to a two-stage conditional probit model (Rivers and Vuong 1988). Like the linear model, the two-stage probit model corrects for endogeneity and provides consistent estimates of the effects of voting in 1998. Unlike the linear probability model, probit ensures that predicted probabilities of voting lie within the permissible range of zero to one. As Table 4 indicates, the probit estimate for the effect of voting in 1998 on voting in 1999 is 1.308. This coefficient implies that if non-voters in 1998 had a 16.6% probability of voting in 1999, voters in 1998 had a 63.2% chance of returning to the polls the following year. These figures are virtually identical to the results presented earlier using a linear model. When covariates are added to the model, the probit coefficient rises to 1.48, and its standard error rises slightly. Again, we easily reject the null hypothesis that voting in 1998 had no effect on voting in 1999 ($p < .05$).

These results suggest that the reason voting behavior is correlated over time is not simply that the background factors that cause people to vote at one point in time reassert their influence during each subsequent election. In addition to the continuities created by

socio-psychological and environmental influences, voting and nonvoting per se appear to create behavioral patterns that persist over time. This conclusion, while supported by the empirical results, requires some caution. Our experimental results measure the effect of voting in one election on participation rates in the next election. The long-term effects, as well as the generality of the experimental findings, are still a matter of speculation. However, the pattern observed here has intriguing testable implications. When patterns of behavior are disrupted—an older, well-educated, partisan skips an election or a younger, poorly educated unaffiliated voter makes her first trip to the polls—subsequent probabilities of voting may be altered dramatically.

Why Do Habits Form?

The statistical evidence presented here and in previous research suggests that voting affects the probability of voting in subsequent elections. Now that this pattern has been observed in a range of experimental and nonexperimental analyses, subsequent studies should begin to investigate various explanations for the apparent persistence in voting behavior. Four hypotheses suggest themselves.

1. The political environment reinforces one's level of political participation. Voters receive much more attention from parties, candidates, and issue activists than do nonvoters. When a registered voter fails to go to the polls, he or she becomes less likely to attract the attention of the campaign, whether through direct mail, phone calls, or canvassing. Voting is self-reinforcing, by this account, because parties and interest groups have an incentive to focus their attention on active voters.
2. Voting alters certain broad psychological orientations known to influence voter turnout, such as feelings of civic obligation, level of partisanship, or interest in politics. This kind of argument is consistent with Finkel's (1985) finding that political participation alters one's sense of political efficacy.
3. Going to the polls alters what Fishbein and Ajzen (1975) call "conative attitudes" toward voting, that is, positive or negative feelings about engaging in the act of voting itself. The registered nonvoter may regard going to the polls with a certain amount of apprehensiveness. (Will I know how to work the voting machine? Will I know which line to stand in?) Like internal efficacy, this orientation concerns one's self-confidence in a political environment, but it does so

with a much higher degree of specificity. Internal efficacy is typically operationalized and measured with items like "politics is too complicated for me to understand," whereas conative attitudes toward voting address the issue of whether the image conjured up by the prospect of voting is attractive or aversive.

4. Civic participation subtly alters the way that citizens look at themselves. Going to the polls confirms and reinforces one's self-image as a civic-minded, politically involved citizen. The more one votes, the more one comes to regard going to the polls as "what people like me do on election day." Conversely, abstention weakens this self-conception and the feelings of obligation that grow out of it. In this respect, abstention desensitizes in much the same way that violations of social norms in general reduce inhibitions about subsequent norm violations (Tyler 1990).

These four hypotheses have quite different empirical implications. The first suggests that the observed persistence in voting behavior is not due to habit in the ordinary sense of the term. Instead of increasing the voter's taste for political participation, voting triggers forces in the environment, which in turn stimulate future voting. Reminiscent of the Alcoholics Anonymous slogan "Bring your body and your mind will follow," the latter three hypotheses each involve the psychological repercussions of voting behavior. Let us briefly take up each in turn, commenting on the ways that they might be tested using experimental and nonexperimental data.

The first hypothesis is the easiest to test. If campaign activity accounts for voting habits, two things must be true. Voting in one election must alter the amount of attention that a citizen receives from campaigns, and this attention must significantly increase the chances of voting in the next election. It is unclear how the former supposition comports with the observed behavior of political campaigns. In jurisdictions where campaigns have easy access to voting records, they sometimes tailor their campaign activity to target "active" voters, although their efforts to increase turnout also concentrate on party registrants who fail to vote regularly. An experimental test would be to survey respondents in a turnout experiment, asking them to describe their campaign contacts prior to the second election, while a nonexperimental test would be to examine the persistence in voting patterns in jurisdictions with varying levels of political communication and mobilization activity. The latter part of this proposition concerns the effectiveness of partisan mobilization. Even the most generous estimates of the effects of face-to-face partisan mobilization (Eldersveld 1956; Miller, Bositis, and Baer 1981) place it well below the fifty percentage-point increase we see in voter turnout as the result of past voting.

More common campaign tactics, such as phone calls and direct mail, seem to have much weaker effects (Gerber and Green 2000). Thus, while campaign attention could in theory account for part of the habit effect, it seems unlikely to account for a substantial portion of it.

Scholars working with survey data have posited a reciprocal causal link between political participation and broad political attitudes. The more one participates, the more likely one is to feel that one's participation is meaningful and important. And the more efficacious one feels, the more likely one is to go to the polls. This proposition is certainly testable in the context of a voter-turnout experiment. Any intervention that raises turnout should also heighten feelings of political efficacy, trust, interest in politics, or other attitudes conducive to voter participation. Thus, a post-election survey should reveal significant differences between treatment and control groups. Again, existing research casts doubt on whether this mechanism is sufficiently powerful to account for the observed degree of behavioral persistence in voting. As Cassell and Luskin (1988) point out, although these political orientations have in some cases changed markedly over time, cross-sectional survey evidence shows them to be fairly limited explanations of voting—too limited to explain either the sizeable trends in turnout since the 1960s or the dramatic shifts we witness in the wake of an experimental intervention.

Less is currently known about conative attitudes concerning voting. Focus group interviews seem to suggest that nonvoters are apprehensive about going to the polls, working the voting machines, and the like (National Association of Secretaries of State 1999). It may be that habits form in large part because, through repetition, people grow comfortable with certain types of activity. Conative attitudes about voting tend not to be assessed in conventional surveys, which generally seek to explain voting by reference to more distal causes, such as partisanship or interest in the campaign. As a result, little is known about how people feel about the act of voting itself or how it changes in the wake of voting. Again, this question could be assessed using post-treatment surveys to gauge the effects of increased participation on conative attitudes.

Finally, habit may be explained as an outgrowth of changed self-perceptions. Those who vote come to think of themselves as voters, while those who abstain shrug off this role and its attendant obligations. One interesting implication of this hypothesis is that the schedule of frequent and often low-salience elections typical of the United States contributes to the breakdown of voting habits by offering many opportunities for abstention. Sleepy local elections, by this logic, are akin to gateway drugs, eroding citizens' sense of themselves as involved participants in elections. This proposition is difficult to test using existing

survey data, which seldom if ever explore the self-images of citizens with respect to voting. Nevertheless, the empirical implications are fairly straightforward: high-turnout elections should increase the proportion of the electorate who describe themselves as "voters," a self-perception that would be expected to persist over time and to predict subsequent voting rates.

Conclusion

This article makes three contributions to the study of political and social behavior. The first is methodological. Throughout the social sciences, habit is adduced to explain behaviors ranging from blood donation to tax compliance to drug use, but rarely has habit been studied in a rigorous fashion. The present study illustrates how randomized field experimentation enables researchers to isolate the causative role of habit. The key insight is that any randomized intervention that produces a change in behavior sets the stage for subsequent investigation of behavioral persistence.

Second, our findings demonstrate the profound influence of current behavior on subsequent behavior, at least in the short term. To put our results in perspective, compare the 47 percentage-point effect of past voting to the effects of leading demographic or social-psychological predictors of voting. For example, in their seminal study of voter turnout in the 1972 election, Wolfinger and Rosenstone (1980, 24, 47) found that a 26 percentage-point gap separated those with some high school education from those with postgraduate education. Similarly, citizens in their 80s voted at rates that are approximately 30 percentage points higher than voters in their 20s, controlling for education and other demographic characteristics (1980, 42, 124). The fact that the estimated effects of habit exceed even these enormous group differences merits further investigation into the magnitude and duration of voting habits.

Finally, our study suggests the importance of taking long-term effects into account when assessing the behavioral consequences of campaigns. To date, studies of voter mobilization have examined only whether contact with campaigns stimulates voter participation in the current election. Our results imply that the long-term effects of campaign contact may be equally important. If the effects of habit decay geometrically over successive elections, a mobilization campaign that stimulates 1000 people to vote in the current election produces an additional 887 votes in elections that follow.¹⁴ Conversely, any change in

¹⁴Using the estimate of .467 in the first column of Table 1, we calculate that the cumulative effects of a raising turnout by one vote to be $.467 + .467^2 + \dots = .467/(1 - .467) \approx .887$.

the political environment that diminishes voter turnout, such as the decline of face-to-face mobilization of voters by parties and candidates since the 1960s (Gerber and Green 2000; Putnam 2000), has cumulative consequences. One thousand votes lost today mean another 887 votes subtracted from future elections.

References

- Beaman, Arthur L., Maureen C. Cole, Marilyn Preston, Bonnel Klentz, and Nancy Mehrkens Steblay. 1983. "Fifteen Years of Foot in the Door Research: A Meta-Analysis." *Personality and Social Psychology Bulletin* 9(June):181–96.
- Brody, Richard A., and Paul M. Sniderman. 1977. "From Life Space to Polling Place: The Relevance of Personal Concerns for Voting Behavior." *British Journal of Political Science* 7(1):337–60.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. New York: John Wiley & Sons.
- Cassell, Carol A., and Robert C. Luskin. 1988. "Simple Explanations of Turnout Decline." *American Political Science Review* 82(4):1321–30.
- Cialdini, Robert B., Melanie R. Trost, and Jason T. Newsom. 1995. "Preference for Consistency: The Development of a Valid Measure and the Discovery of Surprising Behavioral Implications." *Journal of Personality and Social Psychology* 69(2):318–28.
- Eldersveld, Samuel J. 1956. "Experimental Propaganda Techniques and Voting Behavior." *American Political Science Review* 50(March):154–65.
- Finkel, Steven E. 1985. "Reciprocal Effects of Participation and Political Efficacy." *American Journal of Political Science* 29(4):891–913.
- Fishbein, M., and A. Ajzen. 1975. *Belief, Attitude, Intention, and Behavior: An Introduction to Theory and Research*. Reading, MA: Addison-Wesley.
- Freedman, Jonathan L., and Scott C. Fraser. 1966. "Compliance Without Pressure: The Foot-in-the-Door Technique." *Journal of Personality and Social Psychology* 4(2):195–202.
- Gerber, Alan S., and Donald P. Green. 2000. "The Effects of Canvassing, Phone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94(September):653–63.
- Gorassini, Donald R., and James M. Olson. 1995. "Does Self-Perception Explain the Foot-in-the-Door Effect?" *Journal of Personality and Social Psychology* 69(1):91–105.
- Gosnell, Harold F. 1927. *Getting out the Vote: An Experiment in the Stimulation of Voting*. Chicago: University of Chicago Press.
- Green, Donald P., and Alan S. Gerber. 2002. "The Downstream Benefits of Experimentation." *Political Analysis* 10(4):394–402.
- Green, Donald P., and Roni Shachar. 2000. "Habit-formation and Political Behavior: Evidence of Consuetude in Voter Turnout." *British Journal of Political Science* 30(September):561–73.
- Huckfeldt, Robert, and John Sprague. 1992. "Political Parties and Electoral Mobilization: Political Structure, Social Structure, and the Party Canvass." *American Political Science Review* 86(March):70–86.
- Imai, Kosuke. 2002. "The Importance of Modern Statistical Methods for Analyzing Data from Field Experiments: Evaluating Voter Mobilization Strategies." Unpublished manuscript, Harvard University.
- Kraut, Robert E., and John B. McConahay. 1973. "How Being Interviewed Affects Voting: An Experiment." *Public Opinion Quarterly* 37(3):398–406.
- Lake, R.L., and Robert Huckfeldt. 1998. "Social Capital, Social Networks, and Political Participation." *Political Psychology* 19(September):567–84.
- Milbrath, Lester W. 1965. *Political Participation: How and Why Do People Get Involved in Politics?* Chicago: Rand McNally.
- Miller, Roy E., David A. Bositis, and Denise L. Baer. 1981. "Stimulating Voter Turnout in a Primary: Field Experiment with a Precinct Committeeman." *International Political Science Review* 2(2):445–60.
- Miller, Warren E., and J. Merrill Shanks. 1996. *The New American Voter*. Cambridge: Harvard University Press.
- National Association of Secretaries of State. 1999. *New Millennium Project—Phase I: A Nationwide Study of 15–24 Year Old Youth*. Lexington, KY: NASS.
- Niven, David. 2002. "The Mobilization Calendar: The Time-Dependent Effects of Personal Contact on Turnout." *American Politics Research* 30(3):307–22.
- Nownes, Anthony. 1992. "Primaries, General Elections, and Voter Turnout: A Multinomial Logit Model of the Decision to Vote." *American Politics Quarterly* 20(2):205–26.
- Plutzer, Eric. 2002. "Becoming a Habitual Voter: Inertia, Resources, and Growth in Young Adulthood." *American Political Science Review* 96(1):41–56.
- Putnam, Robert C. 2000. *Bowling Alone: The Collapse and Renewal of American Community*. New York: Simon and Schuster.
- Rivers, Douglas, and Q.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.
- Tyler, Tom R. 1990. *Why People Obey the Law*. New Haven: Yale University Press.
- Verba, Sidney, and Norman H. Nie. 1972. *Participation in America*. New York: Harper & Row.
- Wilson, James Q. 1995. *On Character*. Washington, DC: AEI Press.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven: Yale University Press.
- Yalch, Richard F. 1976. "Pre-election Interview Effects on Voter Turnout." *Public Opinion Quarterly* 40(1):331–6.