

Does Campaign Spending Work?

Field Experiments Provide Evidence and Suggest New Theory

ALAN S. GERBER

Yale University

This article reports the results of several field experiments designed to measure campaign effects in partisan contests. The findings suggest incumbent campaigns failed to increase incumbent vote share, whereas the challenger campaign was effective. To understand these and other results, the incumbent's optimal spending strategy was analyzed theoretically. The analysis reveals that if incumbents maximize their probability of victory rather than vote share, campaigns by typical incumbents are expected to produce only minimal improvement in incumbent vote share. The analysis also explains how returns to campaign spending vary with the competitiveness of the election, how incumbent spending can improve the incumbent's probability of victory yet have only minimal effect on incumbent vote share, and why rational spending plans might decrease the sponsor's expected vote. This article demonstrates the wide scope of application for field experiments and provides an example of how experimental findings can serve as a catalyst for generating theories.

Keywords: *field experiment; experiment; campaigns; campaign spending; campaign finance; elections*

In U.S. election campaigns, some candidates have significantly more money to spend than others. How important is this financial advantage? Previous studies have reached conflicting conclusions. Some authors argue that incumbent spending is ineffective but that challenger spending produces large gains (Jacobson, 1978; Abramowitz, 1988). Others claim that challenger and incumbent spending are both effective, and still others maintain that neither incumbent nor challenger spending makes any appreciable difference (Erikson & Palfrey, 2000; Green & Krasno, 1988; Gerber, 1998; Levitt, 1994). Scholars and citizens alike have a large stake in this empirical debate. If campaign spending matters for election outcomes, it is a short step to conclude that campaign finance has a significant influence on who runs, who wins elections, and what the winners do in office.

This article presents new evidence and arguments regarding campaign-spending effects. I report the results of several field experiments designed to measure the effect of partisan campaigns on candidate vote shares. These

experiments, performed with the cooperation of actual candidates, estimate the effect of candidate mailings. The experiments strongly suggest that incumbent campaign spending has a negligible effect on candidate vote margins. In contrast, the experimental study of challenger campaign spending shows the challenger's campaign was highly effective. Ironically, these results follow the basic pattern of spending effects uncovered in one of the earliest observational studies (Jacobson, 1978), a study that has been subject to methodological attack since its findings were first reported 25 years ago.

In his seminal work on campaign-spending effects, Gary Jacobson found that an ordinary least squares (OLS) regression of incumbent vote share on incumbent spending produced a small negative coefficient estimate for the marginal effect of incumbent spending (Jacobson 1978). He concluded that incumbent spending does not improve the incumbent's vote share. Many scholars dismissed the possibility that this regression estimate accurately represented the effect of incumbent spending, and they offered plausible explanations for why OLS "failed."¹ A common critique of Jacobson's findings was that incumbents raise their spending levels in response to strong threats. If the control variables do not fully account for the threat level, candidate spending effects will tend to be biased downward due to a negative correlation between incumbent spending and the regression error.

The experimental results presented here, however, are not vulnerable to this criticism because the level of campaign exposure is varied randomly. The findings are consistent with Jacobson's earlier results, and so the possibility that incumbent spending fails to boost incumbent vote share needs to be taken more seriously. In retrospect, although it is quite likely that the early OLS findings are biased due to spurious correlation, I believe scholars might have been too quick to conclude that the simple regression results were so badly biased that they could be dismissed.

Incumbents appear to take fund raising and campaign spending quite seriously. The experimental results, therefore, suggest a need to reconsider the theoretical relationship between incumbent campaign spending and incumbent vote shares. I propose a simple model to explain why, even if incumbent spending typically fails to increase the incumbent's vote share, campaign spending may still yield an electoral benefit. My argument assumes that candidates select a spending strategy to maximize their probability of victory in the face of uncertainty. I explain that, if incumbents are attempting to maximize victory probability rather than their vote share, the typical incumbent campaign is expected to produce only a minimal boost in the incumbent's vote share. Moreover, it will sometimes be perfectly rational for incumbents to spend money in a manner that on average *decreases* the incumbent's vote share. In sum, the original regression results showing no vote return to incumbent spending may have been somewhat biased, but they may also have been badly misinterpreted.

From a broader methodological perspective, the analysis demonstrates the myriad contributions of field experimentation to knowledge accumulation. I present new empirical findings and theoretical arguments, but the research also has a significant methodological implication. This article illustrates the role experiments can play in advancing knowledge. When observational studies produce results that conflict with widely held beliefs (e.g., Jacobson's [1978] finding that incumbent spending does not boost vote share), scholars are quick to find flaws in the studies' research designs. Many of the alleged defects may turn out to be real on close inspection, but the overall effect of this scholarly practice is to reinforce an excessive intellectual bias against the discovery of unexpected social relationships (or nonrelationships). Experiments offer a powerful check on our natural tendency to reject empirical results that do not accord with common views. When the experimental results clash with accepted opinion, it is a strong warning that our conventional assumptions or interpretations may be flawed.

This article proceeds as follows: Section 1 reviews the literature on campaign-spending effects. I describe the methodological approaches used in a few of the most prominent papers and note that these papers reach radically different conclusions regarding the impact of candidate spending. A review of previous work strongly suggests that, given the difficulty of overcoming important methodological criticisms, additional observational research alone will likely fail to produce consensus. Randomized experiments represent a useful addition to the existing research strategies.

Section 2 reports the designs and results of several field experiments. These experiments, which are among the first of their kind, extend the use of field experimentation significantly beyond non-partisan voter mobilization drives (Gerber & Green, 2000).²

Section 3 describes a simple theoretical model of incumbent spending strategy. I show that the experimental results indicating only small or even negative returns to incumbent spending are entirely consistent with a model where the incumbent chooses a campaign plan that maximizes victory probability. Furthermore, the model can explain the pattern of results found in some previous studies of spending effects and can also generate additional empirical predictions.

Section 4 provides concluding observations on the empirical and theoretical contributions of field experiments. This is obviously a speculative article and more work is required before either the empirical or the theoretical arguments can be considered decisively established. I therefore add important cautions regarding the need for further theoretical study and additional empirical research to replicate the findings and to generalize the research to other political environments and forms of campaign communication.

SECTION 1: THE LITERATURE ON CAMPAIGN-SPENDING EFFECTS

Virtually all of the work on candidate spending effects has been based on nonexperimental evidence. There have been several dozen studies performing regressions of candidate vote totals on candidate spending levels.³ The basic model is

$$Vote_{Inc} = \alpha + \beta_1 f(Spending_{Inc}) + \beta_2 f(Spending_{Chal}) + \beta_3 X + \varepsilon \quad (1)$$

where $Vote_{Inc}$ is the incumbent's share of the two-party vote, $Spending_{Inc}$ is total incumbent campaign spending, $Spending_{Chal}$ is total challenger campaign spending, $f(z)$ stands for an arbitrary function of z , and X represents a set of variables other than campaign spending that are thought to influence candidate vote totals, such as challenger quality or constituency partisanship. The spending variables are entered additively, and the particular function of campaign spending is usually either the spending level itself (i.e., $f(z) = z$) or the natural logarithm of spending.

There has been significant disagreement over how Equation 1 should be estimated. A full discussion of the methodological disputes is beyond the scope of this article; however, the relevant issues will be briefly reviewed. There have been three main approaches. Jacobson and other authors estimate spending effects using ordinary least squares (e.g., Abramowitz, 1988; Jacobson, 1978, 1985, 1990). This technique generates unbiased estimates under the assumption that there is no correlation between the observed spending levels and the regression error. Some scholars question this assumption. They argue that even a careful attempt to include all the variables that influence the election result will inevitably leave some things out, or the exact form of the model will be misspecified. If candidates choose spending levels in response to electoral conditions not adequately reflected in the regression model, OLS will yield biased estimates. This danger has led to two main departures from cross-sectional ordinary least squares estimation.

First, some scholars propose instrumental variables for candidate spending levels and use two-stage least squares (e.g., Ansolabehere & Snyder, 1996; Gerber, 1998; Green & Krasno, 1988). A closely related approach is taken by Erikson and Palfrey (2000), who use a theoretical model to deduce conditions under which candidate spending levels can be treated as exogenous. Erikson and Palfrey argue that for very close elections, the marginal return to spending is insensitive to small changes in election conditions. If the candidate's optimal spending levels are determined by the marginal return to spending, it follows that spending levels for close elections can be treated as basically exogenous. Using this insight, Erikson and Palfrey partition U.S. House elections according to their expected closeness and perform separate OLS regressions for each

subsample. The regression results for the subsample of close elections is interpreted as the true marginal return to spending, because for these elections the estimates are not corrupted by spurious correlations between spending levels and omitted variables.

Second, a panel study approach has been suggested as a way to eliminate the hard-to-observe district-level or candidate-level variables that might be lurking in the error term. This line of attack, pioneered by Steven Levitt (1994), limits the sample to pairs of races in which the same two candidates face each other in a “repeat election.” Because the candidates and the district are the same both times the candidates compete, it is argued that once differences in national conditions across election years and incumbency status are controlled for, any remaining change in vote percentages can be attributed to changes in candidate spending levels. If so, regression of the change in candidate vote totals on change in spending levels yields unbiased estimates (Levitt, 1994).

Table 1 compares the results from work using each of these methodological approaches. The table shows the predicted change in the candidate’s vote share associated with a \$100,000 (1998\$) increase in candidate spending. To permit comparison across the four studies of House elections featured in the table, I use the results from models in which candidate spending appears linearly (with the exception of Erikson and Palfrey [2000], who report the main results using the logarithm of spending only). The Jacobson (1985) model predicts that spending an extra \$100,000 will yield only a slight increase in the incumbent’s share of the vote (about a 0.1% increase), whereas a similar spending increase results in a large boost for the challenger (about a 2.2% increase). Green and Krasno (1988) predict large increases in vote share for both the incumbent and challenger, an approximately 2-percentage-point increase for each. Erikson and Palfrey (2000) find that an extra \$100,000 is worth about 1% for the average challenger and 0.6% for the average incumbent. Levitt (1994) in contrast predicts that increasing spending has only a minimal effect, regardless of whether the candidate is an incumbent or challenger. Spending an additional \$100,000 yields approximately 0.1% for the incumbent and 0.2% for the challenger.⁴

Table 1 illustrates the dramatic differences in estimated spending effects. The coefficient values vary widely across studies, but these differences remain rather abstract; it is very hard to say whether a 2.2% response seems “too large” or a 0.1% response seems “too small.” Table 2 converts the values in Table 1 into the cost per additional vote implied by the Table 1 estimates. The dollar figures listed in Table 2 are the cost of changing the vote margin by one vote. If a campaign activity causes a supporter who would otherwise have stayed home on election day to vote, this changes the vote margin by one vote. If a campaign activity causes a voter to switch candidates, this would change the vote margin by two votes.⁵

Table 2 shows that, depending on which model you believe, the cost of changing the vote margin ranges from moderate to truly exorbitant. Levitt’s

TABLE 1: Effect of \$100,000 Spending Increase on Candidate Vote Share (House Elections)

	<i>Incumbent</i>	<i>Challenger</i>
Jacobson (1978, 1985), OLS, linear model	+0.14	+2.17
Green & Krasno (1988), TSLS, linear model	+1.8	+1.96
Levitt (1994), panel, linear model	+0.07	+0.24
Erikson & Palfrey (2000), Table 1, Column 1	+0.57	+1.07

NOTE: 1998 dollars. The table reports the marginal effect of increasing spending by \$100,000. This amounts to an increase in spending of \$0.21 for each member of the voting age population (using 460,000 as the typical voting age population per district). The Erikson and Palfrey (2000) study did not report the effects of spending using a linear model, so the calculations require specification of a base level of spending. Base spending levels were mean incumbent and challenger spending in 1998, as reported in *Vital Statistics on Congress 2001-2002*. TSLS = two-stage least squares.

TABLE 2: Approximate Cost of Adding One Vote to Candidate Vote Margin

	<i>Incumbent</i>	<i>Challenger</i>
Jacobson (1985)	\$188/vote	\$12/vote
Green & Krasno (1988)	\$15/vote	\$13/vote
Levitt (1994)	\$367/vote	\$110/vote
Erikson & Palfrey (2000)	\$46/vote	\$24/vote

NOTE: 1998 dollars. Calculations are based on 190,000 votes cast in a typical House district. For House elections, this implies that a 1% boost in the incumbent's share of the vote increases the incumbent's vote margin by 3,800 votes.

(1994) "minimal effects" panel study regressions imply that raising the incumbent's vote margin by a single vote costs approximately \$350. Improving the challenger's vote margin by one vote costs \$110. Jacobson's (1985) regressions also indicate it is expensive to increase the incumbent's vote margin, with a cost per additional vote of approximately \$200. However, the OLS model suggests a much lower vote cost for challengers. The vote cost for challengers implied by the OLS model is similar to the two-stage least squares estimates for both incumbents and challengers; the cost of moving the vote margin by one vote is in the \$15 to \$20 range.⁶ Erikson and Palfrey's (2000) results imply that both challenger and incumbent spending is effective, though the cost of a marginal vote is twice as high for the incumbent. An additional vote costs the average incumbent \$46 and the average challenger \$24.

Tables 1 and 2 show that estimates generated by the four different methodological approaches vary widely. If one of the methodologies could be defini-

tively determined to be “correct”, then the estimates produced using that method would trump the alternatives. Unfortunately, each of the estimation approaches featured in Tables 1 and 2 is subject to critique. The OLS estimation relies on the questionable assumption that spending levels are uncorrelated with omitted variables. Two-stage least squares relies on assumptions about the validity of the instruments, and bad instruments can sometimes do more harm than good, producing both bias and large standard errors. Erikson and Palfrey (2000) marshal theoretical arguments for why the estimates produced by the subsample of close elections should be interpreted as the true effects of candidate spending. This conclusion rests on several modeling assumptions, and one especially vulnerable assumption is that a candidate’s ability to fund raise does not depend on the closeness of the election. Levitt’s panel approach focuses on a nonrandom subset of election contests and so there is a risk these elections are atypical. It is also possible that restricting the sample to repeat elections reduces, but does not eliminate, the biases due to omitted variables. If history is any guide, debates about modeling assumptions will not be resolved quickly or conclusively.

The experimental method avoids some of the important difficulties typically faced by other types of empirical work. Field experimentation in particular enjoys several advantages over alternative methodologies. Improving on the regression analysis that dominates the literature on campaign spending effects, field experiments use randomization to generate exogenous variation in the independent variable. Experiments thereby avoid bias due to potential spurious correlation with omitted or mismeasured variables. Since randomization is used to sort the subjects into treatment and control groups, the randomization can serve as a valid instrumental variable. There is no need to rely on debatable assumptions about the validity of instrumental variables. If field experiments can be conducted in cooperation with actual campaigns, the questions of external validity will be minimized.⁷

Despite the fact that field experiments can, in theory, be highly informative, there have been only a handful of such studies. Several field experiments have attempted to measure the effect of nonpartisan and partisan political activity on voter behavior.⁸ The most prominent of these are many decades old. Gosnell (1927) performed a large-scale field experiment that measured the effect of nonpartisan mail on turnout in Chicago. Eldersveld (Eldersveld, 1956; Eldersveld & Dodge, 1954) conducted an experimental campaign on behalf of a local charter reform, and later analyzed the effect of a drive to mobilize apathetic voters in an Ann Arbor municipal election.⁹

There have been only two randomized field experiments that assess the effects of candidate campaign activity on their vote share, neither of which measures the effect of a significant campaign effort in a partisan contest. Miller and Robyn (1975) measured the effect of one campaign mailing (sponsored by the eventual winner) on voter turnout and candidate preference in a barely

contested Democratic congressional primary. Voter turnout and candidate preference were measured by a postelection survey. There were approximately 50 respondents in the treatment and control groups for the candidate preference portion of the analysis. Despite the fact that the study found that both turnout and the sponsor's reported vote share were about 10% higher in the treatment group, the authors conclude the mailing had "no effect" since the change in behavior was far short of standard significance levels. Adams and Smith (1980) investigated the effect of brief persuasion calls on turnout and candidate choice in a special election for a Washington D.C. city council seat.¹⁰ The study reported that the telephone call generated a statistically significant boost in turnout but had no effect on candidate preference.¹¹

SECTION 2: FIELD EXPERIMENTS MEASURING CAMPAIGN-SPENDING EFFECTS

Section 1 demonstrates the unfortunate diversity of results produced by observational studies of campaign-spending effects and describes the difficult methodological challenges limiting the prospects for rapid advance using observational methods. Fortunately, randomized experiments provide a powerful means by which to discern the causative role of campaigns. Thanks to five campaigns that were willing to randomize their direct mailing lists, experiments could be performed to measure the effect of campaign mailings on candidate vote totals. The experiments were conducted during a Connecticut mayoral election, an election for the New Jersey state assembly, an election to the Connecticut state legislature, a Congressional primary election, and a Congressional general election. These campaigns were broadly typical of those that rely on direct mail. None of the state and local candidates had the resources to conduct expensive television campaigns. The incumbent member of Congress, who represents a district in an expensive media market, also preferred to rely heavily on direct mail. Due to the difficulty of orchestrating these experiments, they occur where the opportunity presents itself rather than through a process of sampling randomly over geography and types of elections. These studies therefore do not represent a random sample of all election contests and should be viewed as a first step in the accumulation of experimental case studies.

There were two basic designs. For the three state and local elections, randomization was done at the household level, and the effect of the mailings on voter preferences was measured through a postelection survey. For the congressional primary and general election, randomization was performed at the ward level, and campaign effects were measured by comparing vote totals in the treatment and control wards.¹²

2. 1. HOUSEHOLD-LEVEL RANDOMIZATION

2. 1. 1. Description of the Cases and Experimental Design

New Jersey assembly race. An experiment assessed the effect of campaign mailings in a New Jersey assembly race. The Democratic incumbents (there were two incumbents because New Jersey has double-member assembly districts) were expected to win by a comfortable margin.¹³ The campaign by the incumbents (they conducted a joint campaign) was a series of 6 mailings. The campaign concentrated on Democratic voters, especially those with a record of regular voting. Mirroring the structure of the campaign's targeting strategy, there were three distinct samples of households. Sample 1 consisted of 6,354 households containing Democratic voters with a record of high turnout rates. Sample 2 consisted of 10,200 households with additional Democrats and those Independent voters who voted in some but not most recent elections. Sample 3 consisted of 3,000 district households that did not fall into either Sample 1 or 2. This sample included Republicans and low-turnout Independents. The campaign plan did not originally call for sending any mailings to Sample 3, but these households were added to mailing lists to maintain the original total number of campaign mailings after the creation of Sample 1 and Sample 2 control groups.

The nature of the direct-mail campaign varied across the three samples. Among high turnout Democrat households in Sample 1, 5,654 received the first two and the last two mailings, and 700 households were randomly excluded from the mailings. For Sample 2 (Other Democrats, Independents) 9,200 households got all six pieces of mail, while 1,000 households were randomly excluded from the mailings. For Sample 3, the 3,000 households were divided into three groups: 1,000 households received six mailings, 700 got the first two and the last two mailings, and 1,300 households received no mailings.

The Democratic incumbents ran a positive campaign, touting their legislative activities and their support of broadly popular causes, such as efforts to improve education and children's health. The challengers were not mentioned by name in any of the mailings, even for purposes of comparison. Overall, the text of the incumbents' campaign mailings was positive, with criticism reserved only for very unpopular, impersonal targets such as the HMOs and "big insurance companies." Less detailed information is available about the challenger's campaign. Democratic sources said that the challengers, who were soundly defeated, did not appear to be very active. Unfortunately, I do not know how much, if any, background mail was sent to households by other campaigns.

Voter preferences were measured using a postelection survey. Phone numbers were obtained from the voter lists and supplemented with numbers from Survey Sampling Inc. The experimental design created seven treatment and control groups. For the five smallest treatment and control

groups, all households with phone numbers were provided to the survey firm. For the two largest groups, a random sample of households with phone numbers were provided. The survey firm was given a list of 5,806 households and completed 1,163 interviews.¹⁴

Connecticut House Race. A second set of experiments measured the effect of campaign mailings by the Democratic incumbent in a Connecticut House election. The election was expected to be an easy win for the incumbent.¹⁵ The incumbent's mailing strategy was to target Democrats and Independents. The campaign consisted of three mailings. The mailings were sent to two types of households. Sample 1 consisted of households with at least one registered Democrat or Unaffiliated voter. Sample 2 consisted of households with only registered Republicans. As in the New Jersey case, the voters from the challenger's party were not part of the original campaign plan but were added to replace households excluded due to the creation of a Sample 1 control group. To measure the effect of the mailings, the campaign randomly excluded 694 households from the pool of 3,011 Democrat and Unaffiliated households and replaced these with a random subset of 141 Republican households (from a pool of 411 available households).¹⁶

The incumbent's campaign mailings were positive, and there was virtually no mention of the incumbent's partisan affiliation. None of the pieces named or even referred to the incumbent's opponent. Two of the mailings emphasized the candidate's efforts to improve education and health care. The third mailing presented the candidate as a fiscally responsible legislator who "knows that you work hard for your money . . . and believes that you should keep more of it." The incumbent won the election easily. Sources within the incumbent's campaign described the opposing campaign as minimal. After the election, a survey was conducted to determine voter preferences. The survey firm was provided with telephone numbers for 3,422 households and completed 687 interviews.¹⁷

Connecticut mayoral election. A third experiment assessed the effect of the challenger's campaign in a mayoral race. The mayoral election was expected to be reasonably competitive, though the Republican incumbent was popular and the challenger was the clear underdog.¹⁸ The challenger's campaign consisted of a nine-piece mail program. The campaign sent mailings to a target list that was constructed using party registration and turnout history. The campaign plan was to mail to all Democrats and those unaffiliated voters who were likely to vote. The campaign also targeted newly registered voters, because the newly registered were thought less likely to be habitual supporters of the incumbent. The sample universe for the experiment consisted of households that satisfied at least one of the following criteria: (a) There was a registered Democrat at the address; (b) there was a registered "unaffiliated" at the address who voted in either the 1998 general election, the 1997 municipal elections, or the 1995 municipal election; or (c) the household contained a voter who registered after November

1998. The campaign selected 9,900 households that satisfied at least one of these criteria. Of these households, 1,100 were randomly assigned to a control group that received none of the mailings, leaving a treatment group of 8,800 households that received nine pieces of campaign mail.

The incumbent mayor ran a positive campaign based on his accomplishments in office. By the standards of the community, the challenger appeared to wage a fairly negative campaign. The challenger attacked the incumbent as an incompetent, wasteful manager who was looking out for the special interests. As is common in such campaigns, the negative mailings were sandwiched between more positive mail, so that the first and last pieces were upbeat assessments of the sponsoring candidate that presumably established and then restored a good feeling toward the sponsor of the hard-hitting mail.

In addition to these direct mailings, voters received mail from the incumbent's campaign. The town Republican chair revealed after the election that the incumbent mayor sent five mailings to all households containing a registered voter. This Republican candidate also placed some ads on the local cable television channel. The Republican mailings began arriving approximately 2 weeks before the election, whereas the Democratic mailings began about 3 weeks before the election. In addition to the mailings described earlier, all households with at least one registered voter got a single postcard from the Democratic candidate approximately 1 week before the election. Thus, the control group in the experiment received one Democratic and five Republican mailings.

The effectiveness of the campaign was measured through a postelection survey. Similar to the New Jersey assembly and Connecticut state legislative race, phone numbers were obtained from the campaign's voter file and supplemented by matches obtained from Survey Sampling Inc. The survey firm was provided with the phone numbers for 3,776 households, 1,009 control group households, and 2,767 treatment group households. This constituted all the control group households with phone numbers and a random sample of the treatment group households with phone numbers. There were 846 completed interviews for the treatment group and 288 for the control group.¹⁹

2.1.2. Results of Legislative and Mayoral Election Experiments

Tables 3, 4, and 5 show the effect of the direct-mail campaigns on candidate vote margins. Table 3 presents the results for the Connecticut and New Jersey state legislative races for those households included in the original campaign plans (there are three pairs of columns showing the control and treatment group voting patterns for Sample 1 from the Connecticut House election, and Samples 1 and 2 for the New Jersey Assembly election).

The results show that incumbent spending did not produce a consistent increase in the sponsor's vote margins. The first two columns record the results for the Connecticut House election. In Row 7, I show the incumbent's average vote margin for the control and treatment groups. The vote margin is calculated

TABLE 3: Effect of Treatment on Vote Margin

	<i>CT State Legislature Incumbent, 2000</i>		<i>NJ State Assembly Incumbents, 1999</i>			
	<i>Democrat and Unaffiliated Households</i>		<i>Prime Democrats</i>		<i>Other Democrats, Some Independents</i>	
	<i>Control</i>	<i>3 Mailings</i>	<i>Control</i>	<i>4 Mailings</i>	<i>Control</i>	<i>6 Mailings</i>
2 Republicans	—	—	2.9%	2.5%	6.0%	10.5%
1 Republican	17.1%	20.2%	0%	0.4%	1.2%	1.3%
Split ballot/no vote/ other	18.7%	18.8%	47.8%	44.5%	55.1%	53.2%
1 Democrat	64.2%	61%	3.6%	2.5%	5.4%	2.6%
2 Democrats	—	—	45.7%	50.2%	32.3%	32.5%
<i>n</i>	123	480	138	281	167	314
Average vote margin	0.47	0.41	0.89	0.98	0.57	0.45
Standard error	0.07	0.04	0.09	0.07	0.09	0.07
Treatment effect		−0.06		0.08		−0.12
Standard error		0.08		0.11		0.12

NOTE: For the New Jersey 1999 race, I report vote choice of those respondents who, according to public records, actually voted. For the Connecticut 2000 race, vote choice and participation are based on the survey responses. This accounts for the lower rate of Split/No Vote/Other in Connecticut.

according to the following formula. If a subject reported voting for the incumbent, the vote margin for the ballot is +1; if the voter said they did not vote, or did not report their vote, I code the vote margin as 0, and if the voter reported voting for the Republican challenger, the vote margin is −1. The average vote margin is the average of the vote margins for subjects in each treatment and control group. The treatment effect is calculated by comparing the average vote margin for the treatment group and control group. In the case of the Connecticut House race, the average vote margin for the control group was 0.47, and the average vote margin for the treatment group was 0.41. Comparing the means of these two groups, the average vote margin for the control group is 0.06 lower than that of the treatment group, with a standard error of 0.08. These results imply that if 100 voters similar to those participating in the survey were exposed to the treatment, the sponsor's vote margin would *fall* by 6 votes.²⁰ This suggests that the incumbent's campaign was not effective. This result is not statistically significant, and I attribute the negative point estimate to sampling error.

Columns 3 through 6 of Table 3 show the results for the New Jersey assembly case. The calculations are adjusted to reflect the double member districts. If a ballot selected both Democrats, the vote margin is +2, and if the voter voted for both Republicans the vote margin is −2. For the Prime Democrat sample in New

Jersey (Sample 1), survey respondents who received the four mailings provided, on average, an extra 0.08 votes for the Democratic candidates (0.04 per candidate). This means that for every 25 people receiving the mailings who were like the survey respondents, the sponsors got on average a one-vote increase in their vote margin. In contrast, Columns 5 and 6 imply that those in Sample 2 (other Democrats, Independents) who received the six mailings were *less likely* to provide votes for the Democratic incumbents.

The standard errors for the Connecticut and New Jersey treatment effects are fairly large. Fortunately, supplementary data from the surveys can be used to increase the precision of the estimates. Including controls for reported party identification, ideology, education, age, and race reduces the standard errors a moderate amount.²¹

Regression analysis was performed on a sample formed by pooling the data from the New Jersey Sample 1 and Sample 2 experiments (Table 3, Columns 3 through 6). Ordinary least squares for the pooled sample (using vote margin as the dependent variable and the number of mailings as the independent variable) but including no covariates shows a -0.008 (standard error = 0.016) change in the vote margin per mailing. Including covariates in this regression changes the estimated negative effect to -0.011 ($SE = 0.014$) per mailing. To translate these results back into the marginal effect of the mail treatment, the -0.011 estimate implies that four mailings reduce the vote margin for the incumbents by -0.044 votes, on average, for a subject receiving four mailings.

Performing the same analysis for the Connecticut case produces estimates of -0.021 ($SE = 0.026$) per mailing when covariates are included, and an estimated effect of -0.002 ($SE = 0.024$) per mailing when covariates are included. If we are willing to aggregate the Connecticut Sample 1 and New Jersey Sample 1 and 2 experiments, the standard errors for the treatment effect estimates can be further reduced. Using inverse variance weights and the regression estimates obtained when covariates are included, the pooled treatment effect is -0.009 ($SE = 0.012$) per mailing. This suggests that the campaigns did not improve incumbent vote margins.

Table 3 and the subsequent regression analysis imply that the campaign mailings did not boost the incumbent's vote margin when the mailings were sent to the audience that the campaign originally targeted. The experimental design also permits the study of the effect of sending mail to households not part of the original campaign plan. It would be very reasonable to expect that these supplementary households would be less responsive to the campaign mailings than the households originally selected as strategic targets. The campaigns did not think these voters were worth targeting in the first place, and the campaign mailings were not designed with these voters in mind.

Table 4 shows the effect of the mailings on those voters added to the original campaign plans. The incumbent candidates in the Connecticut state legislature and New Jersey assembly experiments were Democrats, and the households added to the original campaign plans contained mostly Republicans. In Con-

TABLE 4: Effect of Treatment on Vote Margin (Supplementary Groups)

	<i>NJ State Assembly Incumbents, 1999 (Republicans, Low-Turnout Independents)</i>			<i>CT State Legislature Incumbent, 2000 (Republican Households)</i>	
	<i>Control</i>	<i>4 Mailings</i>	<i>6 Mailings</i>	<i>Control</i>	<i>3 Mailings</i>
2 Republicans	13.9%	12.3%	11.2%		
1 Republican	0%	1.8%	1%	48.2%	35.7%
Split ballot/ no vote/other	75.9%	75.4%	69.4%	10.7%	28.6%
1 Democrat	1.9%	1.8%	1%	41.1%	35.7%
2 Democrats	8.3%	8.8%	17.4%		
<i>n</i>	108	57	98	56	28
Average vote margin	-0.09	-0.07	0.12	-0.07	0
Standard error	0.09	0.13	0.11	0.13	0.16
Treatment effect		0.02	0.22		0.07
Standard error		0.15	0.14		0.21

NOTE: For additional details regarding coding, see Table 3 note.

necticut, the “supplementary” treatment group of households (Sample 2) was sent three mailings, and in New Jersey, the supplementary treatment groups created by random assignment (Sample 3) were sent either four or six mailings.

The experiments reported in Table 4 were small, and so the estimated effects have large standard errors. It is nevertheless very interesting to observe that, in all three experiments, the treatment appeared to work for the incumbent. Vote margins for the incumbent were higher in the treatment groups than in the control groups, and in the case of the six-mailings group in New Jersey, the vote margin for the incumbent was much higher for the group receiving the treatment. Regression analysis was performed to improve the precision of the estimates. For the New Jersey case (pooling the 4 and 6 mailing groups), the coefficient estimate without covariates was 0.033 ($SE = 0.023$) per mailing, and with covariates was 0.034 ($SE = 0.021$). For the Connecticut case, the coefficient estimate before adding covariates was 0.024 ($SE = 0.069$). Given the small size of this experiment, including the full set of covariates increased the standard error for the treatment effect. When the set of control variables is restricted to only dummy variables for reported party identification, the coefficient estimate for the effect of mailings was -0.057 ($SE = 0.067$) (note the sign switch). Examining the data further, the change in the coefficient estimate reflects a weak correlation (not statistically significant) between receiving the mailings and Republican party identification. Pooling the estimated effects across all three cases reported in Table 4 and using covariates produces a treatment effect of 0.026

TABLE 5: Effect of Treatment on Vote Margin, Connecticut Mayoral Challenger, 1999

	<i>Control</i>	<i>9 Mailings</i>
Incumbent	31.3%	25.9%
No vote/other	45.5%	47.9%
Challenger	23.3%	26.2%
<i>n</i>	846	288
Vote margin	-0.08	0.004
Standard error	0.025	0.043
Treatment effect		0.083
Standard error		0.050

NOTE: The campaign mailings were sent by the challenger.

($SE = 0.020$) per mailing. Although it is hardly conclusive, the experimental results provide some puzzling evidence that the mailings worked on those who would normally be excluded from the mailings and failed to work on those who were targeted.

The experiments reported in Tables 3 and 4 measured the effect of incumbent spending. Table 5 reports the effects of the challenger's direct-mail campaign. The mailings had a strong positive effect on the challenger's performance. Comparing respondents in the treatment and control group, the incumbent's vote margin was reduced by one vote for each 12 respondents receiving the mailings. This effect was statistically significant at the 5% level using a one-sided test.

The precision of the estimate can be improved by using covariates. The regression coefficient estimating the effect per mailing is 0.009 ($SE = 0.006$) without covariates, and 0.012 ($SE = .005$) when controls are included. Based on these results, there is a reasonable basis to conclude that the challenger's campaign improved the challenger's vote margin by a politically meaningful, as well as statistically significant, amount.

2.2. WARD-LEVEL EXPERIMENTS

2.2.1. Description of the Case and Design

Experiments were performed to measure the effect of campaign mailings on the incumbent's vote margin in a congressional primary and general election. In these experiments, the randomization was done at the ward level rather than at the household level. The vote is publicly available at the ward level, and so there was no need to rely on postelection surveys. Because surveys may be biased due

to nonresponse and inaccurate reporting, the ability to obtain vote totals directly from voting statistics is a clear advantage to working at the ward level.

Congressional primary election. The incumbent member of Congress was expected to win the primary without great difficulty. He faced an opponent he had beaten on two previous occasions. The challenger had won a significant share of the vote in a previous contest, however, and so the incumbent was attentive to the race. In the end, the incumbent won an easy victory.

Participation in party primary elections is typically low, and the campaign targeted its mailings at households with some record of participation in recent primary elections. The campaign categorized households as “casual” or “consistent”, based on primary election voting history.²² A consistent household contained at least one voter who participated in all three of the last primaries, and a casual household had at least one voter who voted in at least one (but not all three) of the previous primaries. The campaign sent out two pieces of mail to the approximately 9,687 consistent households, and four pieces of mail to 28,468 casual households. The incumbent’s campaign reported that the challenger did not send out any mailings but did run spots on local radio for several days prior to the election.

The overwhelming majority of the district’s voters live in one county, and the experiment was conducted in this county. The county is divided into a few hundred wards, and the effect of the campaign was measured by comparing treatment and control wards. After eliminating 49 very large wards, 5 wards were randomly excluded from the mailings to serve as a control group.

Congressional general election. In the general election, the incumbent congressman faced a very weak opponent. Sources within the incumbent’s campaign reported that the challenger did not actively campaign.

The incumbent’s general election campaign concentrated on households with a track record of participation in the general election (at least one voter who voted in three of the past four general elections), with the exception of households where voters also had a strong record of participation in party primaries. Households determined to be “hard D,” based on consistent participation in Democratic primaries, were ignored. The 15,328 households with consistent participation in the Republican primaries (“hard R” households) were only sent one mailing, rather than the three sent to other target households (77,861 households).²³

The experimental design was similar to that used in the primary election. After eliminating large wards and the wards excluded as controls during the primary election, the sample was sorted into six strata based on the candidate’s average vote share in 1998 and 2000. Five wards were randomly excluded from each strata to form a control group, for a total of 30 excluded wards.²⁴

TABLE 6: Effect of Treatment on Vote Margin, Congressional Incumbent 2002

	<i>Primary Election</i>			<i>General Election</i>	
Treated	0.027 (0.014)	0.024 (0.008)	0.028 (0.008)	0.003 (0.008)	0.002 (0.005)
2000 net votes	—	0.43 (0.09)	0.42 (0.09)	—	0.58 (0.04)
1998 net votes	—	0.24 (0.06)	0.16 (0.06)	—	0.23 (0.06)
1996 net votes	—	—	0.19 (0.07)	—	—
R^2	0.40	0.60	0.61	0.87	0.95

NOTE: Dependent variable is vote margin in 2002 (incumbent votes minus challenger votes) divided by the number of registered voters in the ward who are registered with the incumbent's party. Observations are weighted by precinct size. General election regressions include dummy variables for the six sampling strata. Robust standard errors in parentheses.

2.2.2. Results of Congressional Primary and General Election Experiments

Table 6 shows the results of the mail campaign for the primary and general election. The primary election campaign was very effective. The incumbent's vote margin was increased by approximately 2.7% of the total number of registered voters in those wards that received the campaign mailings. The magnitude of this effect is large, even too large, though it is possible that the primary electorate is especially receptive to persuasion.²⁵ The incumbent's general election campaign produced only a very small, statistically insignificant boost in the candidate's vote margin. The point estimate of 0.2% is swamped by the 0.5% standard error.

2.3. HOW MUCH DOES A VOTE COST?

In Section 1, I described the conflicting conclusions reached in the literature measuring the effectiveness of campaign spending. These estimates of the cost of increasing incumbent vote margins in congressional elections ranged widely, from approximately \$15 per vote to \$350 per vote. According to previous studies, challenger spending produced a vote for a price ranging from \$12 to \$110 per vote. It is interesting to compare these previous findings with the experimental results. However, the results of the experiments are not directly comparable to the cost per vote figures for congressional elections, since the campaigns studied here are generally for lower offices.

There are two reasons to expect that the cost of getting votes will be lower in state and local elections.²⁶ First, the candidates are less well known. If campaign

spending wins votes by providing new information, then the marginal contribution of additional information will be greater when voters begin with less information. Second, in elections for national office, voters typically vote their partisan predisposition. This is true for lower offices but less so; as a result, in elections for state and local office, there are more marginal voters who might be induced to switch their votes.

Using the results from the direct-mail experiments, the approximate cost per vote can be calculated. Table 7 shows the estimated cost of increasing the sponsor's vote margin by one vote for each of the cases studied. The table assumes that there are on average 1.5 voters per household, and that the mailings cost 50 cents each to print and deliver. This implies that it costs approximately 33 cents per mailing to reach a voter. The calculations reported in Table 7 are performed using the estimated spending effects as well as the estimated spending effects plus or minus one standard deviation.

There are four main findings. First, the cost per vote for the state and local incumbents is too large to be measured (recall the point estimates are negative), whereas the congressional incumbent's general election campaign spent nearly \$200 per vote. This reflects the low estimated returns to the incumbents' campaign mailings. Second, and in strong contrast, challenger spending was effective. The mayoral challenger spent around \$30 per vote. Adding very appropriate caveats about the potential idiosyncrasy of the specific cases studied here, these two results are qualitatively consistent with Jacobson's (1978) finding that incumbent spending does not improve the incumbent's vote share, whereas challenger spending does.

The conclusion that incumbent spending did not work, whereas challenger spending did, speaks directly to the extensive literature on incumbent and challenger spending effects in general elections. The remaining results, although they address less-well-studied situations, are in some respects just as provocative. A third finding was that the incumbent Congressman's primary campaign was quite effective at increasing his vote margin. As previously stated, there are indications that the estimated effect of the campaign may be exaggerated, but even if the campaign were half as effective as the estimates suggest, there is still a strong contrast with the congressional general election findings. The primary election result is only partially consistent with prior theoretical speculations about when spending might be more effective. The incumbent is well known, which, according to conventional theory, would tend to depress spending effects. A reasonable alternative line of argument is that in a primary election, where partisanship is not an issue, there are many marginal voters, and therefore campaign spending is especially effective.

Finally, Tables 4 and 7 show that when the campaign mailings are sent to voters not originally targeted to receive mailings, the mailings are effective.²⁷ This finding is striking, and it raises the question, were the original targeting

TABLE 7: Estimated Cost Per Vote

	Mailings Required to Produce One Vote, Using Estimated Effect +/- One Standard Deviation				Cost Per Vote, Using Estimated Effect +/- One Standard Deviation	
	Mailings Required to Produce One Vote	Estimated Effect + One SD	Estimated Effect - One SD	Cost Per Vote, Using Estimated Effect	Estimated Effect + One SD	Estimated Effect - One SD
State and local incumbents, campaign targets	—	333	—	infinite	\$109.90	infinite
State and local incumbents, supplementary targets	38	22	167	\$12.70	\$7.20	\$55.00
Mayoral challenger	83	59	143	\$27.50	\$19.40	\$47.10
Congressional incumbent, primary	36	28	50	\$11.80	\$9.20	\$16.50
Congressional incumbent, general election	500	143	—	\$165.00	\$47.10	infinite

NOTE: The calculations for state and local incumbents and the mayoral campaign are based on estimated effects when control variables are included.

plans consistent with the assumption that incumbents do the best they can to win elections? This question is addressed in the next section.

SECTION 3: A MODEL OF INCUMBENT CAMPAIGN STRATEGY

The experimental results suggest that spending does not improve incumbent vote margins. In this section I explore whether it is possible to reconcile the experimental findings with the view that incumbents are well informed and rationally pursue their election goals.

Campaign strategy is in critical respects analogous to an insurance purchase. My analysis elaborates the implications of this observation, which follows directly from using maximization of the probability of victory, rather than vote share maximization, as the candidate objective. It is surprising that this choice of candidate objectives matters. The difference between these alternative objective functions is often treated as a minor nuisance, to be dealt with by a footnote citing literature exploring conditions leading to their equivalence.²⁸

Candidates who maximize their victory probability rationally focus their attention on the scenarios under which their campaign activity might be pivotal in determining the election result. For the typical incumbent, who expects an easy win, this will entail focus on the worst-case scenario. One difference between existing accounts and my own is that previous scholars suggest that the incumbent's attention to the worst case is based on distorted thinking. For example, in his section on incumbent campaigns, Jacobson (2001) very perceptively observes that incumbents "are inspired by worst-case scenarios," but adds that this is "rather than objective probabilities" (p. 91).²⁹

Before applying these ideas to campaign strategy, it is useful to describe a familiar situation in which an individual makes a choice in the face of uncertainty. Consider the decision to buy insurance. At the basis of every insurance contract is uncertainty about future events. The set of possible outcomes are called the "states of the world," and insurance contracts exchange a premium for promised payoffs in different states of the world. When you buy fire insurance, your house may burn down or it may not. In the state of the world where your house burns, you get a payment, and in the state of the world where your house is fine, you get paid nothing. To continue this example, assume that you know there is a 1% chance your \$100,000 house will burn down. Your expected dollar loss from fire in the absence of insurance is therefore \$1,000 (1% times \$100,000). If your house burns down and you are insured, you are paid the value of your house (\$100,000). Given the well-known imperfections in the insurance market, suppose the insurance companies offer fire insurance for a \$1,500 premium. That is \$500 more than your expected loss.

Why might you still buy insurance, even though it is, on average, a money-losing proposition? Insurance transfers money across states of the world, moving dollars from states where there is a low marginal utility of wealth to states where there is a high marginal utility of wealth. In the state of the world where there is no fire, you are relatively rich (you have a house), and when you are relatively rich, the \$1,500 insurance premium is not important. In the state of the world where there is a fire, you are relatively poor and you need money (your house just burned down). Since the insurance policy pays off when your need for money is heightened, by purchasing insurance your expected *utility* is raised. However, in dollar terms, buying insurance is a bad bet. Your expected wealth is lowered because you lose \$500 on average. A lesson to take from this example is that it may make perfect sense to forgo money “on average” in exchange for having more money when you really need it.

The logic of the insurance market can be applied to the incumbent’s spending decision. Consider an election featuring a Republican incumbent. I start with three assumptions about incumbents and their goals. First, the incumbent’s objective is to win the election. If two alternative campaign strategies produce the same probability of victory, the incumbent breaks the tie by selecting the strategy that maximizes expected vote share.³⁰ Second, incumbents pursue a rational strategy and therefore spend their money in the best way possible to achieve re-election. Third, incumbents know the effectiveness of alternative spending strategies.

Assume that the electorate can be divided into two equally sized groups (Republicans and Democrats) and that campaigns can send targeted messages (a targeted message goes only to voters in the targeted group). Assume also that turnout levels are fixed, and so campaign spending is aimed at winning support. To simplify still further, suppose that campaign strategy can be reduced to a single decision: the candidate’s choice of whether to target the campaign at Republican or Democratic voters.

When the voters in a group receive a campaign communication from a candidate, the utility the voters associate with the sponsor increases, according to how persuasive they find the communication. The extent to which campaign spending translates into votes is a function of two factors: the size of the utility boost associated with the campaign communication and the level of support for the sponsor prior to the communication.³¹ If the voters in the target group are initially evenly split between the candidates, a given boost in voter utility can have a large effect, whereas if the target group is already strongly supportive, the increase in support associated with the same boost in utility is assumed to be small (by definition, the boost in support must be near zero when initial support is close to 100%). Last, assume that it is easy for the incumbent to craft effective (utility-boosting) messages for his or her own partisans, and more difficult to persuade voters from the opposite party.³²

3.1. THE INSURANCE PERSPECTIVE AND THE RETURN TO INCUMBENT SPENDING

To examine the theoretical relationship between campaign spending and incumbent vote shares, I consider two political scenarios: the very common case in which the incumbent anticipates an easy victory and the rare case in which the incumbent faces a close election.

Case 1: easy win. As the campaign begins, the Republican incumbent expects an easy re-election. There is always some uncertainty, however, and this is captured by defining three possible states of the world on election day: good (things are unusually better than the incumbent expected), typical (things are as things usually are), and bad (things turned out much worse than the incumbent expected). When the incumbent campaigns, I assume he is uncertain which of these states will actually hold.

Calculating campaign effects requires a benchmark for comparison that specifies candidate performance in the absence of the campaign (referred to as “*ex ante*” performance). Suppose that, if the incumbent engaged in only minimal spending, the incumbent would win 95% of the Republican vote and 45% of the Democratic vote when the state of the world turned out to be typical. This translates into the incumbent’s receiving 70% of the vote overall (Democrats and Republicans compose equal shares of the electorate). If things turn out “good,” the incumbent wins 100% of the Republican vote and 60% of the Democratic vote. That results in the incumbent receiving 80% of the overall vote. If things turn out “bad,” the incumbent gets 60% of the Republican vote and 30% of the Democratic vote and wins just 45% of the vote overall. For purposes of this example, let the probability that the state of the world is good equal θ , the probability of bad equal θ , and the probability of typical equals $1 - \theta$.³³

Should the Republican incumbent target Republican or Democratic voters? This depends on the return to campaign spending. A very simple formulation would have campaigning result in a $\beta_i[100 - V_i]$ boost in the incumbent’s group i vote share, where β_i is a number between 0 and 1, V_i is the vote share given to the incumbent by group i *ex ante*, and $i = D$ or R , representing the Democratic and Republican voters, respectively. The assumption that it is relatively easy for Republican incumbents to inspire Republican voters is captured by assuming that $\beta_R > \beta_D$. For purposes of this case and Case 2 described next, suppose that $\beta_R = 0.3$ and $\beta_D = 0.1$.³⁴

To identify the incumbent’s optimal targeting strategy, we compare the effect of alternative campaign plans on the candidate’s probability of victory. To determine which targeting decision is best, we can restrict our attention to the state (or states) of the world where the campaign’s strategy choice can make a difference in whether the incumbent wins or loses. I will refer to the state of the world where the incumbent’s decision can change the election outcome as the “politically important state.”³⁵ In Case 1, the politically important state is

the “bad” state; in the bad state, the incumbent’s overall vote share *ex ante* does not exceed 50%.

If the incumbent chooses to target Republican voters, the probability of victory increases by θ ; there is a θ % chance that the bad state occurs and the decision to target the Republicans produces a 12% improvement in the incumbent’s vote share among Republican voters. This boosts the Republicans overall vote share 6%, which is sufficient for a 51% to 49% incumbent victory. In contrast, if the Republican incumbent targets the Democratic voters, too few additional votes are won to produce a victory; targeting the Democrats produces only a 3.5% boost in overall vote share in the politically important state. Therefore, under the conditions of this example, the incumbent’s optimal campaign strategy is to target the Republicans.

Targeting the Republicans maximizes the incumbent’s chances of victory. What effect does this decision have on the incumbent’s average vote share? The general formula for the boost in expected vote share from campaign spending is $(1/2) ([1 - \theta] \beta_i [100 - V_i | (\text{state} = \text{typical})] + \theta \beta_i [100 - V_i | (\text{state} = \text{bad})])$, which is the sum of the increase in a targeted group’s *ex ante* vote share in the typical and bad states, respectively, weighted by the probability that the typical and bad states occur, all multiplied by 1/2 (since I assume that each group has an equal share of the electorate). Substituting the Case 1 values for Republican vote share and β_R , the average increase in incumbent vote share from targeting the Republicans is $(1/2) (12\theta + 1.5[1 - \theta])$. The boost in average vote share from targeting the Democrats is $(1/2) (7\theta + 5.5[1 - \theta])$. Assuming that the bad state of the world is very rare (θ is close to zero), which for nearly all congressional incumbents it surely is, the average return from targeting the Republican voters is an approximately 3/4% boost in average vote share. The return to targeting the Democrats is approximately four times higher, a 3% boost in overall vote share.

The analysis shows that if the incumbent chooses the best campaign strategy, there is only a very small average return to spending. Despite the much higher average return to targeting the Democrats, the incumbent’s best choice is to target Republican voters. This follows from the incumbent’s objective: to win the election. That a “rational” incumbent may ignore vote share maximizing strategies and instead select a strategy with minimal average returns exemplifies something much more general. When formulating strategy, the only state of the world that matters for the incumbent is the politically important state, and this state may occur very rarely. Since the incumbent ignores whether or not the strategy that is best in the politically important state also has a large payoff under typical conditions, it is very possible that the optimal strategy will generate average vote shares far lower than the highest possible levels. This logic should be familiar from the insurance example: Just as fire insurance is on average a money losing proposition, but provides a payoff when it is really needed, so campaigning to the Republican voters will generally “lose votes” relative to an alternative strategy. It is, however, still the best strategy, since it wins votes when votes really matter.

Case 2: the close election. How are the conclusions regarding the return to spending altered when the election is close? To investigate this question, I change the *ex ante* vote shares associated with the good, typical, and bad states of nature. For Case 2, let the incumbent's share of the Republican vote be 90%, 70%, and 50% in the good, typical, and bad states, respectively. Let the incumbent's share of the Democratic vote be 40%, 30%, and 20% in the good, typical, and bad states, respectively; the overall incumbent share of the vote is then 65% in the good state, 50% in the typical state, and 35% in the bad state. For simplicity, assume that, as in Case 1, the probability of the typical state is $1 - \theta$, and the bad state is θ .

In Case 2, the "politically important state" is now the typical state of the world. In the typical state, targeting either the Republican or the Democratic voters will produce enough votes to secure victory, whereas neither strategy would provide enough votes to produce a win if the bad state holds. Since targeting the Republican voters generates a larger boost in average vote share for the incumbent, following the tie-breaking criterion set out in the assumptions about incumbent behavior, the incumbent again chooses to target the Republicans.³⁶

In Case 2, the average return to spending when the candidate targets Republican voters equals $(1/2) ([1 - \theta] 0.3[100 - 70] + \theta 0.3[100 - 50]) = (1/2) (9[1 - \theta] + 1.5\theta)$. Assuming the bad state is rare, the average vote return from the optimal targeting strategy in Case 2 is approximately 4.5%. This return to spending is much greater than the 3/4% boost produced by the optimal spending strategy in Case 1.

Why is the average return to incumbent spending so much greater in Case 2 than in Case 1? First, in Case 2 the politically important state is also the typical state. This is the critical difference between the two cases. When the politically important state is also the typical state, as it is when the election is expected to be close, then the strategy that produces the most votes in the politically important state will (almost) by definition be the best strategy for raising average vote share; what is best on average obviously depends heavily on what is best under typical rather than rare conditions. An additional factor producing high average return to spending is that when the election is close, the *ex ante* level of support for the incumbent is lower, and so there is more room for incumbents to improve their vote share.

To summarize the implications of Cases 1 and 2, there are two important conclusions. First, it may be perfectly rational for the incumbent to forgo targeting the voter group that produces the largest average return. In the examples, when the election was not close (Case 1), the incumbent pursued the optimal strategy and yet failed to target the vote-maximizing voter group. In the experimental case studies of incumbent campaign spending effects, all the studies were conducted in elections in which the incumbent expected easy victory. The experimental finding that the incumbent campaigns had minimal effects, while alternative strategies might have produced more votes, is *exactly what we would*

expect from the analysis of Case 1. Second, the model inspired by the experimental results was extended to analyze a political scenario not studied experimentally. The analysis of Case 2, the close election, produced a second important finding: a prediction about how campaign spending effects might vary with incumbent electoral safety. Comparing Cases 1 and 2 suggests that the average return to incumbent campaign spending is higher when the election is expected to be closer.

3.2. BROADCASTING VERSUS TARGETING: CAN RATIONAL CAMPAIGN SPENDING LOSE VOTES ON AVERAGE?

In Jacobson's (1978) seminal work on campaign spending in U.S. House elections, cross-sectional regressions of incumbent vote share on incumbent spending produce a negative coefficient for the effect of spending. This important result is widely viewed as an artifact of model misspecification, but this may be incorrect. Is it possible that incumbent spending may, on average, lower the incumbent's vote share?

Consider the case in which, rather than targeting spending at Democrats or Republicans, candidates select a broadcasting strategy, and the campaign message is broadcast to the entire electorate. When campaigns can selectively target voter groups, the return to targeting each group was a function of β_i , which specified the return from a communications strategy best tailored to appeal to the targeted group. In the case of broadcasting, different broadcasting strategies are characterized by $\beta = (\beta_D, \beta_R)$, a vector of effects.

The important limitation of broadcast communications is that campaign messages are viewed widely and the most effective message for one type of voter may be ineffective or worse for other voters. It is very plausible that some messages can cost you support. If the incumbent runs an advertisement containing a partisan endorsement, or emphasizing work for or outspoken support of a partisan cause, this will bolster support among the incumbent's partisans but reduce support among voters of the opposite party.

To illustrate how an *optimal* broadcast strategy can lose votes, briefly consider how the analysis would change in Case 1 if the incumbent were selecting a broadcast strategy rather than a targeting strategy. To simplify the algebra, suppose that, in the bad state, the incumbent vote share is just shy of 50% (alter Case 1 slightly so that, in the bad state, the incumbent gets slightly less than 65% of the Republican vote and slightly less than 35% of the Democratic vote). Let the incumbent's optimal campaign strategy be a broadcast strategy: $\beta^* = (\beta_D^*, \beta_R^*)$. Assume that this communications strategy is an appeal pleasing to the incumbent's Republican partisans but alienating to Democrats ($\beta_D^* < 0, \beta_R^* > 0$).

Is it possible that the optimal broadcast strategy loses votes on average? By restricting the set of feasible broadcast strategies, we can always assert that any given strategy β^* is the best feasible strategy. Since the incumbent can always

choose not to air the advertisement, a necessary condition for β^* to be an optimal campaign strategy is that it increase the incumbent's probability of victory. For this to be true, it must increase votes in the politically important state (in Case 1, this is the bad state). Inserting the Republican and Democratic *ex ante* vote shares in the politically important state, the broadcast increases the incumbent's probability of victory as long as $\beta^*_D [100 - 35] + \beta^*_R [100 - 65] > 0$, which can be simplified to Condition 1, $(\beta^*_R / -\beta^*_D) > 2$.³⁷ Does satisfying this necessary condition also imply that the broadcast strategy must increase the incumbent's average vote margin? The effect of using any broadcast strategy depends heavily on its effects in the typical state of the world. Assuming that the probability of the bad state is small, the average return to the optimal broadcast strategy β^* is approximately $\beta^*_D [100 - 45] + \beta^*_R [100 - 95]$. This average return is positive only if $(\beta^*_R / -\beta^*_D) > 11$, which is called Condition 2. Combining Conditions 1 and 2 implies that there exists a range of values for β^*_R and $-\beta^*_D$ ($11 > \beta^*_R / -\beta^*_D > 2$) such that, if β^* falls within the range, the best broadcast strategy increases the probability of victory and also on average loses votes.

The calculations show that when the incumbent expects an easy victory (Case 1), it is possible that the best broadcasting strategy increases the probability of incumbent victory while lowering the incumbent's expected vote share. The finding that a rational incumbent may select a strategy that lowers their average vote share is not at all an artifact of extreme assumptions. The essential element for this result is that the best broadcast strategy attracts Republican voters while alienating the Democrats. If the positive effect on the Republicans is sufficiently larger than the negative effect on the Democrats, this will ensure that the broadcast strategy produces an increase in the overall vote share in the politically important state, even if the incumbent starts with a larger share of the Republican vote *ex ante*. The average effect of the broadcast, however, depends on how it works in the typical case. Since the broadcast alienates Democrats, in the typical case the strategy will lose Democratic votes. Since in the typical case there is little if any room to increase already overwhelming support among Republicans, there will be no Republican vote gain sufficient to offset the losses the broadcast produces among Democrats. The net effect in the typical state of the world, and therefore on average, will be vote loss.

In contrast, a vote-losing broadcast strategy will not be used in the political conditions described by Case 2, where the incumbent expects a close election. A vote-losing broadcast would fail the "necessary condition" for an optimal strategy. When a close election is expected, the "important" state of the world very commonly occurs. Strategies tailored to win votes in the politically important state in Case 2 will therefore win votes on average. Finally, it should be noted that the average vote loss result does not stem in any way from miscalculation by the incumbent or uncertainty about the effects of campaign communications.

3.3. ADDITIONAL PREDICTIONS AND FURTHER IMPLICATIONS OF THE MODEL

My arguments about the effects of incumbent spending can be used to interpret existing empirical evidence and generate predictions about the results of new empirical inquiries. They can also help explain the strategic decision making of officeholders engaged in activities other than planning their campaign spending strategies.

In Section 1, I described three of the most important studies of campaign spending effects in U.S. House election: Jacobson (1978), Levitt (1994), and Erikson and Palfrey (2000). As Tables 1 and 2 show, these articles come to very different conclusions about the effect of incumbent campaign spending. Is there any way to reconcile these conflicting studies? The model described here produced two predictions that are relevant to this question: (a) In elections where the incumbent expects an easy victory, the return to spending may be small or even negative, and (b) the return to incumbent spending is higher when the election is expected to be close. A further relevant consideration is that in the overwhelming majority of cases, the incumbent House member expects an easy victory.

Suppose for purposes of argument that the basic regression models predicting incumbent vote share as a function of candidate spending levels and standard control variables produce accurate estimates of candidate spending effects. Under this assumption, Jacobson's early regression findings are essentially accurate estimates of the "average" marginal effect of incumbent campaign spending, where the average is taken over the entire sample of elections in which an incumbent House member faces a challenger (Jacobson's sample). Jacobson's estimation results show (roughly) a small negative average marginal effect on incumbent vote share.³⁸

Jacobson's analysis pools together all contested elections, close and not close, and estimates a single coefficient for the effect of incumbent spending. My arguments suggest that when the election is expected to be an easy incumbent victory, the effect of incumbent spending is near zero or may even be negative. If these typical cases are averaged in with the relatively small number of cases where the election is close, then the coefficient produced by a regression that uses the entire sample of races would be dominated by the cases where the incumbent expects easy victory. The results would show a zero (or negative) return to spending.

The possibility that the Jacobson result is a function of pooling across different kinds of races can be examined directly by partitioning the sample of cases according to the expected closeness of the election, and then running the standard OLS regression for each of the subsamples. Erikson and Palfrey (2000) have already done just that, and they show that incumbent spending effects are much larger when the election is close. Erikson and Palfrey, however, argue that

omitted variable bias is a major problem with the simple regression models, and they interpret their findings in light of this belief. They argue, based on theoretical considerations, that the coefficient estimate produced using the small subsample of close elections is the accurate estimate of the true return to candidate spending. The much lower returns to incumbent spending Erikson and Palfrey observed in the less competitive elections are attributed to increasingly severe omitted variable bias. An alternative interpretation of Erikson and Palfrey's findings is suggested by this article: the estimates for both the close and noncompetitive elections are correct, and they accurately reflect how the average return to spending varies across elections.

It is certainly possible that the true effect of incumbent spending really varies with the closeness of the election *and* that the measured effect of incumbent spending also suffers from omitted variable bias. Can we make any progress in determining the extent to which these different interpretations of the Erikson and Palfrey results are correct? The Levitt (1994) article presents an important opportunity to explore this issue further. Recall that Levitt argues that, by comparing the performance of the same pair of candidates across different elections (the "Levitt subsample" of House elections), the effects of omitted variable bias can be eliminated. This suggests a straightforward test: Do the Erikson-Palfrey results hold up for the Levitt subsample of elections? To perform this test, one would follow Erikson and Palfrey and partition the Levitt subsample according to the expected closeness of the elections. If the results using the Levitt subsample follow the same pattern found in Erikson and Palfrey, this implies that spending effects vary with the closeness of the election even when there is no omitted variable bias. That would suggest that the pattern in incumbent spending effects reported by Erikson and Palfrey represents true differences in average incumbent spending effects across electoral environments, differences that follow the pattern predicted by my theoretical arguments about incumbent strategy. The arguments presented here suggest a prediction: The pattern reported by Erikson and Palfrey will still be present, even when there is no reason to suspect the presence of omitted variable bias according to Levitt's argument.³⁹

My arguments suggest a second empirical test. In noncompetitive elections, the effect of incumbent spending on incumbent vote share may be small or even negative. The effect of spending on the probability of incumbent victory, however, must always be positive, even though the boost in probability of victory will be smaller when the election is less competitive. Empirical analysis can be conducted to determine how the effect of incumbent spending on incumbent probability of victory (rather than vote share) varies across competitive environments and whether the effect of incumbent spending on the incumbent's probability of victory is always positive.⁴⁰

The ideas described here can also be applied to officeholder behavior more generally. Consider, for example, the puzzle of unpopular presidential decisions. Why might a re-election-oriented president push policies that are unpopular? To rationalize this behavior, scholars must identify a benefit to offset the

cost associated with an unpopular decision. One important recent study suggests that the benefit enjoyed by the president is that, by selecting an unpopular course of action, the leader signals a superior understanding of the consequences of alternative actions, which is a very desirable quality in a leader (Canes-Wrone, Herron, & Shotts, 2001). My argument suggests an alternative explanation. Presidents are willing to take actions that lose the support of some voters, in exchange for deepening the level of support they enjoy from others. A president will take actions that lower his approval rating from 65% to 55%, for example, if he knows that this will help make the 55% who approve of him rock-solid supporters. Maintaining broader but more shallow support exposes the president to a greater possibility of defeat if uncertain events break against him. Maximization of vote share may not correspond to maximization of victory probability.

SECTION 4: CONCLUSION

This article makes several contributions. I report new evidence on campaign-spending effects. The experiments suggest that incumbent spending has only a negligible effect on incumbent vote margins. In contrast, the challenger's spending appeared highly effective. Because these experiments are free from many of the methodological hazards that commonly afflict observational studies, this evidence is a useful contribution to the large literature measuring campaign effects. The article builds on the experimental findings to provide a new perspective on incumbent spending strategy. My analysis helps explain how the observed return to campaign spending varies with the closeness of the election, why in typical U.S. House elections incumbent spending may increase the incumbent's probability of victory and yet have only minimal effect on the incumbent's vote share, and how rational spending plans might *decrease* the sponsor's expected vote share. The theoretical analysis also explains the pattern of results found in some of the leading observational studies and produces testable predictions.

It is important to keep in mind the limitations of the evidence and analysis. The experimental evidence comes from only a handful of cases. Additional empirical studies are needed to make the treatment estimates more precise and to gauge the generality of the findings. It is certainly possible that the next set of experimental results will establish that incumbent spending significantly increases the incumbent's vote share. The theoretical arguments explaining the relationship between incumbent strategy and the return to campaign spending are novel and, I strongly suspect, on the right track. They need to be analyzed more rigorously. The examples I provide show that there are conditions under which the theoretical claims hold. Although the intuitions appear to be robust to small changes in the setup, the arguments should be explored in a more general framework to confirm their validity and provide further insights.

There are two additional implications of the research. First, field experiments may have a wider scope of application than initially thought. Partisan campaigns can be studied using field experiments because it is possible to find campaigns interested in participating. It is impossible to know what further opportunities, whether in the study of political behavior or other areas of research, are available to scholars who elect to make inquiries. Second, field experiments can lead to theoretical advances. Field experimentation is (merely) a measurement technology, and how it is used depends on the objectives of the researcher. If one suspects that observational research is subject to large and unpredictable biases, experimental research can be an especially valuable source of “facts worth explaining.” This article is an illustration of how experimental findings can serve as a catalyst for generating new theories.

NOTES

1. This effort was begun by Jacobson (1978) himself, who clearly understood the potential effects of omitted variable bias on the regression results and took some steps to address the issue.

2. Although these studies of campaigns are novel in the academic literature, political parties, interest groups, and consultants sometimes perform “tests” of their campaigns. The results of these tests are kept as trade secrets, and even when hints about the results are shared, the research designs are not made public.

3. For a review of the empirical literature through the mid-1980s, see Jacobson (1985). For a more recent review of the literature, see Squire (1995). I am aware of one paper, which uses a different argument from that presented here, that emphasizes the possibility that the effect of campaign spending may vary with electoral conditions. For details, see Moon (2002).

4. Studies of Senate elections show a similar pattern. OLS studies by Jacobson (1985) and Abramowitz (1988) show that incumbents get a small boost from the increase in campaign spending, whereas the challengers are helped much more. In contrast, Gerber (1998) uses two-stage least squares and finds that additional spending has a relatively large effect for both the incumbent and the challenger.

5. The move from Table 1 estimates to the numbers reported in Table 2 requires some assumptions about how spending improves the candidate vote margin. For these calculations, it is assumed that campaign spending increases vote share solely by persuading a fixed electorate, and so these calculations do not require specification of the closeness of the election. Furthermore, if the election is very close, the mix of persuasion and mobilization does not alter the cost of increasing the vote margin. For cases where the initial conditions favor one candidate, the cost of increasing the vote margin will vary with the relative importance of mobilization of new voters versus persuasion of existing voters. The qualitative features of the table are not very sensitive to which choice is made among reasonable alternative assumptions. The algebraic details are available upon request.

6. Results for Senate elections follow a similar pattern. The traditional OLS results show a cost of around \$100 per vote for the incumbent (around \$200 for a switch from challenger to incumbent), whereas the two-stage least squares estimates are in the \$20 range for both the challenger and incumbent.

7. Field experiments provide a direct measurement of the effect of a given real campaign treatment, and so there is no need to make the difficult, if not impossible, translation from a laboratory treatment and laboratory measurement of experimental effects to a real world campaign activity and its consequences for turnout or candidate choice.

8. My discussion of experiments is restricted to work in which there was random assignment of treatments. This excludes a few studies, which are sometimes labeled experiments, in which researchers either knew about or participated in political activity in one place and then found a politically similar place to serve as a comparison group (Bartell & Bouxsein, 1973; Blydenburgh, 1971; Reback, 1971). In each of these particular studies, it was also the case that the research did not examine the effects of an actual campaign, but rather scholars devised and in some cases also participated in the campaign activity being studied. This comment applies least well to Bochel and Denver (1971), who compared the effects of a multistage canvassing effort for the Labour party (only one stage involved students) on turnout and Labour's vote share from those voters living in two purportedly similar high-rise apartment buildings. The authors do not say how they decided which building to canvass and present no evidence from previous elections regarding the turnout and voting behavior of those in the two apartment buildings.

9. More recent work includes Gerber and Green's (2000) nonpartisan mobilization drive in New Haven, Connecticut.

10. Both candidates were Democrats.

11. In addition to these randomized trials, Jennifer Steen recently conducted an experimental study of the effect of candidate canvassing in a 2002 state legislative race. Results were unavailable at the time this article was prepared. Related work also includes a recent study by Wantchekon (2002), who measures the effect of alternative campaign messages in an election in Benin.

12. A more complete description of the experiments and results will be presented in future work. The results reported here are interim findings.

13. The Democratic incumbents both won re-election with about 65% of the vote. The assembly race was the only contest for state office; there were also a few county offices on the ballot.

14. The fraction of completed interviews for each of the seven groups is Sample 1, no mailings (138/553); Sample 1, four mailings (281/1094); Sample 2, no mailings (167/753); Sample 2, six mailings (314/1444); Sample 3, no mailings (108/848); Sample 3, one mailing (57/460); and Sample 3, two mailings (98/654). In addition, there were 22 interviews that could not be matched to a registered voter who was on the 1999 postelection voter file.

15. The incumbent won with about 75% of the vote.

16. Since candidate preference was measured using a telephone survey, households without a phone number on the voter file were also excluded. These telephone numbers were also supplemented using a phone match performed by Survey Sampling Inc. The initial randomization was performed on a list of 3,697 households. Prior to submitting the list to the survey firm, it was discovered that 275 of these households did not actually have phone numbers. As the randomization was not affected, this has no effect on the experiment, although if this had been discovered earlier, the number of random exclusions from the experimental sample could have been increased.

17. The completion rate by experimental group was as follows: Republican households receiving mailings (28/141), Republican households not receiving mailings (56/270), Democrat and Unaffiliated households receiving mailings (480/2317), and Democrat and Unaffiliated households not receiving mailings (123/694). An additional nine interviews were performed with individuals who answered the phone but could not be identified as voters registered at the address.

18. The incumbent won re-election by approximately 8%. The mayoral election was by far the most important and active contest on the ballot.

19. There were 12 additional completed interviews that could not be matched to a registered voter.

20. In Tables 3 through 5, I assign a value of 0 to voters who do not reveal how they voted. An alternative scheme is to assume that those who refuse to say how they voted selected candidates as other members of their experimental group did. Using this method has minimal effect on the estimates and would not alter the conclusions.

21. This improvement is not more dramatic in part due to the relative homogeneity of the experimental groups prior to randomization.

22. This is a slight simplification of the targeting plan.

23. The campaign also planned to send out two mailings aimed at seniors. It is unclear whether the households in the excluded wards were sent these mailings or not. If not, then the differential in the number of mailings sent to the treatment wards versus the control wards should be increased by approximately 15%. This ambiguity will be resolved once the consultant who handled the campaign returns from military service. Adding these mailings will only reinforce the conclusion that the general election spending did not boost the incumbent's vote margin.

24. In reviewing the data, I found that 1 of the 30 wards that should have been excluded was switched with a ward (in the same town) that should have been included. There is no reason to suppose this was consequential. The two wards are basically indistinguishable; they were not only from the same strata but had nearly identical values for the incumbent's average share of the vote in 2000 and 1998 (57.11% versus 57.15%).

25. My concern that this effect might be overestimated is partially based on preliminary analysis of a parallel household-level experiment suggesting a smaller boost in turnout than measured in the ward-level randomization.

26. This argument is offered with a general caveat that we lack a well-accepted theory of how campaign spending might influence voters, so speculation about relative impact across different types of elections should be read with some skepticism.

27. Recall that the mailings in question emphasized the incumbents' legislative efforts on issues with broad popular support. This result may not have held if the mailings had been partisan attacks. However, the incumbent always has the option of sending more positive mailings to voter groups outside his or her partisan base, and the evidence suggests that this would be a successful vote-producing strategy.

28. Maximizing the probability of victory is the most intellectually satisfying objective function, though most theoretical articles have candidates maximize vote share, since this is much easier technically. In most applications, it is a very minor issue.

29. Jacobson also writes: "Because of uncertainty, members tend to exaggerate electoral threats and overreact to them" (Jacobson, 2001, p. 91). My argument is also compatible with Fenno's description of how members of Congress aim to expand and then solidify their re-election constituencies (Fenno, 1978).

30. This tie-breaking rule simplifies the construction of examples, but it is not essential to any of the arguments.

31. For "negative campaigning," replace *utility* with relative utility.

32. This assumption is important and I think it is also reasonable. Candidates preparing ads aimed at their own partisans have a lot of material to work with: They can either remind or inform their own partisans about their endorsements and activities (a list of actions and votes filled with appealing partisan accomplishments and position taking) or else attack the opponent. Candidates trying to appeal to members of the opposite party have much less to work with, since it is unlikely they have a record of accomplishments appealing to voters from both parties. Although it is going outside the model (since the idea of targeting is that the message goes only to the intended audience), the set of sensible communications to voters outside your party is further constricted by the danger that appeals that are too accommodating may alienate members of your own party or be viewed by all voters as craven or deceptive. It is also possible that cross-party appeals are accepted uneasily and dismissed as not credible when they are faced with even minimal response. The important consequence of this assumption is that incumbents are encouraged to focus on their own partisans. Adding endogenous turnout might also serve this purpose.

33. Setting the probability of the good state equal to zero simplifies the exposition and has no substantive implications for the arguments.

34. The example could be constructed using much smaller differences in relative effectiveness of campaigning if, rather than assuming the vote change associated with campaigning was a function of the distance between initial vote share and 100, it was assumed that relative utility for the incumbent candidate among partisan groups was distributed according to a symmetric distribution with a peak at the distribution's mean value (e.g., probit, normal). In that case, the marginal change in votes

associated with a given shift in the location of the distribution would be a function of the initial cutoff point for voting for the incumbent, and by symmetry, the marginal change in incumbent vote share associated with a shift in the distribution's location given a base vote share of x percent would be locally equal to the marginal change for a group where the initial vote share was $(100 - x)$ percent. For the empirically relevant cases, this alternative formulation of the return to spending would boost the incumbent's relative return from campaigning among Republicans.

35. In the cases considered here, only one state of the world will be "politically important."

36. It should be clear that, because targeting the Republican voters generates more votes than targeting Democrats in the politically important state (in the typical state, targeting the Republican voters boosts overall vote share by 4.5%; targeting the Democratic voters boosts vote share by 3.5%), minor adjustments to the example, such as specifying the probability of states of the world on which the incumbent's share of the overall vote would fall between 45.5% and 46.5%, would eliminate reliance on a tie-breaking rule.

37. It might be reasonable to specify the effect of a communication that alienates the Democratic voters as $\beta_D[V_i]$ rather than $\beta_D[100 - V_i]$. This changes the algebra but does not alter any substantive conclusions.

38. As stated in the introduction, Jacobson did not argue that incumbent spending decreased incumbent vote share but rather concluded that incumbent spending was not effective at raising incumbent vote share.

39. Furthermore, a comparison of the estimates produced using the Levitt (1994) subsample and the entire sample for each of the partitions (that is, a comparison of the estimates produced using the Levitt subsample and the Erikson and Palfrey [2000] estimates) might help provide a measure of the relative importance of omitted variable bias versus true differences in the average effect of incumbent spending.

40. This analysis, as well as tests of the earlier prediction, might also wish to take into account the possibility of decreasing marginal returns to spending. For the second prediction, there is also still the problem of omitted variable bias.

REFERENCES

- Abramowitz, A. (1988). Explaining Senate election outcomes. *American Political Science Review*, 82(June), 385-403.
- Adams, W., & Smith, D. (1980). Effects of telephone canvassing on turnout and preferences: A field experiment. *Public Opinion Quarterly*, 44, 389-395.
- Ansolahehere, S., & Snyder, J. M., Jr. (1996). *Money, elections and candidate quality*. Unpublished manuscript, Massachusetts Institute of Technology.
- Bartell, T., & Bouxsein, S. (1973). The Chelsea Project: Candidate preference, issue preference, and turnout effects of student canvassing. *Public Opinion Quarterly*, 37, 269-275.
- Blydenburgh, J. C. (1971). A controlled experiment to measure the effects of personal contact campaigning. *Midwest Journal of Political Science*, 15(May), 365-381.
- Bochel, J. M., & Denver, D. T. (1971). Canvassing, turnout, and party support: An experiment. *British Journal of Political Science*, 1, 257-269.
- Canes-Wrone, B., Herron, M., & Shotts, K. (2001). Leadership and pandering: A theory of executive policy making. *American Journal of Political Science*, 45, 532-550.
- Eldersveld, S. J. (1956). Experimental propaganda techniques and voting behavior. *American Political Science Review*, 50, 154-165.
- Eldersveld, S. J., & Dodge, R. W. (1954). Personal contact or mail propaganda? An experiment in voting turnout and attitude change. In D. Katz, D. Cartwright, S. Eldersveld, & A. M. Lee (Eds.), *Public opinion and propaganda*. New York: Dryden.

- Erikson, R. S., & Palfrey, T. R. (2000). Equilibria in campaign spending games: Theory and evidence. *American Political Science Review*, 94, 595-609.
- Fenno, R. F. (1978). *Home style: House members in their districts*. Boston: Little, Brown.
- Gerber, A. (1998). Estimating the effect of campaign spending on Senate election outcomes using instrumental variables. *American Political Science Review*, 92, 401-411.
- Gerber, A. S., & Green, D. P. (2000). The effects of canvassing, telephone calls, and direct mail on voter turnout: A field experiment. *American Political Science Review*, 94, 653-663.
- Gosnell, H. (1927). *Getting out the vote: An experiment in the stimulation of voting*. Chicago: University of Chicago Press.
- Green, D. P., & Krasno, J. S. (1988). Salvation for the spendthrift incumbent. *American Journal of Political Science*, 32, 884-907.
- Jacobson, G. C. (1978). The effect of campaign spending in congressional elections. *American Political Science Review*, 72, 469-491.
- Jacobson, G. C. (1985). Money and votes reconsidered: Congressional elections, 1972-1982. *Public Choice*, 47(1), 7-62.
- Jacobson, G. C. (1990). The effects of campaign spending in House elections: New evidence for old arguments. *American Journal of Political Science*, 34(May), 334-362.
- Jacobson, G. C. (2001). *The politics of congressional elections* (5th ed.). New York: Longman.
- Levitt, S. (1994). Using repeat challengers to estimate the effect of campaign spending on election outcomes in the House. *Journal of Political Economy*, 102, 77-98.
- Miller, R., & Robyn, D. (1975). A field experimental study of direct mailing congressional primary campaign: What lasts until election day? *Experimental Study of Politics*, 5, 1-37.
- Moon, W. (2002). *The paradox of less efficient incumbent spending: Theory and tests*. Unpublished manuscript, Department of Political Science, University of California, Los Angeles.
- Ornstein, N. J., Mann, J. E., & Malbin, M. J. (2002). *Vital statistics on Congress, 2001-2002*. Washington, D.C.: AEI Press.
- Reback, G. L. (1971). The effects of precinct level voter contact activities on voter behavior. *Experimental Study of Politics*, 1, 65-97.
- Squire, P. (1995). Candidates, money, and voters: Assessing the state of congressional elections research. *Political Research Quarterly*, 48(4), 891-917.
- Wantchekon, L. (2002). *Market for votes: Evidence from a field experiment in Benin*. Unpublished manuscript, Politics Department, New York University.

ALAN S. GERBER is a professor of political science at Yale University, where he has taught since receiving his Ph.D. in economics from MIT in 1994. He heads the Center for the Study of American Politics at Yale University. His research interests include electoral politics, campaign finance, representation, voter turnout, and research methodology.

Request Permission or Order Reprints Instantly

Interested in copying, sharing, or the repurposing of this article? U.S. copyright law, in most cases, directs you to first get permission from the article's rightsholder before using their content.

To lawfully obtain permission to reuse, or to order reprints of this article quickly and efficiently, click on the "Request Permission/ Order Reprints" link below and follow the instructions. For information on Fair Use limitations of U.S. copyright law, please visit [Stamford University Libraries](#), or for guidelines on Fair Use in the Classroom, please refer to [The Association of American Publishers' \(AAP\)](#).

All information and materials related to SAGE Publications are protected by the copyright laws of the United States and other countries. SAGE Publications and the SAGE logo are registered trademarks of SAGE Publications. Copyright © 2003, Sage Publications, all rights reserved. Mention of other publishers, titles or services may be registered trademarks of their respective companies. Please refer to our user help pages for more details: <http://www.sagepub.com/cc/faq/SageFAQ.htm>

[Request Permissions / Order Reprints](#)