B.J.Pol.S. (2005) **36,** 269–290 *Copyright © 2006 Cambridge University Press*
doi:10.1017/S0007123406000159 *Printed in the United Kingdom*

Filling in the Blanks: A New Method for Estimating Campaign Effects

JOHN GEER AND RICHARD R. LAU*

Scholars have invested a great deal of effort in trying to estimate the impact of political campaigns on the public. While progress has been made, one fundamental problem continues to plague our attempts to study campaigns: the lack of good, detailed data about the behaviour of candidates. In the United States, for instance, the presidential battle is a national struggle that unfolds locally on the stages of fifty states, yet most data collection efforts have treated presidential elections as if they were national contests, conducted (implicitly) in an identical manner across the entire country. Using available (national) data as a baseline and *theory* to predict plausible variations from that baseline, the authors devise a method for simulating variation in presidential campaigns across states and over election years–-one of the crucial missing pieces of the puzzle. Their method generates a *range* of plausible effects, which is often narrow enough to shed light on important hypotheses. It can be employed whenever data are available at a more aggregate level than is desirable. This method is then applied to assess the debate over the impact of attack advertising on turnout. This approach suggests that campaign negativism *stimulated* (rather than demobilized) turnout in presidential elections from 1980 through 2000.

Campaigns matter to the practice of democracy. As William Riker forcefully reminds us, 'campaigns … are a main point – perhaps *the* main point – of contact between officials and the populace over matters of public policy'.[1] But estimating precisely how much this 'contact' matters *for voters* has proved difficult. Using a variety of data and theoretical frameworks, scholars have examined the influence of campaigns on the electorates' preferences on voters'[2] willingness to go to the polls,[3] and on whether campaigns can

* Department of Political Science, Vanderbilt University; and Department of Political Science, Rutgers University, respectively. The authors would like to thank various colleagues at North Carolina, Princeton, Rutgers, Stony Brook and Vanderbilt Universities, and in particular Jane Junn, Geoff Layman, Bruce Oppenheimer and Gerry Pomper, who have patiently listened to and offered valuable comments and suggestions on earlier presentations of the arguments made here. A special note of thanks goes to Larry Bartels who helped with this project in numerous ways.

[1] William Riker, *The Strategy of Rhetoric* (New Haven, Conn.: Yale University Press), p. 2.

[2] James E. Campbell. *The Presidential Campaign* (College Station: Texas A&M University Press, 2000); Robert S. Erikson, Joseph Bafumi and Bret Wilson, 'Was the 2000 Presidential Election Predictable?' *PS* (2001), 815–19; Steven E. Finkel, 'Reexamining the "Minimal Effects" Model in Recent Presidential Campaigns', *Journal of Politics*, 55 (1993), 1–21; Andrew Gelman and Gary King, 'Why are American Presidential Election Polls so Variable when Votes are so Predictable?' *British Journal of Political Science*, 23 (1993), 409–51; Richard Johnston, Michael G. Hagen and Kathleen Hall Jamieson, *The 2000 Presidential Election and the Foundations of Party Politics* (New York: Cambridge University Press, 2004); D. Sunshine Hillygus and Simon Jackman, 'Voter Decision-Making in Election 2000: Campaign Effects, Partisan Activation, and the Clinton Legacy', *American Journal of Political Science*, 47 (2003), 583–96; Thomas Holbrook, *Do Campaigns Matter?* (Thousand Oaks, Calif.: Sage, 1996); Richard Johnston, André Blais, Henry Brady and Jean Crete, *Letting the People Decide* (Stanford, Calif.: Stanford University Press, 1992); Richard R. Lau and Gerald Pomper, 'Effectiveness of Negative Campaigning in U.S. Senate Elections, 1988–98', *American Journal of Political Science*, 46 (2002), 47–66; Richard R. Lau and Gerry Pomper, *Negative Campaigning: An Analysis of U.S. Senate Elections* (Lanham, Md.: Rowman, Littlefield, 2004).

[3] Ansolabehere and Iyengar, *Going Negative*; Finkel and Geer, 'A Spot Check'; Paul Freedman and Kenneth M. Goldstein, 'Measuring Media Exposure and the Effects of Negative Campaign Ads', *American Journal of Political Science*, 43 (1999), 1189–208; Richard R. Lau and Gerald Pomper, "Effects of Negative Campaigning

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

inform and educate voters.[4] But despite a large (and growing) body of work on the general subject of 'campaign effects', we still face a number of major hurdles before we can understand how campaigns work.

The biggest problem facing researchers studying political campaigns is a *lack of detailed evidence about the nature of the campaigns themselves*. As Hillygus and Jackman recently observed, 'the chief impediment to a more comprehensive research agenda has been the limitation of appropriate data'.[5] In US presidential elections, a key problem is that the Electoral College requires the campaign to be conducted within fifty-one different settings. Yet our data collection efforts on the nature of the campaigns have tended to be national in scope.[6] This tendency has led to the implicit assumption that the campaign is conducted in an identical manner throughout the country, which we know is not true. We should, and can, do better.

The 2000 and 2004 presidential elections, in fact, have offered important developments in this regard. Efforts by Kathleen Jamieson, Doug Rivers, Ken Goldstein and others provided a much more detailed reading of these contests than previous presidential elections. Such efforts, if they continue, provide the chance to learn a great deal about future campaigns. But these important undertakings do not offer solutions to our lack of information about earlier campaigns and, of course, there remain holes in the current base of evidence about campaigns.

The purpose of this article is two-fold. First, we develop a new method that provides a way to fill in some of the blanks in our understanding of past presidential elections. This technique is actually quite general, and can be employed whenever data are available at a more aggregate level than theory suggests is desirable. Secondly, we use this method to test one of the most discussed hypotheses about campaigns – the so-called demobilization hypothesis. We turn first to our method and then apply it to the question of whether negative political advertisements depress turnout in elections.

OVERVIEW OF METHOD

Our core idea is simple, although its execution is more complicated. Therefore, we begin with a brief overview of our method in an effort to communicate our plan of attack clearly.

We start with existing national-level data from recent US presidential campaigns and treat them as a *baseline* – that is, a reasonable approximation of how candidates would have run had the campaign been a truly national affair. We know, of course, that candidates

(F'note continued)
on Turnout in U.S. Senate Elections, 1988–98', *Journal of Politics*, 63 (2001), 804–9; Lau and Pomper, *Negative Campaigning*.

[4] Larry M. Bartels, 'Messages Received: The Political Impact of Media Exposure', *American Political Science Review*, 87 (1993), 267–85; Craig L. Brians and Martin P. Wattenberg, 'Campaign Issue Knowledge and Salience: Comparing Reception from TV Commercials, TV News, and Newspapers', *American Journal of Political Science*, 40 (1996), 172–93; Kim Fridkin Kahn and Patrick J. Kenney, *The Spectacle of U.S. Senate Campaigns* (Princeton, N.J.: Princeton University Press, 1999); Darrell West, 4th edn, *Air Wars* (Washington, D.C.: Congressional Quarterly Press, 2005).

[5] Hillygus and Jackman, 'Voter Decision-Making in Election 2000', p. 583.

[6] We are, of course, not the first to notice this problem or to try and solve it. It is important, therefore, to acknowledge some of the earlier efforts by scholars to collect important data about political campaigns; see, for example: Steven Rosenstone, *Forecasting Presidential Elections* (New Haven, Conn.: Yale University Press, 1983); Thomas Holbrook, 'Presidential Elections in Time and Space', *American Journal of Political Science,* 35 (1991), 91–109; James E. Campbell, 'Forecasting the Presidential Vote in the States', *American Journal of Political Science,* 36 (1992), 386–407.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

vary their campaign according to the context of each state. Were this variation across states random, we would be just as well off assuming one national campaign. But *theory* informs us that this variation is not random. It reduces an infinitely large set of possible variations to a smaller and more manageable set of *plausible* variations. Our method permits us to identify this smaller set of plausible moves by the candidates.

Our approach rests on the premise that candidates seek to maximize their chances for victory. From this obvious starting point, we develop a set of decision rules that provide guideposts to the behaviour of candidates. We know, precisely, the incentives for carrying each state because we know how many electoral votes each victory would bring. We also know, with slightly less precision, a candidate's chances of winning the electoral votes of any state, because those odds can be estimated reliably with a number of different techniques, mostly involving variations of the normal vote. These two factors, we assume, influence the decision of how much effort should be invested in trying to win each state.

We then employ data from a coding of the political advertisements aired by the major party presidential candidates to ground our estimates of how they behave empirically. We view these data as 'baseline' measures of each candidate's campaign strategy. We then specify a set of theoretically grounded decision rules that adjust candidates' potential advertising strategies to the actual context of each state**.** Thus, even though we do not know the actual decisions made by candidates in the different states, we can, through our baseline data and the set of campaign decision rules, estimate a *range* of possible decisions. Thus our method does not purport to provide a single estimate of what candidates did, but rather a range of plausible campaign efforts.

To increase the chances that our estimates will be narrow enough for hypothesis testing, we go one step further. If the different campaign strategies are believed to be equally likely, standard statistical techniques allow one to estimate a mean and confidence interval for the plausible effects of the campaign. If the investigator believes some of the possibilities are more likely than others, the estimates can be refined through a procedure such as Bayesian Model Averaging which will take those prior beliefs into account.[7] This tactic helps sharpen the analysis and, therefore, increase the potential import of our method.

The remainder of this article is divided into four major sections. First, we briefly describe our baseline data. Secondly, we describe our method in detail. In so doing, we present some evidence that strongly suggests that we have in fact captured the strategic moves of presidential candidates. Thirdly, we apply our method to the current controversy over whether 'attack advertising' demobilizes the electorate. Finally, we discuss the implications of our findings and method.

BASELINE DATA: CANDIDATES' POLITICAL ADVERTISEMENTS

Presidential campaigns are large, multi-faceted undertakings. Our focus here is to estimate the effects that campaigns have on the public, not activists or members of the news media. Candidates seek the public's support in many ways, but the most visible and important of these actions involves the content of the paid advertising the candidates air on television. Advertisements not only constitute the set of appeals candidates offer to voters, but these bits of propaganda represent the single most costly part of the campaigns for the national office.[8] The fact that politicians invest such a huge proportion of their resources into

[7] Larry M. Bartels, 'Specification Uncertainty and Model Averaging', *American Journal of Political Science*, 41 (1997), 641–74.

[8] Frank Luntz, *Candidates, Consultants, and Campaigns (*New York: Blackwell, 1988); West, *Air Wars*.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

television advertisements underscores the importance they hold for these presidential hopefuls.

Fortunately, good data are available to us. Geer has compiled a detailed content analysis of 60 second or shorter presidential advertisements aired in the general election, between 1960 and 2000.[9] By using these data, we assume that candidates adopt a baseline campaign strategy that is reasonably reflected in the political commercials they air on television. One way to think about this assumption is that candidates have at their disposal a set of appeals they can make to the public in their quest for votes. Candidates, then, choose to emphasize a subset of these appeals depending on the context of the particular contest they are seeking to win. Thus, candidates might run one set of advertisements heavily in certain states and run different spots in other states, and/or some advertisements early in the campaign and others late in the campaign, depending on their beliefs about what approach will win the most votes.

While Geer's coding scheme captures a large number of key variables allowing one to assess multiple 'effects' of campaigns, we will limit our attention in this article to whether appeals in advertisements were self-promotional, and thus were 'positive' (for example, my policies have lowered unemployment); or attacked the opponent, and thus were 'negative' (for example, my opponent's policies have increased unemployment). As will be seen, this focus helps to limit further the number of campaign strategies that must be considered. Our basic measure, then, and the starting point of all of our data estimates, is the proportion of all appeals made by the major-party presidential candidates in every election campaign from 1980 up to 2000 that were negative (see Appendix Table A). We limit ourselves to the more recent election years because there is some suggestion that before this date most televised advertisements in presidential elections were aired on national, rather than local, media.[10] These data give us the chance to examine the effects of attack advertising on turnout in presidential elections.

MODELLING STRATEGIC BEHAVIOUR

While every candidate will behave in slightly unique ways, there is an underlying and widely shared goal in campaigns: the desire to win the election. That goal provides the opportunity to generate a set of considerations, representing the range of behaviour available to candidates. But to do so, we have to build a framework that captures the electoral context facing politicians. We start with estimating the chances candidates have to win a state. From there, we move to a discussion of the incentives for allocating campaign resources. These incentives reflect the constraints the Electoral College system imposes on presidential campaigns. With the strategic context sketched out, we then develop the set of campaign decision rules that are relevant for the particular hypothesis we seek to test. These rules flow directly from our assumption that candidates react to different electoral contexts rationally – i.e. with an eye to winning the election – although they reflect

[9] John G. Geer, 'Campaigns, Competition, and Political Advertising', in John G. Geer, ed., *Politicians and Party Politics* (Baltimore, Md.: John Hopkins Press, 1998); John G. Geer, *Attacking Democracy: A Defense of Negativity in Presidential Campaigns, 1960–2000* (Chicago: University of Chicago Press, 2006).

[10] In personal interviews both Carter Eskew, Gore's media consultant in 2000, and Alex Castellanos, Bush's media consultant in 2000 and 2004, made this claim about advertising buys in the early decades of television. See Geer, *Attacking Democracy,* for a complete discussion of these data.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

different beliefs about how best to accomplish that goal. When combined, these rule-based decisions result in an overall 'campaign strategy' that varies across the different electoral contexts – i.e., across the states.

Estimating Candidates' Beliefs About Their Chances in a State

A key starting point for our method is to determine the chances of winning particular states (or more precisely, the proportion of the state's votes the candidate expects to win). The ideal way to estimate candidates' expectations would be to have their polling data, or short of that, independent 'trial heat' polls from every state conducted some reasonable period of time before the election (for example, in August). Such data do not exist for a majority of the elections under consideration here, and to increase the generality of the presentation of our method, we will proceed in another direction.

Since we do not know exactly candidates' estimates of their support in each state, we generate three reasonable estimates of those chances. By so doing, we strive to measure a range of plausible expectations that captures the actual views of the candidates during these elections. All of these approaches are a close cousin to normal vote analyses.[11]

1. One possibility is *strictly historical* and very objectively calculated. It assumes a candidate will expect to do as well as his party's candidate did in the state over the past two presidential elections, with the immediately prior election counted twice as much as the election eight years earlier. We adopt this weighting scheme in the belief that the most recent election is the one that will figure more prominently in politicians' calculations.
2. A second estimate builds in an *optimistic bias* for all candidates. Here we assume that every candidate believes he will be able to do as well as the best his party's candidate had done in a state over the past two elections, plus 1 point.[12]
3. A third possibility seeks to approximate what a candidate with a very good pollster could do. We start with the actual vote share won by the two major candidates in each state, but add to those totals a random component with a mean of 0 and a standard deviation of 2.5 per cent, to approximate the typical margin of error associated with most polling data.

The resulting three estimates of expectations appear to be sound, correlating strongly with state polls available from Shaw in 1992 and 1996 (average Pearson's $r = 0.72$).[13] As a further check, we compared our measures of candidate expectations with Shaw's report of how consultants from 1992 and 1996 judged the competitiveness of the states. These judgements yielded a similar correlation (average $r = 0.73$). These findings suggest that we have tapped fairly accurately into this aspect of campaign decision making.

[11] Philip E. Converse, 'The Concept of a Normal Vote', in A. Campbell, P.E. Converse, W.E. Miller and D.E. Stokes, eds, *Elections and the Political* Order (New York: Wiley, 1966), pp. 9–39.

[12] In addition, we build in a 5-point boost in expectations for favourite-son presidential candidates, and a 2-point boost for favourite-son vice-presidential candidates. This approach squares with the work of Michael Lewis-Beck and Tom Rice, *Forecasting Elections* (Washington, D.C.: C.Q. Press, 1992).

[13] Daron R. Shaw, 'The Effects of TV Advertisements and Candidate Appearances on Statewide Presidential Votes, 1988–1996', *American Political Science Review*, 93 (1999), 345–62.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

Allocating Limited Campaign Resources Across States

There are at least two related factors beyond the chances of winning that go into the decision of how many campaign resources to devote to each state, both of which are shaped by the Electoral College: (1) the number of electoral votes there are to win – i.e., the *incentives* for winning the state; and (2) the need to secure 270 votes in the Electoral College. The first factor structures the payoff for winning a state. But candidates cannot simply focus on the large states, nor always try to win every statewide contest. In any given election year a number of states might, for all practical purposes, be impossible to win. So candidates may focus on those states that are winnable and contain enough electoral votes to capture the White House.

Incentives. The Electoral College makes the incentives for a candidate to win a state in US presidential elections easy to model. Very simply, incentives increase as the number of electoral votes gained by winning a state increase. The incentives for winning a state are identical for each candidate within an election year, although they shift slightly over time with reapportionment. We provide two alternatives here, one which says incentives for winning a state are proportional to the number of electoral votes it has, the second which says resources should be devoted in a manner which disproportionately benefits the larger states[14] – the same proportion as above, but raised to the power of 1.5.

Building a majority. The size of the state is just one variable that affects the calculations of candidates. The other is the need to win 270 electoral votes. We consider three different decision rules whereby beliefs about how to win a presidential election could affect the attention the campaign pays to a given state. Once again, we develop multiple possibilities to capture the many options open to candidates.

1. The first assumes that candidates believe *all states and all votes are equally valuable*, and try to win as many votes as they can in every state. Hence the perceived closeness of an election in a state should not have any effect on candidate behaviour. This rule advocates spending in proportion to the number of electoral votes in a state. This alternative seems less plausible to us than the others, particularly given limited campaign resources. But there were candidates in recent US history who had a reasonable chance of winning every state (Nixon in 1972, Reagan in 1984), and apparently tried to do so; again, we include this rule for the purpose of casting a broader net.
2. Another alternative is to concentrate one's local efforts on the states that are most likely to guarantee an *Electoral College majority*, and to ignore all other states. This rule assumes that candidates spend proportionally to the number of electoral votes within this first category (or disproportionately, according to the 1.5 rule), and spend nothing in the second group of superfluous states.[15] Dukakis's alleged 'nineteen state' strategy in 1988 would be a good example of this approach.
3. A third possibility is based roughly on Shaw's findings from the 1988–96 presidential campaigns.[16] Shaw suggests that states are placed by candidates into *five roughly*

[14] Claude S. Colantoni, Terence J. Levensque and Peter C. Ordeshook, 'Campaign Resource Allocation under the Electoral College', *American Political Science Review*, 69 (1975),141–54.

[15] This would not mean that residents in a 'superfluous' state would see no advertisements from a candidate, but rather that they would see only those purchased on the national networks.

[16] Shaw, 'The Effects of TV Advertisements'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

equal-size categories, ranging from their own 'base', to states marginally in their camp, to 'battleground' states, to states marginally in their opponent's camp, and finally to their opponent's 'base'. This rule tells candidates to allocate approximately 40 per cent of their campaign resources to battleground states, 20 per cent to each of the two 'marginal' categories of states, and the remaining 20 per cent to states belonging to the two base categories (devoting slightly more resources to their own base than to their opponent's base).

Thus we have modelled two factors that go into the decision of how much effort to expend trying to win a state's electoral votes. This decision is important to model not only in its own right but also because it should influence the success of any advertising strategy adopted in the state (to which we turn next). In general, the more advertising aired by a candidate in a state, the greater the potential influence of the advertisements, both because a higher proportion of the electorate should be exposed to the candidate's advertising, and (through repeated exposures) because a higher proportion of the electorate should notice and process the message the candidate wants to deliver.

We can approximate the effect of campaign spending by treating the measure of local 'campaign effort' as a weight. The details are reserved for the Appendix. For simplicity, this weight is standardized to have a mean across states of 1.0 for each candidate in every election year, although the range varies between 0 and almost 3. Again, the problem is not that we lack decent data about general indicators of campaign effort (such as spending at the national level), but that we lack (with a few exceptions across this period) decent data about campaign effort at the state level.[17]

Since much of our approach rests on our ability to measure campaign effort, we looked for any data that might show that candidates do indeed make strategic decisions along the lines we posit. Fortunately, we were able to secure data about which states presidential candidates visited in the 1988 to 1996 elections.[18] These data have a direct strategic component that should provide a good test of our approach. Candidates have only limited time in a campaign, just as they have only limited money to spend on advertising, and must use it wisely. While we would not expect spending on local television advertising to correlate perfectly with campaign visits to a state, the two should be strongly related. Thus our attempt to approximate spending with the 'campaign effort' variables should correlate positively with visits – if we are doing things right.

We correlated our six measures of campaign effort (two estimates of the incentives for winning each state, multiplied by three strategies for building winning majorities) with the share of visits by presidential candidates to the states. In the Democratic case, the average correlation is 0.63 and is pretty stable across these various indicators. For Republicans, the correlation drops slightly to an average of 0.55. This decline is attributed to a low correlation between visits and our 'optimistic' measures of effort. In these cases, the correlation drops to about 0.3, the correlations among the other four measures hovering

[17] Implicitly, we are thus assuming that all candidates had approximately the same amount of campaign resources within and across elections. In US presidential elections, particularly after the onset of public financing – the period we happen to be considering – this is a reasonable assumption. But the assumption of equal campaign resources across parties would not be true for all elections, and it would certainly be possible to substitute more differentiated indicators of resources available to different candidates with our general method. In practice, this would simply mean changing the mean value to which the weights are standardized.

[18] We are grateful to Daron Shaw for sharing his data on campaign visits with us.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

in the 0.65 to 0.70 range. Overall, these relationships provide confidence that we are measuring effort with some reasonable degree of success.

Deciding How to Campaign in Each State

There is still the problem of estimating how, as opposed to how much, a candidate will campaign in each state. There are many dimensions along which candidates could (and probably do) vary their advertising strategies across the states. Fortunately, the empirical problem at hand typically limits the relevant dimensions to one or two. To illustrate our approach, we examine the recent controversy over whether attack advertising demobilizes the electorate. This question has currency in the field and stands as one of the more interesting debates in the literature.[19] This focus limits our interest in variation in campaign advertising strategies to the single dimension of how positive or negative the overall tone of the campaign should be.

After reading the literature and consulting our colleagues, we have generated five decision rules that candidates might adopt which vary the proportion of positive or negative themes that would be emphasized in a state, given differential expectations for winning the state. In practice, each of these 'ad-mix strategies' involves increasing or decreasing the proportion of attack advertising from a baseline – the proportion of negative themes coded by Geer for each candidate.[20] This measure is what Finkel and Geer call 'advertising tone'.[21] The baselines can be viewed as reflecting shifting norms as to how much campaign negativism is appropriate and/or effective, along with individual proclivities specific to different candidates and their campaign advisers. (See the Appendix for details.)

1. *Devise one global or national advertising strategy for a campaign, and use it everywhere*. While this is a very simplistic rule, we have included it, in part, as a benchmark to judge the more complicated and realistic strategies. But if all advertising were at the national level – as most of it was, in the early years of television advertising – this is exactly the strategy a candidate would be using.
2. *If the election is close, go negative.* The second rule assumes candidates believe negative advertising is effective, and will increasingly employ it in states where the election is expected to be close.[22] This tactic reflects the current conventional wisdom about negative political advertising.[23]
3. *If behind, attack; if ahead, stay positive*. The third rule assumes candidates adjust their national advertising strategies towards a heavier use of negative themes in states where

[19] Stephen Ansolabehere, Shanto Iyengar, Adam Simon and Nicholas Valentino, 'Does Attack Advertising Demobilize the Electorate?' *American Political Science Review*, 88 (1994), 829–38; Stephen Ansolabehere, Shanto Iyengar and Adam Simon, 'Replicating Experiments Using Aggregate and Survey Data: The Case of Negative Advertising and Turnout', *American Political Science Review*, 93 (1999), 901–9; Finkel and Geer, 'A Spot Check'; Freedman and Goldstein, 'Measuring Media Exposure and the Effects of Negative Campaign Advertisements'; Lau and Pomper, 'Effects of Negative Campaigning on Turnout in U.S. Senate Elections, 1988–98'; Martin P. Wattenberg and Craig L. Brians, 'Negative Campaign Advertising: Demobilizer or Mobilizer?'; Richard R. Lau, Lee Sigelman, Caroline Heldman and Paul Babbitt, 'The Effects of Negative Political Advertising: A Meta-Analytic Assessment', *American Political Science Review*, 93 (1999), 851–75.

[20] Geer, *Attacking Democracy.*

[21] Finkel and Geer, 'A Spot Check'.

[22] Scott Thomas. 'Negative Advertising Theory of Campaign Expenditures', in W. Mark Crain and Robert D. Tollison, eds, *Predicting Politics: Essays in Empirical Public Choice* (Ann Arbor: University of Michigan Press, 19ninety), pp. 245–56.

[23] Lau *et al*., 'The Effects of Negative Political Advertising'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

they believe they are behind, but use relatively more positive themes in states where they believe they are ahead. This strategy assumes that candidates believe that negative advertisements are effective, but risk voter 'backlash'. Candidates who are behind have little to lose, however, and thus are increasingly willing to use negative advertising.[24]

4. *Compare current expectations to traditional performance*.[25] This rule holds that if a candidate expects to do well in a state where the party traditionally does well, the candidate stays positive; but if the candidate expects to do well in a state where the party traditionally does poorly, the candidate goes relatively more negative to try to ensure that victory. In toss-up states, the rule says go positive if your party currently controls the White House, but go negative if you are from the out-party. If a candidate expects to do poorly in a state where the party traditionally does well, the strategy says go negative to try to recapture your base; but if a candidate expects to do poorly in a state where the party traditionally does poorly, this rule says the candidate should stay positive to avoid further alienating voters.
5. *Treat issue and trait themes differently.* The last rule holds that if you expect to do well in a state you go mainly with issue themes, while if you expect to do poorly in a state you disproportionately emphasize trait or character themes.[26] Rather than using the overall difference between the percentage of positive and negative themes, this strategy employs distinct percentage differences calculated separately for issue and trait themes.

Generating the Estimates

In summary, we posit that the strategic context facing presidential contenders is a function of four variables: the perceived likelihood of winning a state, the size of a state (i.e. the number of electoral votes it has), a combination of the first two as candidates decide how much effort to expend campaigning in each, and a strategic decision on how to campaign in each state – which again is conditioned on the perceived likelihood of winning the state. All four of these variables can be considered *decisions* that candidates must make as they attempt to devise a winning campaign strategy. Together, they yield 3 (different measures of expectations) $\times$ 2 (rules about how to consider the number of electoral votes a state has) $\times$ 3 (rules about ways to build a majority in the Electoral College) $\times$ 5 (rules for devising the best ad-mix strategy for a state) = ninety distinct campaign strategies for each candidate. Since the demobilization hypothesis concerns the joint effects of the negativism of all candidates' campaign strategies, we averaged together our estimates of the local strategies (i.e., the decision of how much to attack the opponent) of the Democrat and Republican candidate in each state.

[24] Karen E. James and Paul J. Hansel, 'Negative Advertising: The Malicious Strain of Comparative Advertising', *Journal of Advertising*, 20 (1991); 53–69; James A. Karrh and David H. Halpern, 'Nothing to Lose? Assessing the Impact of Competitive Position on Responses to Negative Political Advertising' (paper presented to the annual conference of the American Academy of Advertising, St Louis, 1997); Stergios Skaperdas and Bernard Grofman, 'Modeling Negative Campaigning', *American Political Science Review*, 89 (1995), 49–61.

[25] This strategy requires a measure of 'traditional' performance that is distinct from 'current expectations'. The difficulty comes in because we have estimated candidates' current expectations from the historical record – the outcomes of the past two presidential elections. We therefore define 'traditional performance' as having a longer timeframe, and measure it with the simple average of how well a party has done in the state over the past four presidential elections.

[26] Richard M. Perloff and Dennis Kinsey, 'Political Advertising as Seen by Consultants and Journalists', *Journal of Advertising Research,* 32 (1992), 53–60.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

We do not pretend that we have captured all of the factors that go into devising campaign strategies, or even that we have modelled the most plausible alternatives of the factors we have identified. We could easily generate more alternatives for each of the four factors discussed above, and perhaps come up with additional factors to be considered. But we are not trying to model every conceivable campaign strategy that might be employed. Our goal is to capture a fairly broad spectrum of what, plausibly, might have been the actual campaign strategies utilized by the major presidential candidates in past US presidential elections.

Our analytic strategy is as follows. We specify a basic equation to predict turnout, including an estimate of the relative negativism of the presidential campaign advertising aired in a state during a particular election and the relative effort the two candidates (combined) put into winning the state. We run this same equation ninety different times, once with every possible combination of the four factors discussed above. This approach yields ninety distinct estimates of the impact of advertising tone on turnout, and ninety distinct estimates of the effect of effort. We do not know which (if any) of those ninety estimates is correct; but with ninety different estimates of a parameter, we can calculate the range of effects advertising tone has on turnout. If the observed range centres around 0, we would conclude that negative advertising neither mobilizes nor demobilizes the electorate. But if the results tend to cluster above or below 0, our results will provide evidence about whether attack advertising stimulates or demobilizes turnout.

AN ILLUSTRATION OF THE METHOD: DOES NEGATIVE ADVERTISING DEMOBILIZE THE ELECTORATE?

Table 1 presents the results of a time-series cross-sectional analysis predicting state-level turnout for the presidential elections between 1980 and 2000.[27] For our measure of turnout, we employ McDonald and Popkin's estimates, which are somewhat higher than "official" levels because they correct for the incarcerated population (and several other small populations) who are not eligible to vote.[28] Our model has eight predictors, including state-level turnout from the previous presidential election, the expected closeness of the election,[29] national turnout levels (to control for year effects, such as Perot's candidacy in 1992), dummy variables indicating whether there was also a gubernatorial or senate election in the state, our estimate of candidate spending in the state, our estimate of the campaign's tone (with more negative campaigns scored high), and finally the interaction of spending and tone, to test the hypothesis that the effect of a campaign's tone is magnified by higher spending levels. The table averages across the ninety different estimates of this model.

All predictors should have a positive sign, with the possible exception of advertising tone (and its interaction with spending). If the demobilization hypothesis is correct, tone

[27] Nathaniel Beck and Jonathan N. Katz, 'Taking Time Seriously: What to Do (And Not to Do) with Time-Series-Cross-Sectional Data in Comparative Politics', *American Political Science Review*, 89 (1995), 634–47; Nathaniel Beck and Jonathan N. Katz, 'Nuisance vs. Substance: Specifying and Estimating Time-Series-Cross-Section Models', *Political Analysis*, 4 (1996), 1–36.

[28] Michael P. McDonald and Samuel L. Popkin, 'The Myth of the Vanishing Voter', *American Political Science Review*, 95 (2001), 963–74.

[29] We devise a measure of projected closeness based on our estimates or projected proportion: Projected Closeness = |Projected Proportion − 0.5|. Thus for each measure of projected proportion of votes in a state, we have a related measure of projected closeness.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

TABLE 1 *Average Results Predicting State-Level Turnout in US Presidential Elections, 1980–2000*

	B	s.e.
Turnout in prior election	0.14**	0.05
Closeness of election	0.03	0.02
National turnout	1.31***	0.09
Governor's election in state	0.40	1.67
Senate election in state	0.39	0.32
Campaign effort (spending)	0.17	0.35
Advertising tone	0.03*	0.01
Tone × spending	− 0.00	0.01
Constant	− 25.98***	6.31
R^2 within states	0.53	
R^2 between years	0.52	
R^2 overall	0.35	

Notes: Results come from a cross-sectional time-series analysis, averaged across ninety different estimates of the same model. $N = 306$. Reported significance levels are based on mean *t*-values. $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

should have a negative sign the way we have coded things, indicating that 'negative' campaigns decrease turnout while 'positive' campaigns increase it. Ansolabehere *et al.*, for example, estimated that turnout in positive senate campaigns in 1992 was about 4 per cent higher than turnout in the most negative senate campaigns that year.[30] Other researchers have found the opposite effect, however,[31] and the literature as a whole offers no clear guidance.[32]

As can be seen in Table 1, all of the control variables have their expected positive sign, although only year differences in national turnout levels, and the lagged measure of state-level turnout, prove to be statistically different from 0, on average. (The expected closeness of the election is also significant with a one-tailed test.) But the main variable of interest is advertising tone, and its effect is overwhelmingly *positive*. On average, more negativism in a campaign's advertising *increased* the probability of voting. Figure 1 displays the ninety estimates of the effects of advertising tone. Two of the ninety are slightly negative, but the remaining eighty-eight estimates all have a positive sign. This result runs counter to the demobilization hypothesis. The average size of the advertising tone coefficient is 0.026, which means that a 10 per cent increase in the negativism of a campaign's advertising increases turnout about 0.25 per cent.

The interaction of advertising tone and spending is, on average, negative but trivial in magnitude. Since spending has been scaled to have a mean of 1.0 across states, we can simply add together to coefficients for advertising tone and the tone by spending interaction

[30] Ansolabehere *et al.*, 'Does Attack Advertising Demobilize the Electorate?'

[31] Finkel and Geer, 'A Spot Check'; Freedman and Goldstein, 'Measuring Media Exposure and the Effects of Negative Campaign Advertisements'; Kenneth M. Goldstein and Paul Freedman, 'Campaign Advertising and Voter Turnout: New Evidence for a Stimulation Effect', *Journal of Politics*, 64 (2002), 721–40; Lau and Pomper, 'Effects of Negative Campaigning on Turnout in U.S. Senate Elections, 1988–98'; Lau and Pomper, *Negative Campaigning*.

[32] Lau *et al.*, 'The Effects of Negative Political Advertising'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

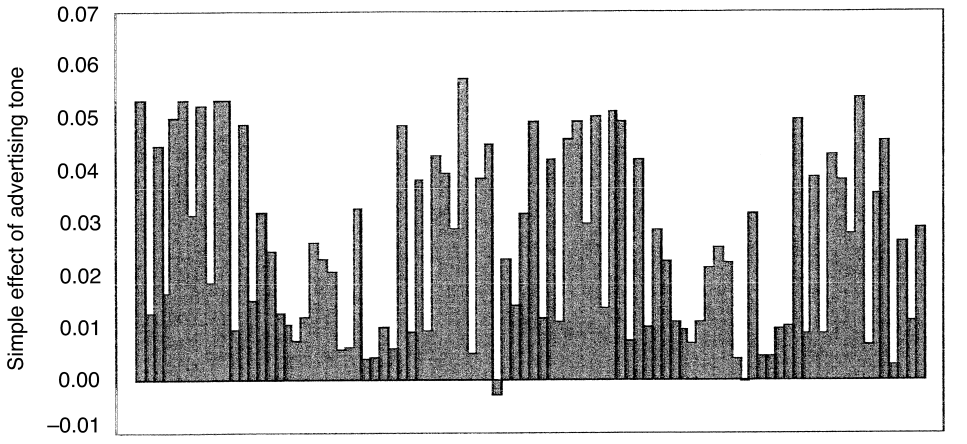


Fig. 1. Ninety estimates of the effect of negative advertising tone on turnout
Note: Figure shows a random display of ninety different estimates of the effect of (negative) advertising tone on turnout. Eighty-eight of the ninety estimates show a positive relationship – that is, negative advertising *stimulates* turnout.

to get an estimated total effect of advertising tone. Every one of these estimated total effects is positive, although the estimates range from barely above 0 to almost 0.055.

If we believed that all of our ninety models were equally likely, we could stop now and conclude that, in all likelihood, the effect of negative advertising in presidential elections was to produce a slight increase in the probability of voting. As indicated in our presentation of the different decision rules which generate the ninety different estimates, however, we do not believe all of these possibilities are equally likely. Fortunately, recent advances in Bayesian Model Averaging allows for more refined hypothesis testing that can move beyond the 'equally likely' assumption, a procedure we illustrate briefly below.

BAYESIAN MODEL AVERAGING

Our goal in creating the ninety different hypothetical campaign strategies examined above was to bracket the plausible range of strategies the presidential candidates might have used over the past several decades. In trying to be fairly comprehensive, however, we knew on going into this exercise that some of the decision rules – and thus some of the resulting campaign strategies – were not very likely. This fact makes using standard statistical procedures to summarize our results problematic, since they assume random sampling (and therefore an equal probability for each of the observations). Of course it is also the case that each of the ninety estimates do not fit the observed data equally well. Bayesian Model Averaging (BMA) can be used to refine estimates in situations like this.[33]

Recall the basic logic of Bayes' well-known theorem, which we have translated into terms that fit our particular problem. First, we assume a sample space *M* of all plausible

[33] Bartels, 'Specification Uncertainty and Model Averaging'; C. Chatfield, 'Model Uncertainty, Data Mining, and Statistical Inference', *Journal of the Royal Statistical Society*, 158 (Series A, 1995), 419–66; Adrian E. Raftery, 'Approximate Bayes Factors and Accounting for Model Uncertainty in Generalized Linear Models', *Biometrika*, 83 (1996), 251–66; Adrian E. Raftery, David Madigan and Jennifer A. Hoeting, 'Bayesian Model Averaging for Linear Regression Models', *Journal of the American Statistical Association*, 92 (1997), 179–91.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

models that could explain a phenomenon of interest. We then partition that space into a set of mutually exclusive models M_i. Next, we observe some relevant data D. Following from the definition of conditional probability, Bayes theorem states that the conditional probability of M_i given the data D is:

$$\text{pr}(M_i|D) = \frac{\text{pr}(D|M_i)*\text{pr}(M_i)}{\sum_{i=1}^{k} \text{pr}(D|M_i)*\text{pr}(M_i)}.$$

In words, the conditional probability of any particular model M_i given the observed data D is equal to the uncertainty in the *data* given the model (that is, how well model M_i fits the data) times the probability of the particular model, divided by the sum of the uncertainty in all of the different models, again given the data, times their probability. The Bayesian approach allows for (indeed, requires) information from outside of the observed data D to be part of the inference process: our prior or subjective probability of each of the different M_is.[34]

BMA is usually employed when there is some uncertainty as to the proper model to examine.[35] In such situations, researchers typically begin with some fixed dataset and examine several different model specifications before they decide on the one to report that, in the light of the data and prior expectations, seems to the researcher to be the most plausible.[36] But then the one 'reported' model is interpreted as if it were the only model examined, when in reality its probability should be discounted according to the number of models actually examined. The Bayesian approach argues that rather than reporting the results of a single model and pretending it is 100 per cent likely, it is better to present a weighted average of all of the models examined, where the weights are determined by the prior probability of each of the models, and their observed fit (i.e., the posterior probability conditional on the observed data). Parameter estimates and standard errors are produced which 'more honestly reflect the observed variation in the results across a range of plausible models'.[37]

Of course we are typically not interested in the overall 'model' itself but rather in some specific parameter(s) of that model, say a regression weight associated with a crucial predictor X_i. The data are comprised of an independent variable y and a matrix of independent variables $\boldsymbol{X}$. The number of columns in $\boldsymbol{X}$ (the number of predictors included in the model) varies according to the different models, but all presumably include the theoretically crucial predictor(s) X_i. The model itself (that is, the other predictors included in the model) is a nuisance variable that we want to get rid of, which is accomplished by integrating it out of the equation.

Our case is somewhat different, but it fits a Bayesian framework equally well. We begin with a problem (explaining turnout) where theory is well developed. There is general consensus about the basic model employed, and we examine the same model over and

[34] The immediate problem of determining the prior and conditional posterior distributions of the infinite number of conceivable models in M is solved by restricting attention to a finite number j of the most plausible models.

[35] See Jennifer A. Hoeting, David Madigan, Adrian E. Raftery and Chris T. Volinsky, 'Bayesian Model Averaging: A Tutorial', *Statistical Science*, 14 (1999): 382–417; Bartels, 'Specification Uncertainty and Model Averaging'. See, also, Larry M. Bartels and John Zaller, 'Election Forecasting Models: A Recount', *PS*, 34 (2001), 9–20; Simon Jackman, 'Estimation and Inference via Bayesian Simulation: An Introduction to Markov Chain Monte Carlo', *American Journal of Political Science*, 44 (2001), 369–98, for several recent political science examples.

[36] Edward E. Leamer, *Specification Searches* (New York: Wiley, 1978).

[37] Bartels, 'Specification Uncertainty and Model Averaging'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

TABLE 2 *Estimated Prior Probabilities of Alternative Campaign Decisions*

Estimating candidates' beliefs about their chances in a state	
Strictly historical method	0.36
Optimistic bias	0.29
Polling approximation	0.35
Incentives	
Proportional to number of electoral college votes	0.44
Disproportional, 1.5 rule	0.56
Strategy for building a winning majority	
All votes equally valuable	0.07
Necessary/Superfluous categories of states	0.28
5-fold categorization of states	0.65
Advertising-mix strategies	
Use one global strategy everywhere	0.14
If election close, go negative	0.19
If behind, attack; if ahead, be positive	0.32
Compare to traditional performance	0.15
Treat issue and trait themes differently	0.20

Note: The prior probability of each unique campaign strategy is the product of the probabilities associated with each alternative decision rule.

over again. Hence $\boldsymbol{X}$ always has the same number of columns. Instead, there is some uncertainty in the data, at least in the variance or specific values of our two crucial, theoretically interesting independent variables X_i and X_j (campaign effort and advertising tone). We can still think of this as uncertainty about which model M_i is appropriate, remembering that the uncertainty involves the values of particular independent variables rather than which independent variables belong in the model. Clearly, *a priori* some of those values (i.e., some of the campaign decision rules) seem more likely than others; and *a posteriori*, some of them do a better job of explaining turnout than others. Once again the model – that is, the particular values that X_i takes on – is a nuisance variable we want to get rid of, concentrating on the values of the parameters B_i and B_j associated with those two variables. Thus we can use a Bayesian approach to incorporate prior beliefs about the relative likelihood of the difference decision rules more accurately to summarize our findings.

We began by estimating the *a priori* probability of each of the different alternative rules for each of the four 'decisions' that went into devising the different campaign strategies. To increase the generality of the prior probability judgements, we also asked ten experts on American political campaigns to make similar estimates. These experts were a mixture of academics and campaign professionals. The mean judgement of our dozen experts is reported in Table 2. The prior probability of each of the ninety different models was calculated by multiplying these four probability judgements together. Notice that the probability of the different alternatives for each decision rule must sum to 1.0, so that the sum of the prior probabilities associated with the ninety different models will also be 1.0.

To calculate the mean and variance of the posterior distribution of the parameters of interest, one must first calculate the conditional posterior model probabilities ($PROB_{Post}$) of each of our ninety different models. As described by Bartels,[38] in the case of ordinary

[38] Bartels, 'Specification Uncertainty and Model Averaging'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

TABLE 3 *Bayesian Model Averaging Adjusted Results*

	B	s.e.	95% C.I.
Campaign effort (spending)	0.19	0.32	− 0.44 to 0.81
Advertising tone	0.06	0.01	0.03 to 0.08
Tone × spending	− 0.01	0.01	− 0.03 to 0.01

Note: Results are based on all ninety original estimates summarized in Table 1.

least squares (OLS) regression this is a fairly simple two-step procedure that involves first calculating a 'Bayes factor' (*BAYES*) for each model, which translates prior model odds into posterior model odds, and then applying the formula:

$$PROB_{Post(i)} = \frac{BAYES_i \times PROB_{Prior(i)}}{\Sigma BAYES_i \times PROB_{Prior(i)}}.$$

We provide a few more details in the Appendix, and refer the interested reader to Bartels.[39] The mean of the unconditional posterior distribution for any parameter B_j of the multiple regression is then simply

$$\text{MEAN}_{\text{Bj}} = \sum_i PROB_{Post(i)} \times B_{ij},$$

and its variance is

$$\text{VAR}_{\text{Bj}} = \sum_i PROB_{Post(i)} \times \text{VAR}_{\text{Bij}} + \sum_i PROB_{Post(i)} \times (B_{ij} - \text{MEAN}_{\text{Bj}})^2.$$

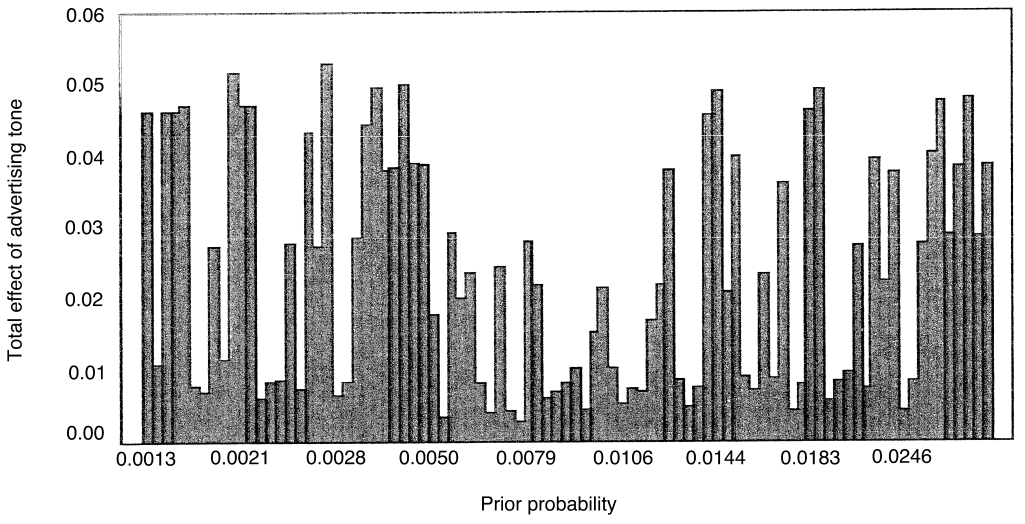
Applying these two formulae, Table 3 presents a summary of the BMA results for the parameter of interest here – the estimated coefficients for advertising effort, advertising tone and their interaction.

As can be seen in Table 3, the BMA provides more powerful estimates of the effect of all three of our crucial variables on turnout. The estimate for the effect of advertising tone is twice as large as the simple, unweighted average of all ninety coefficients estimated, and the standard error is only slightly greater. The end result is a much more powerful test of a hypothesis, and thus gives more confidence in our final estimates. In this particular example no estimate goes from being nonsignificant to being significantly different from 0 (or vice versa), but we have other instances in our own research where greater power resulting from BMA has done just that. Unfortunately, BMA guarantees neither larger mean estimates nor more powerful tests of hypotheses. In our case, the adjusted mean becomes larger because the lowest estimates of the effect of Advertising Tone had, on average, lower estimated *prior* probabilities, and because they fit the data less well, resulting in lower estimated *posterior* probabilities.

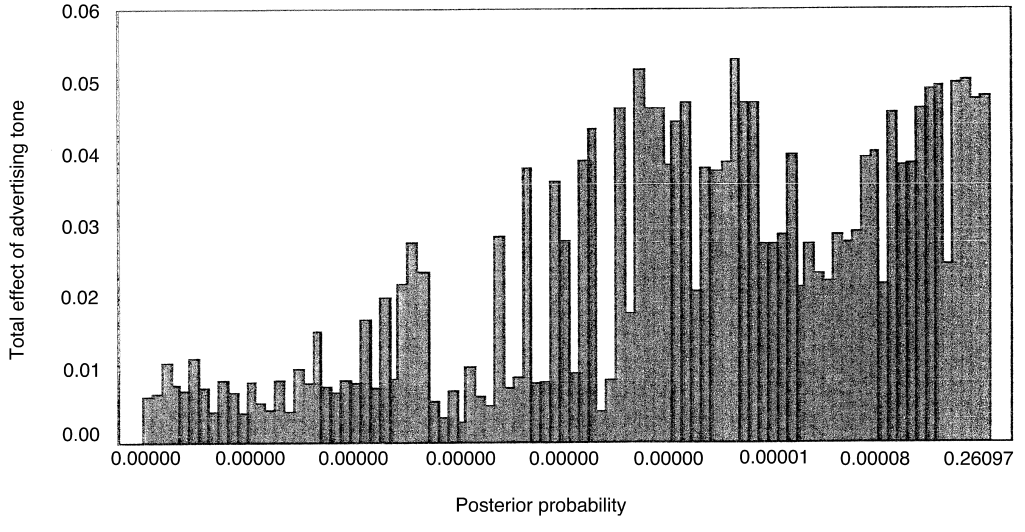
Figure 2 provides a graphical account of this process for the ninety estimates of the *total* effect of Advertising Tone (that is, the 'main effect' of tone, plus its interaction with spending). The figure sorts the ninety estimates by both the prior probabilities assigned to each of them by our panel of experts (the top half of the figure) and their posterior probability (bottom half), which is based on both the prior probabilities and the observed

[39] Bartels, 'Specification Uncertainty and Model Averaging'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press



(a) Sorted by prior probability



(b) Sorted by posterior probability

Fig. 2. Effect of Bayesian Model Averaging on estimates of the total effect of negative advertising tone on turnout

fits. One can see why the estimated coefficients becomes stronger once they are weighted by their posterior probability. The largest positive coefficients are now mostly on the right side of the figure (indicating they are receiving the highest weights), while the smaller coefficients are all on the left half of the figure (because they receive the smallest weights).

To summarize the substantive findings from this empirical example of our method, after considering the prior expectations of a panel of experts as to the relative likelihood of the different campaign strategies considered here and the observed fit of the different models to actual data, we should be very confident that the true effect of negative advertising on

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

turnout in presidential elections over the past two decades was to provide a modest positive (i.e., mobilizing) boost to the probability of voting. For every one point increase in the combined negativism of presidential candidates' campaign advertising, turnout increases by about 0.03 to 0.08 of 1 per cent – employing the full 95 per cent confidence interval (which would be more consistent with the thrust of a Bayesian approach). In a state the size of Florida, to pick a not totally random example, this would mean an increase of somewhere between 954 to 4,293 voters for every *one* point increase in the combined negativism of the two major candidates' campaigns.

Robustness of Findings

To present and illustrate our method in the most general way possible, we have pretended that we have no data on the differential campaign effort or advertising strategy employed in the different states by any of the candidates across the entire 1980–2000 period. This is not true, however. Goldstein has gathered data on the actual campaign advertising shown in the seventy-five largest media markets in the United States by the different candidates during the 2000 presidential election, and these data would provide very precise measures of what we are trying to estimate, for about forty states in this one most recent presidential election. Shaw has presented data on the number of 'ad-buys' (i.e., the number of advertisements shown) by the Democratic and Republican candidates in the fifty states between 1988 and 1996.[40] A true Bayesian analysis – and just plain common sense – would take advantage of as much information as possible. Thus, if our interest were primarily on the effects of negative advertising on turnout rather than the methodological technique itself, and we were willing to limit attention to either the 2000 election or the 1988–1996 elections, we should take advantage of this additional information in generating our estimates. Our purpose here is to present the general technique, but we can use this additional data to test the robustness of our findings.

We will employ Shaw's data, and thus limit attention to the 1988, 1992 and 1996 presidential elections. Shaw has no data on the nature of the campaigning in the different states, but he has very good data on the campaign effort expended by the different candidates in each state. If we treat Shaw's data on what we have been calling campaign effort as certain, we can do away with the six different estimates we had for that concept. We are still left with uncertainty about the candidates' expectations of winning each state, and the particular advertising strategy employed, but this reduces our original ninety estimates of the effect of advertising tone on turnout to fifteen (three methods for estimating expectations multiplied by five methods for devising an advertising strategy). Presumably these fifteen estimates should be much closer to the true figure for the 1988–96 period, because they are based on much less uncertainty. The question we ask here is, how well did our method do in coming up with those same estimates, but without the precise evidence on campaign effort in each state? This will provide a broad test of the robustness of our procedures.

Table 4 presents the relevant results. The pattern of results is identical to our original estimates presented above in Table 3, although the effect of advertising tone is a bit stronger with our more refined estimates. Whether based on all ninety of our original estimates or the new restricted set of fifteen estimates incorporating more precise information on

[40] Shaw, 'The Effects of TV Advertisements and Candidate Appearances on Statewide Presidential Votes, 1988–1996'.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

TABLE 4 *Robustness Test: BMA Adjusted Results, Assuming Campaign Effort Known with Certainty*

	B	s.e.	95% C.I.
Campaign effort (spending)	0.27	0.47	− 0.75 to 1.28
Advertising tone	0.24	0.11	0.01 to 0.47
Tone × spending	− 0.14	0.07	− 0.29 to 0.01

Note: Results are based on fifteen estimates of turnout in the 1988–96 US presidential elections, where the campaign effort within states is based on Shaw's (1999) ad-buy data.

campaign effort, however, the substantive conclusion would be the same: negative campaigning has a small *mobilizing* effect on turnout. More importantly for present purposes, these new data provide strong evidence that our method is doing exactly what we hoped it would – capturing the *range* of campaign activities in which candidates engage.

DISCUSSION

When we first presented this method to some of our colleagues, one told us (with a smile on his face) that we had found a way to do what 'we theorists always knew you empiricists wanted to do – just make-up the data'. We will plead guilty, but only to a somewhat reduced charge. Our method is designed for those situations where we are missing crucial pieces of information. The point of this article is that even in those bleak situations, we do not often work in a vacuum. In the case of political advertising, we usually do not know which advertisements were shown in what locations, nor how often they were shown, but we do have a rough idea of the overall campaign candidates wanted to run from the themes in the advertisements they produced. From such an empirically grounded starting point, *theory* tells us – if not in detail, then at least in generalities – how overall campaign strategy may have varied across distinct electoral contexts. That knowledge provides the leverage to generate a *range* of plausible estimates. The better the theory one uses, the better the estimates. In the absence of a good theoretical understanding of the subject in question, our method cannot be usefully applied. Thus, we plead guilty not to making up the data so much as making up *variance* in the data.

While our approach could be refined even further, it is important to stress the potential value of the approach as already developed in this article. We have shown that we can generate precise enough findings to test important hypotheses. Yet we have only scratched the surface of data and questions that might be examined in future research. One could, for example, use the approach to see how campaign strategies might affect the *outcome* of elections, not just the level of participation. It should also be possible to see if particular themes in candidate's advertising 'prime' the electorate to think more about such considerations when voting. And it would be possible to test for differences in the effectiveness of certain campaign strategies over time. We also have a wealth of aggregate data from elections dating back to the 1800s. Our approach provides an opportunity to examine these campaigns. It would be possible to estimate a campaign's tone from newspaper coverage, for example, which might be used instead of television advertising

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

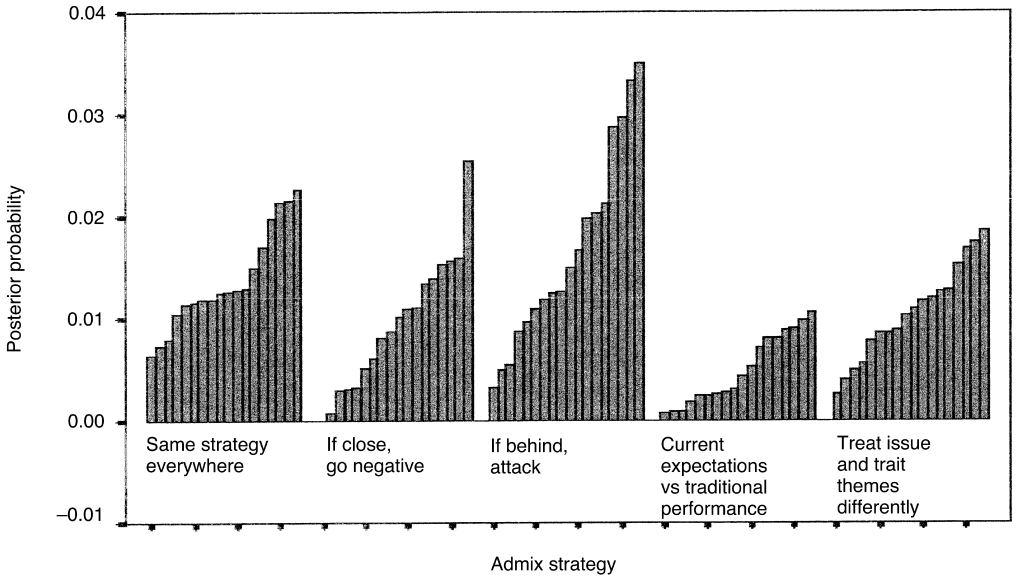


Fig. 3. Posterior probabilities associated with different ad-mix strategies
Note: Figure shows a display of the ninety different posterior probabilities, sorted by the five ad-mix strategies.

for elections prior to 1952. Paid advertising in newspapers could provide further evidence. The availability of good aggregate and survey data extends beyond our borders, of course, making it possible to undertake some cross-national research.

Our results also offer the opportunity to assess the likelihood that each of the different advertising strategies was actually employed by candidates over the six presidential elections considered here. If the reader will indulge our illustrating the last point, Figure 3 plots not the different estimated effects of the campaign, but their *posterior probability*, sorted by the five different hypothesized advertising strategies.[41] As can be seen in the figure, overall two of our proposed strategies, 'use the same strategy everywhere' (with a total posterior probability of 0.247) and 'if behind, attack' (total posterior probability of 0.301) appear to be the most likely, while one of the strategies, 'compare current expectations to traditional performance' (total posterior probability of 0.090) resulted in noticeably lower posterior probabilities than the other four. The first two groups differ significantly from the last, according to the Tukey *post hoc* test.

This method does not, nor should it, stand alone. The advancement of any field requires the development of new methods that open up windows of discovery. We need to be innovative as we seek to learn more about how campaigns or other important institutions work, especially when we are trying to fill in important gaps in our empirical understanding. The bottom line is that we can use our knowledge of a particular subject as leverage to secure additional insights about the topic. This general approach can be applied to topics where we have a strong theoretical understanding of the topic, but often

[41] The probabilities have been subjected to a natural log transformation to reduce heteroscedasticity between groups, but then rescaled to have the usual properties of a probability distribution – all values between 0 and 1, and the sum of all values equal to 1. 'Blank' cases have been inserted into the figures to help make the groups more distinct.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

lack key pieces of information about the detailed empirical workings of it. It is often hard to go back and fill in the empirical record of events that have already taken place. Our approach, at least, has the benefit of not just ignoring this all too common a problem.

APPENDIX

Baseline Advertising Tone Data

The data in Table A reflect the tone of the appeals appearing in each presidential candidate's advertisements over the past twenty years, as coded by Geer. If every appeal were positive, the figure would be + 100; if every appeal were negative, the figure would be − 100. The predominance of positive numbers in the table reflects the fact that, over the years, the appeals of both the Democratic and Republican candidates have been largely positive. We reverse these codes for analysis, because most researchers are concerned with how *negative* a campaign is. All but the fifth advertising mix strategy considered here employed figures from the two 'Overall' columns of the table. That fifth strategy, which distinguishes between issue and trait appeals, utilized figures from those respective columns in the table.

TABLE A *Advertising Tone by US Presidential Candidates, 1980–2000*

	Democratic Candidates			*Republican Candidates*		
Year	Issue appeals	Trait appeals	Overall	Issue appeals	Trait appeals	Overall
1980	56.6	30.2	51.4	9.3	32.0	25.6
1984	− 18.2	62.5	27.3	41.1	55.9	66.0
1988	− 2.2	18.4	24.6	− 7.2	51.4	41.8
1992	11.4	13.0	27.2	13.0	− 50.0	3.2
1996	0.0	37.6	25.1	− 34.8	5.7	− 5.0
2000	5.8	20.4	12.8	30.1	− 13.9	30.8

Note: Original data were coded by Geer (see John G. Geer, *Attacking Democracy: A Defense of Negative Advertising in Presidential Campaigns* (Chicago: University of Chicago Press, 2006)). They represent the number of positive appeals minus the number of negative appeals (in each category), divided by the total number of appeals, times 100. Hence the higher the table entry, the more *positive* the mix of campaign appeals. These data represent the baseline or 'national' advertising strategies utilized by each presidential candidate in each election year. They were turned into proportions and reversed for analysis, so that high values of our measure of advertising tone reflect more *negative* advertising. Note that Issue Appeals and Trait Appeals do not provide an exhaustive breakdown to total appeals. Some appeals were categorized as 'Values', which proved to be almost universally positive as applied to the political advertisements.

Calculating 'Campaign Effort' Weights

The weight reflecting how much total 'effort' (or resources) each candidate was presumed to expend in each state was computed as the product of the 'incentives' they were presumed to associate with winning each state, and a number associated with their presumed strategy for winning an electoral majority. Let us call that latter number the WM (for 'winning majority') weight.

Incentives were set to equal $EC_j/270$ (where EC_j is the number of electoral votes associated with winning state j) if incentives for winning a state are assumed to be proportional to the number of electoral votes; or $(EC_j/270)^{1.5}$ if incentives are assumed to be disproportionately biased in favour of the larger states. In

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

either case, the weight is rescaled to give it a mean across states of 1.0. (There are fifty-one 'states' in our calculations, which includes the District of Columbia.)

There were three alternative decision rules for building a majority in the Electoral College. Each involved a different method for determining how to devise a winning majority, and thus a different WM weight.

—The first rule assumes that candidates believe that all votes are equally valuable, and they therefore want to campaign locally in every state. Thus the WM equals 1.0 for every state (including Washington, D.C.) and sums to fifty-one per candidate within election year.

—The second rule says candidates put states into two categories, those that are most likely to produce the required 270 Electoral College votes, and the remaining 'superfluous' states.[42] To determine these two categories, all states were ranked (for each candidate in each election year) according to the proportion of votes the candidate was expected to win in the state – which in turn was determined by one of the three methods for estimating candidates' beliefs about their chances of winning each state. Beginning with the most likely state to be in each candidate's column, and preceding down the list in order, the number of Electoral College votes associated with each state was summed until the total equalled (or exceeded) 270. All of these 'necessary' states (plus the next state on the list, to provide some margin of safety) were assigned a WM of $51/k$, where k is the number of states in the 'necessary' category. All of the remaining 'superfluous' states were assigned a WM of 0. Within election year, the weights sum to 51.

—The final strategy says candidates place states into five roughly equal size categories, those in their own base, those marginally in their camp, battleground states, those marginally in their opponent's camp, and their opponent's base. For each candidate we assigned states to these five different categories, using the same ranking described above. Then we assigned WM weights of 0.6, 1, 2, 1 and 0.4 to the five categories of states. Again, the sum for each candidate within election year is 51.

Local advertising effort then equals Incentives × WM for each candidate, and the variable Campaign Effort employed in the analyses was the average of these two values across the two candidates. Campaign Effort thus has a mean of 1.0 across all states within every election year.

Adjusting Baseline Advertising Tones

We proposed five different decision rules for devising an Advertisement Tone to be employed in each particular state. In every case, the rule is based on adjustments from the baseline data reported in Table A.

—*Devise one global advertising strategy and use it everywhere* involves no adjustment from the baseline. The advertising tone is assumed to be identical across states, but does vary across candidates and across election years.

—*If the election is close, go negative* calls for an adjustment of the baseline tone according to the expected closeness of the election in the state. Operationally, if the two candidates were expected to be within 2 points of each other (e.g., 50.9–49.1), advertising tone in the state was presumed to be 50 points more negative than the baseline advertising tone. When the difference between candidates was expected to be 2–4 points, advertising tone was 25 points more negative than baseline; states with an expected difference of 4–6 points employed the baseline figure; states with an expected difference of 6–8 points were presumed to be 25 points less negative than the baseline; and finally states where the difference between candidates was expected to be 8 points or more (such as 46–54), tone was presumed to be 50 points less negative than baseline.[43]

—*If behind, attack; if ahead, go positive* presumed that candidates would go 33 points more negative than baseline whenever they expected to lose a state, go 33 points less negative than baseline when they expected to win by 5 or more points, and otherwise employed a baseline level of negativism in their statewide advertisements.

[42] Note that candidates may not actually believe (according to our estimates) that they are going to receive a majority of votes in all of the 'necessary' states. Assuming they nonetheless want to win the election, they should still concentrate their efforts on the states which are most likely to produce the required 270 Electoral College votes, according to this strategy.

[43] In a few election years, the more extreme 50 point adjustments resulted in predictions of the percentage of negative advertising of less than 0 or greater than 100 per cent. In such cases presumed values were reduced to their logical extremes.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press

—*Compare current expectations to traditional performance* was the most complicated strategy. 'Traditional performance' was the mean of how well the party's candidate did in the state over the past four election cycles, while 'current expectations' was based on various combinations of the results from the past two elections only, as determined by the three different measures of candidates' beliefs about winning the state. If both traditional performance and current expectations predicted a victory by more than 3 points in the state, the strategy says use an advertising tone 33 points less negative than baseline; if traditional performance is to lose in the state, however, but current expectations are for a victory of more than 3 points, the strategy is employ an advertising tone 33 points more negative than baseline. If tradition dictates a loss in a state of greater than 3 points and current expectations are following that tradition, the strategy says be 33 points less negative than baseline; but if a candidate traditionally wins a state but currently expects to lose by more than 3 points, he should be 33 points more negative than baseline. Finally, whenever current expectations are for a very competitive election in the state (3 points or less in either direction) and your party currently holds the presidency, be 33 points less negative than baseline, but if you are the out-party, be 33 points more negative than baseline.

—Finally, *treat issues and traits differently* employed data from the 'Issue Appeals' and 'Trait Appeals' columns of Table A. This strategy says to use proportionately more issue appeals, the better you expect to do in a state. Operationally, we assumed the candidate used only issue-based appeals when they expected to win a state by more than 7 points; used 80 per cent issue appeals and 20 per cent trait appeals when they expected to win by between 3 and 7 points in a state; used 60 per cent issue appeal and 40 per cent trait appeals when the election was expected to be within 3 points; use 40 per cent issue appeals and 60 per cent trait appeals whenever they expected to lose by 3 to 7 points; and used 20 per cent issue appeals and 80 per cent trait appeals whenever they expected to lose by more than 7 points. These figures reflect the general predominance of issue-based appeals over trait-based appeals in the data. Notice that this strategy results in very different levels of negativism across election years, even for states that fall into the same expectation category.

https://doi.org/10.1017/S0007123406000159 Published online by Cambridge University Press