Forecasting U.S. National Elections

JAMES E. CAMPBELL
University at Buffalo, SUNY

JAMES C. GARAND
Louisiana State University

The subject of this book is the scientific forecasting models of American national elections—principally presidential elections—and the application of these models in both the recent 1996 election and the upcoming election of 2000. The book is built around the October 1996 special issue of *American Politics Quarterly.* That issue for the first time assembled in one scholarly journal a group of election forecasting models and reported their specific forecasts of the national vote well *before* election day. Each of these models offered a national vote forecast at least 2 months before election day and in many cases more than 3 months before the votes were cast. Each presidential vote forecast was also accompanied by some assessment of the degree of certainty in its prediction of which presidential candidate would receive a majority of the national two-party popular vote.

Forecasting anything, from tomorrow's weather to the stock market, is generally a risky business, and forecasting a national election outcome based on the vote decisions of nearly 100 million American voters is not for the faint of heart. Nevertheless, the scholars whose work is assembled in this volume thought that it could be done with some considerable, if less than perfect, certainty and quite publicly stuck their necks out in the special issue of *APQ* with their specific forecasts several months before the 1996 election.

The first and major section of this volume reproduces the six presidential and one congressional election forecasting articles originally published in that special issue of *APQ* and adds to the mix a chapter with a seventh presidential model.[1] Each of the original forecasting articles is reproduced in unrevised form as it appeared in the October 1996 issue of *APQ* but includes an appended new "After the 1996 Election" section that evaluates the 1996 forecast and further analyzes the forecast model in light of the 1996 election experience.

The general purpose of these postmortem sections is to discuss what has been learned about the forecasting models as a result of the 1996 election. Each addresses three aspects of how a particular model performed in the 1996 election and how that performance affects the specification and the estimation of the model for use in future elections. The first subject of each postmortem is to assess the accuracy of the model's forecast for the 1996 election and to compare it to the model's previous performance. Second, after making whatever revisions in the model that were deemed appropriate (and most make no revisions whatsoever), each forecasting model is updated by including the 1996 election and reestimating the model's coefficients. Third, each postmortem provides a contingency analysis for its updated forecasting model so that readers can readily use the model themselves to forecast the upcoming election of 2000 and to see more concretely how different factors affect the forecasts.

The last chapter of the first section departs from the presidential focus to offer a model predicting the results of voting for the U.S. House of Representatives. As difficult as presidential election forecasting is, because of the staggered Republican realignment and its delayed and gradual (until 1994) filtering into congressional politics, congressional election forecasting has been even more difficult. Bypassing some of the more remote factors affecting congressional elections (and thereby many of the forecasting difficulties related to the realignment), Lee Sigelman and Robert Erikson have developed a congressional forecasting model based on the "generic" House vote question similar to the trial-heat question posed in presidential contests: "If the elections for Congress were being held today, which party's candidate would you like to see win in your congressional district?" As one would well expect, responses to this "generic House vote intention question" asked early in campaigns have been historically related to the vote cast some months later.

The second section of this book contains two chapters taking broader perspectives on the methodological and theoretical issues confronting election forecasting. Nathaniel Beck focuses on the specific issue of how to evaluate and report the uncertainty about the accuracy of any specific or point forecast of the national vote. Although forecasting is an exacting science, like all sciences involving human behavior, it is inexact. We cannot be perfectly certain of any vote prediction generated by any forecast model. The question that Beck addresses is how to best appraise the degree of uncertainty that we should have about any par-

ticular vote forecast. In the second chapter, taking a broad perspective on election forecasting, Jim Campbell discusses four issues confronting forecasting: (a) the general track record of the forecasts for their accuracy and their overall performance in the 1996 election, (b) the issue of how to evaluate uncertainty about a forecast, (c) the contribution of election forecasting to explanatory theories of elections (and, relatedly, the different concerns of explaining and predicting elections), and (d) the potential general impact of forecasting on how political science research is conducted. He finds (a) that the forecasting models as a group have been quite accurate, (b) that there are still unresolved issues about how to appraise uncertainty but that these suggest that uncertainty about forecasts may be somewhat overstated, (c) that presidential election forecasting models have contributed to explanatory electoral theories in several ways, and (d) that interest in forecasting may have the beneficial side effects of encouraging political science research to be more concerned about directly observable aggregate-level political subjects, about precision in its findings, and about the completeness of its explanations.

The final portion of the book, the appendix, includes all of the data that were used in the estimation of the original and updated versions of the various forecasting models. These data are provided so that the reader can replicate, compare, and extend the analyses reported in each of the forecasting model chapters.

Before heading directly into the chapters addressing each of the individual forecasting models, it might be useful to provide some general background on the development of the field of election forecasting and some of the issues concerning it. After reviewing the history of forecasting research, we will review its most recent history—the performance of the forecasting models in the 1996 elections. What did the various forecasting models predict in the 1996 election, and how accurate were they? With that background, we then will turn to two fundamental issues that those in this field are often asked about. The first of these is, Why forecast? Why should anyone be interested in election forecasting? The second question is about the possible impact of forecasting. What are the potential implications of election forecasting for democratic governance through the electoral process?

SOME BACKGROUND ON ELECTION FORECASTING

Although election forecasting outside of political science has a long and colorful history replete with bellwethers, straw polls, and punditry, systematic election forecasting within political science has a relatively brief history. Despite Louis Bean (1940, 1942, 1948) blazing a forecasting trail in the 1940s, a significant body of research on election forecasting only began to accumulate in the late 1970s.

The recent wave of forecasting models began with Lee Sigelman's (1979) analysis of the connection between presidential approval ratings and subsequent election results (also see Brody & Sigelman, 1983), Steven Rosenstone's (1983) model of presidential election results in the states, and Michael Lewis-Beck and Tom Rice's (1984a, 1984b) adaptation of Edward Tufte's (1978) approval rating and economic performance model to forecast both congressional and presidential elections. Alan Abramowitz (1988) amended the Lewis-Beck and Rice approval and economy model by appending a "time for a change" variable to it, and Jim Campbell and Ken Wink (1990), in following a lead from Lewis-Beck and Rice (1985), built a model around the trial-heat poll question ("if the election were held today . . .") and economic conditions. Lewis-Beck and Rice (1992a) significantly amended their initial model by adding indicators from the presidential primaries and prior congressional elections. Building on Rosenstone's state-level model, Campbell (1992) also constructed a model to predict the presidential vote in the individual states based on the trial-heat polls, economic indicators, and a number of state-level variables. With the exception of the state models developed by Rosenstone and Campbell (and reanalyzed by Gelman & King, 1993), all of the forecasting models have been national time-series models.

Scholars from sister disciplines to political science have also been drawn to the forecasting arena. Yale economist Ray Fair (1978, 1982, 1988) developed a national time-series model comparable to the national political science models, with two important exceptions. His model did not include any measure of public opinion, and because of this, it was estimated over nearly twice as many elections as the political science models (covering elections since 1916). Economic conditions and incumbency are at the core of Fair's model. Historian Allan Lichtman covered even a longer span of electoral history (since 1860) with his forecasting rules (Lichtman, 1996; Lichtman & DeCell, 1990). Lichtman identified 13 "keys" or indicators to the presidential election and devised a decision rule to predict the winning presidential candidate based on the number of keys favoring each party's candidate. From the standpoint of political science, it is tempting to criticize Fair's model for even attempting to predict election results without a reading of public opinion and Lichtman's 13 keys for the subjectivity and crudeness of its indicators, as well as its oversimplification of how they are combined to predict simply a winner or loser (rather than the candidates' portion of the vote). Nonetheless, it is heartening to see a diversity of approaches and a lively interest in the enterprise.

For any model, it is one thing to develop it in the abstract and another to have it tested under fire during an election year. Five of the models produced forecasts prior to the 1992 election: (a) Fair's economy incumbency model, (b) Lewis-Beck and Rice's original approval economy model, (c) Lewis-Beck and Rice's amended and more complex model that incorporated internal party divisions (as reflected in presidential primary results) and partisan trends (as reflected in

prior midterm results), (d) Abramowitz's approval-economy-time for a change model, and (e) Campbell and Wink's trial-heat and economy model (Campbell, 1993b; Campbell & Mann, 1992; Greene, 1993). Of this crop, the election offered a clear verdict of three winners and two losers. Each of the three winners—the original Lewis-Beck and Rice model, the Abramowitz model, and the Campbell and Wink model—were accurate to within one percentage point of the vote in predicting the Clinton victory and offered forecasts by or before Labor Day. Each was a good deal more accurate than the polls at the time that the forecast was made and more accurate than the seat-of-the-pants predictions made just days before the election by most of the pundits (Campbell, 1993b). One of the two losers was Lewis-Beck and Rice's amended and more complicated model, which called for a narrow Bush win and missed the two-party vote by more than 5 percentage points. The biggest miss was produced by Fair's model, which predicted a solid majority (about 56%) for President Bush, an error of a whopping 9 percentage points.

After the 1992 election, some of the established presidential models were revised or otherwise amended, and all were updated to include the case of the 1992 election. Three of these are presented in this volume. The models by Jim Campbell and by Alan Abramowitz are fundamentally unchanged. Michael Lewis-Beck and Charles Tien substantially revised the earlier Lewis-Beck and Rice model by dropping two variables (midterm seat change in the House and presidential support in the primaries) and adding a third (the public's sense of national peace and prosperity) to the core model of presidential approval and economic conditions.

In addition, and indicative of the growing interest in forecasting, a number of new and intriguing models were developed for 1996 and are also presented in this volume. Thomas Holbrook offered an adaptation of the three-variable Abramowitz model, examining an average of second-quarter presidential approval ratings and a survey-based assessment of economic conditions. Christopher Wlezien and Robert Erikson developed a two-variable model based on presidential approval ratings and a cumulative index of leading economic indicators. Brad Lockerbie devised a three-variable model using both objective and poll-based indicators of the economy and a "time for a change" variable. Finally, encompassing a much longer series of elections than any of the other forecasting models, Helmut Norpoth constructed a five-variable model that included the votes in the prior two presidential elections, the candidate's primary strength, and two indicators of economic conditions.

As we have already noted, the forecasts of six presidential election models were published in *APQ* prior to the 1996 election and are reproduced with an update and reanalysis in this volume. The 1996 forecasts of these six models plus a new model by Brad Lockerbie are presented in Table 1 along with the difference between each forecast and the actual 1996 vote.

Table 1 Point Forecasts of the National Two-Party Presidential Vote in 1996

Forecaster and Basis of the Model	*Predicted Two-Party Presidential Vote for In-Party (Clinton)*	*Forecast Error*	*Forecast Lead Time in Days Before Election Day*
Abramowitz	56.8	+2.1	92
Approval, gross domestic product, and term	(—)		
Campbell	58.11	+3.8	61
Trial-heat and gross domestic product	(96%+)		
Holbrook	57.2	+2.5	124
Approval, survey on personal finances, and term	(99.5%)		
Lewis-Beck and Tien	54.82	+.1	68
Approval, gross domestic product, and survey on peace-prosperity	(95%+)		
Lockerbie	57.56	+2.8	89
Disposable income, consumer sentiment, and term	(99%)		
Norpoth	57.1	+2.4	92
Prior votes, gross domestic product, inflation, primary	(90%)		
Wlezien and Erikson	56.0	+1.3	92
Approval and leading economic indicators	(89%)		
Mean absolute error		2.1	

NOTE: An estimate of the certainty about the forecast prediction of the plurality vote winner is presented in parentheses. Except for Campbell's model, all predictions were of the conventional two-party vote for the in-party candidate (Clinton). Campbell's prediction was of the two-party vote, with minor-party votes divided evenly between the major-party candidates. The actual two-party vote for Clinton was 54.737% (54.261% calculated with an even division of third-party votes). Abramowitz later revised his forecast to use a different approval reading. This change reduced the 1996 error in this model to 1.2 percentage points. See Abramowitz in this volume. Also, the Wlezien and Erikson forecast also depended on which quarter of economic data were used. The gross domestic product, presidential approval, and term variables differ from model to model in the time period examined. For further details regarding the models, see the relevant chapters in this volume.

How did the models do in 1996? In general, the models did fairly well. Each of the seven models correctly predicted that President Clinton would be reelected and that the election would not be close. The mean error of the models was only about 2 percentage points of the vote. Of all the models, Lewis-Beck and Tien's forecast was indisputably the most accurate, missing the actual vote by a mere tenth of a percentage point. Wlezien and Erikson's forecast was also quite accurate, missing the actual vote by only 1.3 percentage points. With the exception of Campbell's model, which was wide of the mark by 3.8 percentage points, the remaining models were all within 3 percentage points of the vote. This is about the magnitude of error that one might expect from a national poll conducted around election day—only these forecasts were made several months beforehand.

One additional aspect of the 1996 performance of the forecasting models is especially noteworthy: Without exception, each of the models forecast a larger Clinton vote than he actually received. The mean two-party vote prediction for Clinton was 56.8%, 2.1 percentage points higher than his actual 54.7 two-party vote. The 1996 election was not close, but it was substantially closer than most of the models (and the polls) led us to believe. One would expect that the forecasts would bracket the actual vote, but in 1996 they were all on the high side.[2] The reason for the different models with good track records for accuracy universally missing the vote on one side is unclear. At this point, we can only speculate whether unanticipated campaign events or some other reason caused the general overprediction of Clinton support and underprediction of Dole support. Although this pattern of the forecast errors is intriguing, it should not distract us from the fundamental points about the 1996 forecasts: All of the models correctly forecast a Clinton victory, a couple of the models were quite accurate, and most were within a couple of percentage points of the vote.

With this background and recent record of election forecasting in place, we now turn to two questions about forecasting. First, why forecast elections? Second, does forecasting elections undermine their democratic function?

WHY FORECAST ELECTIONS?

Of the many forecasting issues, perhaps none is as basic as its purpose. Why the great interest in forecasting elections? Some serious political scientists consider forecasting beyond the bounds of the scientific study of elections and voting behavior. The charge is that science explains and prediction is not explanation. Others say that if you cannot control the phenomenon, there is no use in forecasting it. Still others suggest that if forecasting allows candidates to control elections, it undermines the electoral process. So, why forecast elections?

At one level, election forecasting is of interest for many of the reasons that we are interested in forecasting anything. We forecast some things to control them

and to prepare for what is to come for the other things that we cannot control. We do not forecast the weather to control it; we forecast it to prepare for it. Likewise, we may want to begin preparing for the election result. Perhaps, too, at the margin we can control events, especially if we act in a timely manner. However, this does not mean that elections necessarily are any more subject to manipulation than they always have been. It is hardly news to politicians that they stand a better chance of reelection when the economy is prosperous and the public is otherwise satisfied. Forecasting models may reveal more about the precise levels of satisfaction and prosperity that are required for reelection and about exactly when these conditions are most electorally important, but this is sharpening existing knowledge rather than providing new knowledge. Moreover, as Key (1966) suggested long ago, voters are not fools, and democracy must work on the premise that enough voters are knowledgeable and skeptical of politicians that they can detect election-inspired conversions and policies.

As in most things, a principal reason for forecasting is curiosity. We want to know what has happened as soon as it happens, if not sooner. People sit up into the wee small hours of the night watching televised coverage of elections to see who won and who lost. Networks take great pride in being first to declare winners and losers. Maybe we should not get carried away with curiosity, but that is like railing against breathing. It is so ingrained in people that we might as well accept it. Curiosity is inevitable, and because forecasting will be done one way or another, it ought to be done as well as it can be. If elections will be forecast, we ought to do what we can to get it right.

Apart from curiosity, election forecasting may be of interest because it can provide useful information to those involved in legislating public policy. The president plays a particularly important role in the policy process, in marshalling support or opposition to a piece of legislation and his decision of whether to veto legislation passed by Congress. If one side of a policy dispute (legislators and interest groups) knows that it will have a political ally or a political adversary in the White House in the foreseeable future, this knowledge may affect its willingness to compromise on legislation. If it appears that one side in a dispute will lose a friendly White House in the upcoming election, it may be eager to strike a deal now rather than wait and face a less hospitable political climate later on. Conversely, if it appears that the likely new president will be an ally replacing an opposing president, then there may be reason to be less flexible in compromising or reaching an agreement with others. For instance, conservatives believing that a liberal president will shortly be replaced by a more conservative successor may be less willing to compromise their views because they think that their political position will be strengthened in the next election. If they wait, they may get more of what they want; so why rush to take less?

Election forecasting may also be of use to anyone observing and attempting to make sense of the political campaigns. Campaigns run on two tracks. The gov-

erning track involves the policy debate between the candidates and may affect or reveal what a candidate would do in office. The second track, and the one that ironically is the one most closely followed, is the political track. The political track of the campaign affects or reveals who will be elected. Increasingly, most observe the political track of the campaign by following the ups and downs of the candidates in the national preference polls, and the media have increasingly followed suit in what critics have decried as horse race coverage of the campaign. The focus is on who is ahead and who is behind, who is moving up, and who is fading back, rather than who stands for what. Once the campaign is in full gear, polls are conducted and released daily, and voters watch a veritable electoral electrocardiogram, a zigzag-like plot of poll numbers week after week until it all mercifully ends on election day. Critics of the media and the public's interest in the polls have argued that the overemphasis on poll watching has distracted the electorate from the campaign's real business of providing voters with useful substantive information about what the candidates would do if elected and why they would do it. The poll-dominated campaign, by this view, has debased the democratic process, exchanging the educational function of campaigns for their entertainment value as contests.

On first blush, it might appear that scientific election forecasting is merely a high-tech continuation of this debasement. Without question, the focus of forecasting is on the political aspect of campaigns rather than on their governing aspect. That said, scientific election forecasting, by setting the endless ups and downs of the poll numbers in some perspective, may ultimately help to diminish interest in horse race coverage of the campaigns and free up more attention for the broad policy differences between the candidates. Why should we spend much time reading or reporting about the daily perturbations of the polls when the forecasts have already revealed where they will inevitably lead? What better way to discourage the viewing of elections as matters of entertainment than to reveal the election story's ending?

Election forecasting may also be useful in providing a baseline from which to evaluate the effects of the particular events that occur in every campaign. The forecasting models are based on the usual patterns observed in past elections. For instance, a certain level of economic growth over a particular period before the election usually translates into a certain effect on the vote. To the extent that the vote in an election deviates from that expected by the forecasts, it may mean that something unique to that campaign made a real difference.

Finally, election forecasting should be of interest because it may contribute to understanding elections in a larger sense, beyond the events of a single election. Both successes and failures in forecasting models can shed some light on what affects elections. In the concluding chapter of this volume, Jim Campbell discusses several potential contributions of presidential election forecasting to explanatory theories of elections.

In particular, the forecasting models have illuminated the nature of the effects of both partisanship and presidential incumbency. Presidential election forecasting models have also both stimulated and guided research by Gelman and King (1993) and Holbrook (1996) into whether and how presidential campaigns have affected the results of elections.

DO ELECTION FORECASTS UNDERMINE DEMOCRACY?

One commonly voiced criticism of election forecasting is that it undermines the electoral process as an instrument by which citizens can control their government. From this perspective, election forecasting discourages potential voters from voting, reduces the sense of efficacy among those who do vote (because the election may be regarded as a settled matter before a vote is ever cast), and provides information to candidates so that they can manipulate voters. The complaints against election forecasting of discouraging and dispiriting voters are similar in some respects to those leveled about the media's reporting of exit poll results before the polls have closed on election day.

These are serious charges against forecasting. Although these complaints are generally made as a piece, they are really quite separable. The claim that election forecasts depress turnout may have merit, although the effects of forecasting (if any) on turnout are probably quite small. Even though most voters do not seriously believe that their vote would make *the* difference between a presidential candidate winning or losing, a belief that the results of the election are a foregone conclusion may diminish the efforts of potential voters to get to the polls. A busy potential voter on election day who is convinced that one candidate has all but won the election may decide that it is not worth expending the extra effort to get to the polls. This applies to both supporters and opponents of the candidate who has been forecast to win the election.

The effect of forecasting on the efforts of potential voters to turn out is probably quite small because most potential voters who intend to vote are motivated by strong preferences for one candidate over another rather than a delusion that their vote will decide the election. If the vote is cast to express a preference or to do what the voter can easily do for the favored candidate, whether that candidate has been forecast to win or lose the election should not matter. In addition, it seems likely that those who are most likely to be aware of a forecast are also more likely to be attentive to politics, have strong opinions, and are not easily dissuaded from voting. Moreover, for those who might be aware of the election forecasts and whose turnout decision might be influenced by them, the information provided by the forecasts should not necessarily have a demobilizing effect. In fact, when the models forecast a very close election or are in conflict over the likely election outcome, they might actually serve to increase turnout somewhat.

The claim that election forecasting may alienate voters from the political system and undermine their sense that their actions may have an effect would seem to be based on a false notion about forecasting: the idea that the accuracy of the forecasts is somehow independent of what voters decide to do. Ultimately, the electorate decides the election, and the accuracy of the forecast models is entirely dependent on what the voters say about their intentions (in the form of approval ratings and trial-heat polls) and the consistency of their reactions to the circumstances surrounding an election with their reactions to similar circumstances (such as economic growth) in past elections. In short, there is no good reason for voters to feel disempowered by forecast models because the accuracy of these models is entirely dependent on the voters themselves. Forecasts based on the reported and anticipated reactions of voters may highlight the relatively small impact that any one voter has, but this is a problem stemming more from voting in a large electorate than a problem of forecasting itself.

The final antidemocratic claim against forecasting is that it provides candidates with information that they can use to manipulate voters. There seems to be little merit in this charge. First, the factors that make an election predictable are not necessarily or readily subject to manipulation themselves. The largest component of most forecasting models is a measure of precampaign public opinion. It is hardly particularly helpful to candidates to tell them that they ought to have public opinion on their side at the outset of the general election campaign. The other components of most forecasting models are either beyond manipulation (the status of incumbency) or also not easily manipulable (the election year economy). Moreover, by no stretch of the imagination do candidates learn what helps in an election from the forecasting models. It is obvious that a candidate who is popular at the outset of a campaign, has the advantages of incumbency, and has favorable economic conditions behind him or her stands a better chance of election.

If anything, by revealing that the principal factors affecting an election are set in place well before the fall campaign gets under way, forecasting models should indicate to candidates the futility of attempting to manipulate voters late in the campaign. What candidates ought to learn from the forecasting models is that voters make up their minds about how they will vote over a longer period than that encompassed by the campaign, and this should enhance the role of elections as a instrument of democracy.

NOTES

1. The seventh presidential model was developed by Brad Lockerbie, whose chapter follows the format of the other forecasting chapters. The first portion is a model estimated without knowledge of what would happen in the 1996 election, and the second section is a

post-1996 reappraisal of the model. Although this chapter did not appear in the October 1996 issue of *American Politics Quarterly,* the basis for Lockerbie's forecast, the actual forecast, and the degree of certainty were published in the compendium of 1996 forecasts at the end of the October *APQ* issue.

2. Although all of the forecasts in the *APQ* compendium overpredicted the Clinton vote, one forecast not included underpredicted the Clinton vote. Based on the incumbent running for a second consecutive term for his party, a predicted third-quarter growth rate in the per capita gross domestic product of 2.1%, an absolute annual inflation rate of 3% through the third quarter of the president's term, and eight quarters of "good economic news," Ray Fair's equation, estimated over the 20 elections from 1916 to 1992, predicted a two-party vote for Clinton of 49.5%. According to this model, the election was a toss-up. As already noted, an earlier version of Fair's model in 1992 predicted a Bush victory. For more details, see Fair (1996).