

Article: “On Forecasting the Presidential Vote”
Author: Christopher Wlezien
Issue: Mar. 2001
Journal: *PS: Political Science & Politics*



This journal is published by the American Political Science Association. All rights reserved.

APSA is posting this article for public view on its website. APSA journals are fully accessible to APSA members and institutional subscribers. To view the table of contents or abstracts from this or any of APSA's journals, please go to the website of our publisher Cambridge University Press (<http://journals.cambridge.org>).

This article may only be used for personal, non-commercial, or limited classroom use. For permissions for all other uses of this article should be directed to Cambridge University Press at permissions@cup.org.

On Forecasting the Presidential Vote*

"He's breaking every political scientist's heart."

Mark Shields, commenting on the Gore campaign for The Newshour with Jim Lehrer on election eve, November 6, 2000

It did not look good for Gore leading up to Election Day. Five days before the election, not a single national poll had Gore ahead. Indeed, the polls at that time showed Bush leading by four percentage points on average (see Table 1). Bush was campaigning in California, and the rumors were that he already had shut down his own internal polling. Even on election eve, only two polls (CBS and Zogby/MSNBC) showed Gore ahead. In these final preelection polls, Bush was leading by 1.7 percentage points on average, and with a pooled N of over 10,000 respondents! At midnight, as Election Day began, the Iowa Political Stock Market vote share market indicated a commanding Bush lead of

4.5 points and the winner-take-all market put the effective probability of a Bush victory at 71%. Some political commentators were writing Gore's

epitaph (Kurtz 2000). His one hope, it appeared, was to win the Electoral College vote despite losing the popular vote.

As everyone now knows, Gore won the popular vote. As of this writing (December 2000), he has received 50.2% of the two-party vote. Clearly, the result surprised most political observers. Except for the few pollsters mentioned above and members of the Gore campaign, it seemed that the only people expecting Gore to win just prior to the election were a handful of political scientists (see, e.g., Bernstein 2000; Stout 2000). A couple of these election forecasters were even the objects of one of Rush Limbaugh's Election Day monologues. They were "liberal Democrats," presumably out of touch with the electoral forces at work. They could not see the obvious, impending Bush victory. As it turns out, these political scientists were right—sort of.

A Brief Retrospective on Forecasting the Presidential Vote

Forecasting the presidential election on election eve is relatively easy, at least since the advent of polling. That is, the outcome becomes clear(er) in the polls as the election cycle evolves (Campbell and Wink 1990; Erikson and Wlezien 1996). By Election Day, the polls give a pretty good reading of the result, though even the final preelection polls are imperfect, as in 2000. As one steps back from the election into the summer and spring, however, the polls offer much less information. Forecasting during this period is more difficult.

For a number of elections, political scientists have been forecasting the presidential vote early in the year, at least before the fall campaign begins in earnest (see Abramowitz 1988; Lewis-Beck and Rice 1984, 1992).¹ In general, the forecasters build on what we have learned from models that account for election outcomes after the fact (e.g., Erikson 1989; Hibbs 1987; Tufte 1978). That is, they rely either explicitly or implicitly on a "referendum" model of sorts, in which voters—those in the middle—vote either to stay the course or to change based on the performance of the incumbent administration. The point is simple: Good performance by the sitting president advantages the candidate of the president's party *and* disadvantages the candidate of the out-party by making the case for change more difficult to make.

All of the forecasters focus on the incumbent party candidate's share of the two-party presidential vote, ignoring other candidates.² As for the predictors themselves, each of the forecasting models, with the exception of Norpoth's most recent model, includes some measure of economic performance or perceptions, typically for the election year itself, though some take a longer view (Wlezien and Erikson 1996). Each also includes some measure of public opinion, and here the differences among the models are greater. Reflecting the referendum perspective, most focus on judgments of the sitting president and use presidential approval itself, following Brody and Sigelman (1983). Others

by
Christopher Wlezien,
University of Houston

use indicators of candidate approval, such as performance in early primaries (Norpoth 1996a) or actual trial-heat poll results from around Labor Day (Campbell 1996). Some of the models also include a

differ in important ways, the forecasting models have a lot in common.³ This is evident in the pattern of forecasts over the years.

By 1992, the field included four of the currently active forecasters. In July or August before the election, they provided their forecasts, which were published well after the fact in *The Political Methodologist* in the spring of 1994. These forecasts are shown in Table 2. They were, for the most part, fairly similar and quite accurate. Three of the four forecasters predicted a Clinton victory and each was within 0.6 of a percentage point of his actual share. Lewis-Beck and Rice predicted Bush would win, though this was based on their new-and-improved model. Using their original, more basic model, they predicted a Clinton victory and came within one percentage point of the actual vote. All told, 1992 was a pretty good year for political science forecasters. Campbell (1993) even noted that “weather forecasters should be so accurate.”

In 1996, a few others joined the fold. As in 1992, the group provided forecasts in July or August. The seven forecasts then were published in *American Politics Quarterly* (vol. 24, 4[October]), this time in advance of the election. These are also shown in table 2. The predictions for 1996 were slightly less similar than in 1992, and also less accurate on balance. What is most important is that all of the forecasters predicted a Clinton victory and that all of the forecasts were too high, by about two percentage points on average. Clinton just did not do as well as expected. Nevertheless, the models did perform better than the final preelection polls, which overestimated Clinton’s vote share by about three percentage points. Perhaps, to paraphrase Campbell, “pollsters should be so accurate.” The 2000 election is a different story.

TABLE 1.
The Bush Lead in Preelection Polls, 2000 (%)

Polling Organization	Five Days before Election Day	One Day before Election Day ^a
Voter.com/Battleground	9	5
Hotline	7	–
Marist College	5	–
ABC	3	3
<i>The Washington Post</i>	3	3
NBC/ <i>Wall Street Journal</i>	3 ^b	3
CNN/ <i>USA Today</i> /Gallup	6	2
Pew Research Center	–	2
<i>Investor's Business Daily</i> / <i>Christian Science Monitor</i> /TIPP	5	2
International Communications Research	–	2
<i>Newsweek</i>	2	–
Fox/ <i>Opinion Dynamics</i>	0	–
Harris	–	0
CBS	2 ^b	–1
Reuters/MSNBC	3	–2
Average Bush Lead	4.0	1.7

Source: PollingReport.com (www.pollingreport.com).

^aPolls released one or two days before Election Day.

^bPolls released four days before Election Day.

TABLE 2.
Presidential Vote Forecasts, 1992–96: Incumbent Party Percent of the Two-Party Vote

Forecasters	1992 ^a	1996 ^b
Abramowitz	46.7	56.8
Campbell	47.1	58.1
Holbrook	–	57.2
Lewis-Beck and Rice	51.5	–
Lewis-Beck and Tien	–	54.8
Lockerbie	–	57.6
Norpoth	–	57.1
Erikson and Wlezien	46.8	–
Wlezien and Erikson	–	56.0
Actual Vote	46.5	54.7

^a*The Political Methodologist*, spring 1994.

^b*American Politics Quarterly*, October 1996.

third variable, such as the amount of time a particular party has controlled the White House (Abramowitz 1988; Holbrook 1996; Lockerbie 2000). The point of this description is not to detail all of the differences among the models but to show that, although they

the forecasters to supply early predictions of the 2000 presidential vote (Kaiser 2000). Those that could, given their models, did. The forecasts are shown in Table 3. The pattern is quite similar to that for 1996. All of the forecasters predicted a Gore victory and the

specific forecasts differed quite a lot, ranging from 53.5% or so to 59.6% of the two-party vote. The average predicted Gore vote share was 56.3%. At the time, Bush was leading in the polls by six percentage points on balance, by as much as nine points in some polls. In contrast with 1996 and even 1992, the forecasts were directly at odds with the polls. Still, it was only May, and the polls at that point in the election cycle tell very little about the ultimate outcome (Campbell and Wink 1990).

At the American Political Science Association Annual Meeting in August, the same scholars plus the others provided what one might call their “final” forecasts. Again, all predicted a Gore victory, and the average predicted vote share was 56.0%, only slightly less than in May. The variance of the forecasts was larger, however, and the specific predictions ranged from 52.8% to 60.3%. Still, the statistical probabilities of victory associated with the forecasts all were over 90%—in effect, if the election were run 100 times, Gore would win 90 times or more. There is no doubting the consensus among the forecasters: The election was Gore’s to lose. In the wake of the Democratic convention and Gore’s subsequent rise in the polls, this expectation seemed more reasonable.

Gore did win the popular vote. In one sense, then, the forecasters were right yet again. In another sense, they were not quite right, since Gore almost lost. Although some were closer than others, all of the forecasts were too high, by an average of 5.8 percentage points using the August numbers. This is a very large “error” given the past performance of the models, though there is some precedent for it, at least given the performance of the Wlezien-Erikson model (1996; forthcoming). This can be seen in Table 4, which contains our out-of-sample forecasts—what we would “forecast” now—for each of the 13 elections between 1952 and 2000. The model includes two variables: (1) a measure of the weighted average growth in leading economic indicators from the first quarter of an election cycle through the thirteenth quarter, the first quarter of the election year; and (2) the most current quarterly reading of presidential approval. The model can be estimated at any point throughout the election

year, but I focus here on the quarter 13 model, which uses information available in the early spring.

As is clear in the table, the model works very well in some years but less well in others. The mean absolute error for the out-of-sample forecasts is 2.3 percentage points over the 13 elections (2.0% when excluding

TABLE 3.
Presidential Vote Forecasts, 2000: Gore Percent of the Two-Party Vote

Forecasters	May ^a	August ^b
Campbell	–	52.8
Abramowitz	53–54	53.2
Norpoth	–	55.0
Wlezien and Erikson	56.1	55.2
Lewis-Beck and Tien	56.2	55.4
Holbrook	59.6	60.3
Lockerbie	–	60.3 ^c
Mean Forecast	56.3	56.0

^aKaiser (2000).

^bPresented at the 2000 Annual Meeting of the American Political Science Association in Washington, DC.

^cThis forecast differs from the one (52.9%) presented in Washington because of a calculation error.

TABLE 4.
Out-of-Sample Forecasts Using the Wlezien-Erikson Model, 1952–2000

Year	Actual	Incumbent Party Vote (%)		Error
		Predicted		
1952	44.6	45.2		–0.4
1956	57.8	57.1		0.7
1960	49.9	52.5		–2.6
1964	61.3	59.8		1.5
1968	49.6	51.7		–2.1
1972	61.8	54.2		7.6
1976	48.9	50.8		–1.9
1980	44.7	39.8		4.9
1984	59.2	60.0		–0.8
1988	53.9	53.3		0.6
1992	46.5	46.6		–0.1
1996	54.7	53.9		0.8
2000	50.2	55.8		–5.6
Mean Absolute Error				2.3

Note: The out-of-sample forecast for each year is based on an estimated model that excludes the particular election.

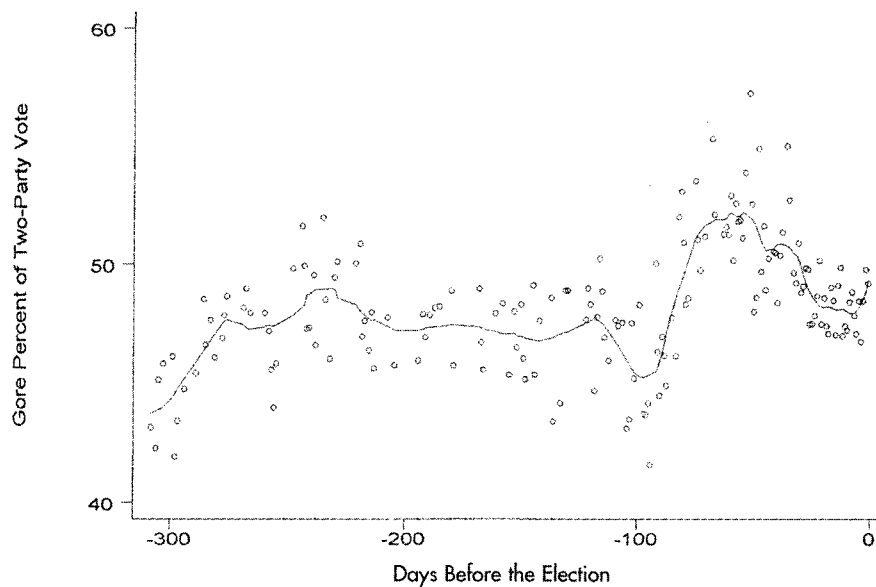
2000). The error for 2000 (5.6%) is, thus, well above average, but by no means unique. In 1972, the out-of-sample error (7.6%) is actually larger. In that year, Nixon did much better than the model predicts. In another year, 1980, the error (4.9%) is just slightly less

than in 2000. The point is that the model has been “wrong” before. Uncertainty surrounds the forecasts—we’re not predicting death and taxes after all. In effect, we forecast a distribution of outcomes. We knew this going into the 2000 election. In August, we estimated Gore’s probability of winning the popular vote to be just above 90%.⁴ The actual election outcome was thus within the 95% (one-tailed) confidence interval. This said, we nevertheless were off by a lot in 2000. What happened? Where did our model and the others’ go wrong?

Observations on the 2000 Election

Some answers come from the forecasting models themselves. After all, their performances in 2000 differed quite a lot. What does this tell us? There are at least two lessons. First, regarding the economy, things may have seemed better than they really were.

FIGURE 1
Trial-Heat Polls Aggregated by Date, 2000



Models relying on objective measures of economic growth (Abramowitz’s, Campbell’s, Wlezien and Erikson’s) outperformed those using subjective measures (Holbrook’s, Lockerbie’s), and the one that uses both (Lewis-Beck and Tien’s) was in between (see Table 3). Of course, while this may explain why some forecasts were worse than others, it does not explain why all of the forecasts were too high. Second, it may have been “time for a change,” as Abramowitz has argued. The Democrats had held the White House for two terms, so it may be that some voters wanted to change course. Although the exact underpinnings are unclear, there is a certain intuitive appeal to this line of thinking, and Abramowitz’s model performed better than all others in 2000, excepting Campbell’s. It only takes us part of the way, however. Witness

Abramowitz’s forecast, which was off by three percentage points.

There are other explanations from “outside” the models, some of which have become usual suspects in the postelection commentary. First, there was the Nader factor. Nader clearly “took” votes from Gore. The now-familiar analysis of exit polls suggests that the net effect was less than one percentage point, however. Of the 2.7% respondents who voted for Nader, 47% said they would have voted for Gore, 21% for Bush, and 30% would have stayed home.⁵ Second, Gore was not an incumbent. Perhaps, therefore, he did not get full credit for the country’s good times. It is hard to gainsay the point, though it is important to note that incumbency does not seem to matter very much historically, which explains why it is not in the models to begin with (Wlezien and Erikson forthcoming). Either way, it is difficult to separate the effects of incumbency from Abramowitz’s “time for a change.” The two literally mirror each other—one “on,” the other “off”—since 1952, with the exception of a single election, 1992. Take your pick.

Notice that these possibilities either could have been (and in some cases, were) reflected in the forecasting models or otherwise have fairly small consequences. Taken together, they only take us so far, and leave a good amount (at least 2 percentage points) unaccounted for. Clearly, something else was at work, something forecasters did not—and maybe could not—anticipate well in advance of the election. Perhaps what the candidates did during the campaign made the difference.

One possibility is that Gore did not campaign on the Clinton record of accomplishment. From this perspective, Gore did not take credit for the economic and policy performance of the last four years and so he did not get credit, at least fully. Instead, he stepped away from this arena, where he was advantaged, into one where he was on more equal footing—the policy future. This is a tempting explanation, and there is some basis for it. That is, the success of the forecasting models in previous election years probably tells us as much about what the candidates did in their campaigns as it does about the importance of the economy and other aspects of presidential performance. Credit for a good economy does not happen magically; candidates make it happen. Perhaps, therefore, a Gore campaign centered around “*Extending Prosperity; Expanding Prosperity*” would have made a meaningful difference. While tempting, it is a less than satisfying explanation, at least from the forecaster’s point of view. It implies that there was nothing really wrong with the models after all. If nothing else, it probably takes the models a little too seriously.

Another possibility is that Gore shifted to the left on various issues (and that Bush drifted toward the middle). From this perspective, Gore not only stepped into the arena of the policy future where he was on equal footing, but actually placed himself at a disadvantage. This also is a tempting explanation, and there is some empirical basis for it. Exit polls suggest that the public placed Gore further from the political center compared to Bush.⁶ There also is additional support for the basic Downsian notion in the out-of-sample forecasts for previous elections (see Table 4). Recall that the biggest errors were in 1972, when Nixon did much better than predicted against the very liberal McGovern, and in 1980, when Carter did better than predicted against the conservative Reagan. Then again, Nixon and Carter were incumbents (and Gore was not) and it also was not “time for a change” in those two years (and it was in 2000). There is a lot to disentangle here. Still, let us assume that Gore shifted left and that it mattered. How far did he shift? How much difference did it make? Does it tell the rest of the story?

It may be that we can piece together the 2000 election using these different explanations. Of course, there are yet other possibilities. Perhaps the “character” issue played an important role. Regardless, what is perfectly clear is that the 2000 presidential vote reflected other things besides the economy, presidential approval, or even the “time for a change.” That Campbell’s model, which relies on polls from the Labor Day period, was the most accurate suggests that what mattered on Election Day was largely in place at the beginning of the fall campaign. That his forecast still was off, by 2.6 percentage points, implies that something also happened during the fall.

The View from the Polls


We may be able to gain insight into how the 2000 campaign unfolded with the benefit of the many polls conducted throughout the year. Figure 1 displays the basic data. It shows Gore’s percentage share of the two-party vote (ignoring Nader and Buchanan) in national trial-heat polls from the beginning of the year through Election Day, along with a smoothed series. The observations in the figure are the daily poll-of-polls. They represent Gore’s share for all respondents aggregated by the mid-date of the reported polling period. The procedure used exactly follows Erikson and Wlezien (1999).⁷ It allows readings for 173 separate days during 2000, 59 of which are in the period after Labor Day.

As can be seen in Figure 1, the daily poll results do exhibit pattern over time but also a lot of noise. Some of the noise is mere sampling error. There also are

other sources of error (see note 7). The smoothed series offers more clarity.⁸ Notice that Gore began the year well behind Bush and gained through the spring, when his support settled at around 47%, until the conventions. Also notice the predictable convention bounces, out of which Gore emerged in the lead heading into the fall. Up to this point, the election looked like 1988 all over again. The sitting vice president is running, the economy and presidential approval are favorable—he is behind in the polls early in the year, and then gains the lead for good after the party conventions. Things were playing out as the forecasters might expect. But then the parallel with 1988 stops. Gore’s support leveled off and began to drop, and continued to do so until just before Election Day, when his support rebounded.

Figure 2 zooms in on the post-Labor Day period. There is a lot of noise in the polls, partly simple sampling error. The early outlying observation is the

FIGURE 2
Trial-Heat Polls Aggregated by Date,
Labor Day to Election Day



notable *Newsweek* poll and the spike about 35 days out reflects Gallup’s fairly well-known flip-flop, the apparent result of their likely voter screen, which relied on “attention paid to the campaign”.⁹ The smoothed series nevertheless indicates real movement during the fall campaign—small, but significant. Gore entered the fall with just above 52% of the “vote,” which held until about mid-September, when the decline began. His support bottomed out at about 48% during the week before the election, and then turned back. Notice that the polls *in the field* at the very end of the campaign captured virtually all of Gore’s last minute surge. In these polls, the election was a dead heat.

In one sense, what we observed was fairly predictable, for researchers have shown that poll leads tend to

shrink leading up to the election (Campbell 2000; Wlezien and Erikson 1998). But Gore's lead didn't just shrink, it disappeared altogether and he actually fell behind. This is not common, though it is not unprecedented.¹⁰ Something happened during the fall to change voters' preferences. Perhaps this was the simple consequence of *voter learning* about, say, Gore's apparent left-shift. This is a reasonable interpretation, if only very general, and fits nicely with much political science research (see Alvarez 1997; Gelman and King 1993). But it is not obviously correct. Even if it is correct, this interpretation of Gore's decline does not help analysts understand why he gained at the very end and ultimately won, at least in the popular vote. Did the "fundamentals" kick in just in time? Did the revelation of Bush's DUI take its toll? Or was it something else altogether? All of this is beside the point. However one interprets the movement in voter preferences after the fact, what happened during the fall campaign was not predictable in advance.

On Election Forecasting

Election forecasting models work pretty well for a reason. The models capture (or reflect) information

Notes

* The article is based, in part, on research presented at the 2000 annual meeting of the Southern Political Science Association, which was held in Atlanta, and at a conference at Baylor University. Portions of the data were collected under a grant from the National Science Foundation (SBR-9731308). I thank Jose Bocanegra and Jeff May for assistance with data collection and Henry C. Adams Jr., James Campbell, James Glaser, Richard Johnston, and Tim Nokken for comments.

1. Political scientists do not have a monopoly, however. The economist Ray Fair (1978, 1982, 1988) has been forecasting for quite a while using a number of different models, all based mostly on economic performance. More recently, the historian Allan Lichtman (1996) has been forecasting using his thirteen "keys" to the White House. For a nice summary of the history of election forecasting, see the introductory essay in Campbell and Garand (2000).

2. The assumption is that the other voters would split proportionately or otherwise not vote.

3. For a description of the different models and analysis of previous performance, see Campbell and Garand (2000).

4. The estimate is based on our forecast of 55.2%, the standard error of the forecast, and the t-distribution with nine degrees of freedom. The standard error of the forecast is larger than the standard error of the estimate because it takes into account the standard errors of the coefficients (see also Beck 2000).

5. Of course, while it may have had little consequence for our forecasts, the Nader vote was enough to cost Gore Florida and the presidency, but evidently so was the Palm Beach butterfly ballot (see Greg Adams and Chris Fastnow's research and related links at <<http://madison.hss.cmu.edu/palm-beach.pdf>>).

6. According to exit polls, 43% of voters said Gore was "too liberal" while only 34% said Bush was "too conservative." If Gore really did shift left, he may have done so in response to Nader's candidacy. It also may be that Gore was responding to the demands of the major Democratic Party constituencies. Of course, it may be that he was simply reflecting his own agenda.

7. The data were drawn from PollingReport.com <www.pollingreport.com>. Where multiple results, reflecting different sampling universes, were reported for the same polling organizations and dates, data for the universe that best approximates the actual voting electorate is used, e.g., a likely voter sample instead of a

registered voter sample. Where there is overlap in polls, usually in tracking polls, conducted by the same survey houses for the same reporting organizations, only separate samples are used here. For example, where a survey house operates a tracking poll and reports three-day moving averages, poll results for every third day are used. This leaves 297 separate national polls. Since most polls are conducted over multiple days, each poll is dated by the middle day of the period the survey is in the field. No adjustment to the polls is made for differences in question wording, sampling universe, survey house, or the various likely voter screens and weighting procedures. Some of these differences are of real consequence, however, especially the practices of different survey houses (see Erikson and Wlezien 1999). Poll results vary from day to day in part because the polls reported on different days are conducted by different houses. Thus, the data presented here offer only a very coarse picture of voter preferences in 2000.

8. The series was generated using the "lowess" (locally weighted scatterplot smoothing) procedure (see Jacoby 1997). This procedure creates a new value for each time point based on the results of regressions using a designated number of surrounding data points. To generate the new value, predictions from these regressions are weighted based on their temporal distance from the particular point in question. Lowess tends to follow the data quite well, though the degree to which it does depends on the bandwidth, or percentage of time points in the full series, one uses to generate the smoothed values. For this exercise, a bandwidth of 0.1 was used.

9. On October 1, Gallup reported a dead heat in its three-day tracking poll. Three days later, Gore was ahead by 11 points. Two days after that, where one-third of the respondents in the two polls were the same, Bush was ahead by eight points. Party identification in the three polls shifted from 35-34% Democrat to 38-30% Democrat and then to 38-30% Republican.

10. The most notable exception is 1948, when Truman surprised most election observers by defeating Dewey. The lead changed a lot in 1980, when Reagan began the fall with a miniscule lead, lost it for a while, gained it back, lost it briefly once again, and then regained the lead for good. In 1960, the lead also shifted hands, though Kennedy began the fall with only a one-point advantage and his vote share varied within a narrow range, between 49 and 51.5%.

References

- Abramowitz, Alan I. 1988. "An Improved Model for Predicting Presidential Outcomes." *PS: Political Science and Politics* 21(December): 843-47.
- Alvarez, R. Michael. 1997. *Information and Elections*. Ann Arbor: University of Michigan Press.
- Beck, Nathaniel. 2000. "Evaluating Forecasts and Forecasting Models of the 1996 Presidential Election." In *Before the Vote: Forecasting American National Elections*, ed. James E. Campbell and James C. Garand. Thousand Oaks, CA: Sage.
- Bernstein, Alan. 2000. "Political Forecasters Sticking with Gore Despite All the Polls." *Houston Chronicle*, November 2, 13A.
- Brody, Richard, and Lee Sigelman. 1983. "Presidential Popularity and Presidential Elections: An Update and Extension." *Public Opinion Quarterly* 47(Fall): 325-28.
- Campbell, James E. 1993. "Weather Forecasters Should Be So Accurate: A Response to 'Forewarned Before Forecast.'" *PS: Political Science and Politics* 26(June): 165-66.
- , 1996. "Polls and Votes: The Trial-Heat Presidential Election Forecasting Model, Certainty, and Political Campaigns." *American Politics Quarterly* 24(October): 408-33.
- , 2000. *The American Campaign: U.S. Presidential Campaigns and the National Vote*. College Station: Texas A&M University Press.
- , and James C. Garand, eds. 2000. *Before the Vote: Forecasting American National Elections*. Thousand Oaks, CA: Sage.
- , and Kenneth A. Wink. 1990. "Trial-Heat Forecasts of the Presidential Vote." *American Politics Quarterly* 18(July): 251-69.
- Erikson, Robert S. 1989. "Economic Conditions and the Presidential Vote." *American Political Science Review* 83(June): 567-73.
- , and Christopher Wlezien. 1996. "Of Time and Presidential Election Forecasts." *PS: Political Science and Politics* 29(March): 37-39.
- , and Christopher Wlezien. 1999. "Presidential Polls as a Time Series: The Case of 1996." *Public Opinion Quarterly* 63(Summer): 163-77.
- Fair, Ray. 1978. "The Effect of Economic Events on Votes for President." *Review of Economics and Statistics* 60:159-72.
- , 1982. "The Effect of Economic Events on Votes for President: 1980 Results." *Review of Economics and Statistics* 64:322-25.
- , 1988. "The Effect of Economic Events on Votes for President: 1984 Update." *Political Behavior* 10:168-79.
- Gelman, Andrew, and Gary King. 1993. "Why Are American Presidential Campaign Polls So Variable When Votes Are So Predictable?" *British Journal of Political Science* 23:409-51.
- Hibbs, Douglas A. Jr. 1987. *The American Political Economy*. Cambridge, MA: Harvard University Press.
- Holbrook, Thomas. 1996a. "Reading the Political Tea Leaves: A Forecasting Model of Contemporary Presidential Elections." *American Politics Quarterly* 24(October): 506-19.
- , 1996b. *Do Campaigns Matter?* Thousand Oaks, CA: Sage.
- Jacoby, William. 1997. *Statistical Graphics for Univariate and Bivariate Data*. Thousand Oaks, CA: Sage.
- Kaiser, Robert. 2000. "Is This Any Way to Pick a Winner?" *The Washington Post*, May 26, A1.
- Kurtz, Howard. 2000. "The Premature Post-Mortems Are Starting" <www.washingtonpost.com/wp-dyn/articles/A47391-2000Oct31.html>. *Washington Post Online*, Oct. 31. Accessed: January 4, 2001.
- Lewis-Beck, Michael S., and Tom W. Rice. 1984. "Forecasting Presidential Elections: A Comparison of Naïve Models." *Political Behavior* 6:9-21.
- , and Tom W. Rice. 1992. *Forecasting Elections*. Washington, DC: CQ Press.
- Lichtman, Allan. 1996. *The Keys to the White House, 1996: A Surefire Guide to Predicting the Next President*. New York: Madison
- Lockerbie, Brad. 2000. "Election Forecasting: A Look to the Future." In *Before the Vote: Forecasting American National Elections*, ed. James E. Campbell and James C. Garand. Thousand Oaks, CA: Sage.
- Norpoth, Helmut. 1996. "Of Time and Candidates: A Forecast for 1996." *American Politics Quarterly* 24:443-67.
- Stout, David. 2000. "Experts Once Certain, Now Say Gore Is a Maybe." *The New York Times*, November 7, A23.
- Tufte, Edward R. 1978. *Political Control of the Economy*. Princeton: Princeton University Press.
- Wlezien, Christopher, and Robert S. Erikson. 1996. "Temporal Horizons and Presidential Election Forecasts." *American Politics Quarterly* 24(October): 492-50.
- , and Robert S. Erikson. 1998. "The Timeline of Presidential Election Campaigns." Presented at the Annual Meeting of the American Political Science Association, Boston.
- , and Robert S. Erikson. Forthcoming. "After the Election: Our Forecast in Retrospect." *American Politics Research*.