Pre-Election Patterns

The evolution of voter preferences

By Christopher Wlezien

and Robert S. Erikson

Figure 1

Dramatic Poll Shifts in 1992

Questions:

Variations on: If the presidential election were being held today, would you vote for... or...?

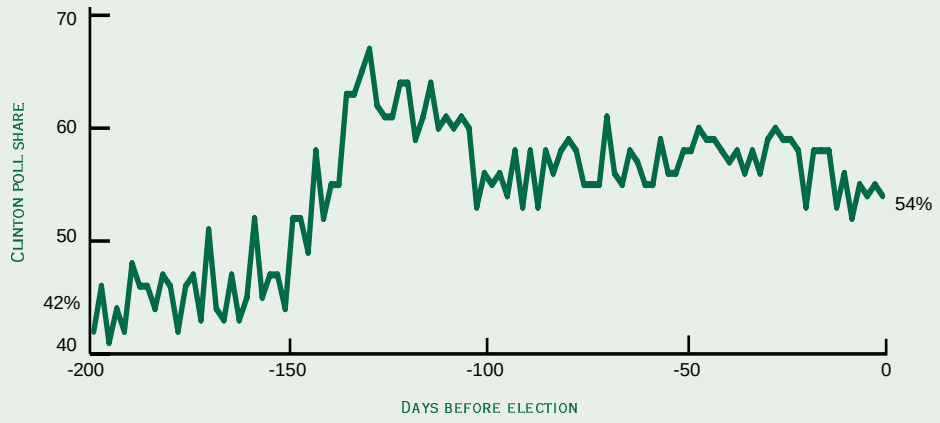

The growth of pre-election polling is well known. In 1968, there were some 22 published "trial-heat" polls pitting the two major-party candidates against each other. By 1980, the number of presidential polls exceeded 100. For 2000, results for more than 500 polls of the Bush-Gore "vote" are listed on the well-known website, PollingReport.com.

All of this ado can't be about nothing, and, indeed, a close analysis of pre-election polls reveals patterns in the evolution of voter preference in presidential elections since 1944. These patterns hold important implications for the stagers of both presidential and congressional campaigns.

Tracking the movement of opinion in sample surveys is not a simple task, however, as trial-heat results represent a combination of true preferences and survey error.

Survey error takes a number of forms. The most basic manifestation is sampling error. Since all polls contain some degree of sampling error, we will observe changes from poll to poll even when the division of candidate preferences is constant and unchanging. This problem is well known, but not easy to address: we cannot separate sampling error from reported poll preferences. Sampling error is random.

Survey results also reflect departures in practice from simple random sampling, which might include clustering, stratifying, weighting, and the like. For election polls, the major source of such design effects relates to the polling universe. Determining who will vote on Election Day is not easy; all we can do is estimate the voting population.

And so survey organizations typically use "screens" to identify likely voters. Some weight their samples by selected distributions of party identification or other variables so as to approximate the likely voting electorate. How one screens and weights, of course, has

Christopher Wlezien is Reader in Comparative Government and Fellow of Nuffield College, University of Oxford. Robert S. Erikson is professor of political science, Columbia University. *A more detailed version of this article appears in the current issue of the Journal of Politics.*

Figure 2

Little Movement in 1996

Questions:

Variations on: If the presidential election were being held today, would you vote for... or...?

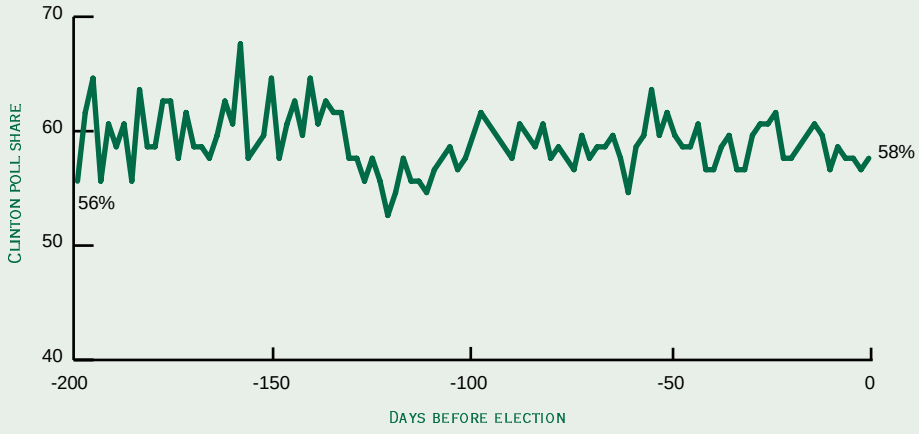

important consequences both for poll margins at each point in time *and* for the variance in results over time.

When dealing with polls from different survey organizations, house effects also are a problem. Results can differ depending upon the ways in which survey houses collect data, train interviewers, and cope with refusals. As with design effects, poll results will vary from day to day because the polls reported on different days are conducted by different houses. The effects can be quite pronounced.

Putting aside survey error, the remaining changes in poll results reflect actual movement in voters' preferences.

As part of our ongoing project of estimating the effects of campaign effects on modern presidential elections, we attempted to locate all national polls that mentioned the actual Democratic and Republican nominees for the 15 elections between 1944 and 2000. In these polls, respondents were asked how they would vote if the election were held today or, less frequently, who they would like to see win.

The bulk of the data were drawn from the Roper Center's iPOLL database, but other sources also were used, including *The Gallup Report*, *Public Opinion*, and *Public Perspective*. For 1996, the data were drawn primarily from the now-defunct *PoliticsNow* website, supplemented by data from *Public Perspective* and the Roper Center. For 2000, the data were taken entirely from PollingReport.com.

Where multiple results reflecting different sampling universes were reported for the same polling organizations and dates, we used data for the universe that in theory best approximated the actual voting electorate. For example, where a survey house reported poll results for samples of registered and likely voters, we used the data for the latter.

We also removed all overlap in polls—typically tracking polls—conducted by the same survey houses for the same reporting organizations. For example, where a survey house operated a tracking poll and reported 3-day moving averages, we only used poll results for every third day. This weeding process left 1,429 separate national presidential preference polls between 1944 and 2000.

Since most polls are conducted over several days, we dated each by the middle day of the period the survey was in the field. For polls that were in the field for an even number of days, we rounded up the fractional midpoint; for instance, for a poll in the field four days, we centered the poll on the third day.

Using this method, the 1,429 polls allowed readings (often multiple) for 984 different days from 1944 to 2000. We then generated a daily poll-of-polls for the 15 election years. The numbers represented the Democratic share of the two-party vote intention—ignoring all other candidates—for all respondents aggregated by the mid-date of the reported polling period. Wherever possible, respondents who were undecided but leaned toward one of the candidates were included in the tallies.

What did these data tell us once they had been identified and so painstakingly arranged? To begin with, we observed that the polls exhibited a lot more volatility in some years than in others. In 1992, for example, the polls shifted dramatically during the last 200 days of the campaign. This can be seen in Figure 1. Conversely, in 1996, the polls did not move much at all, as is clear in Figure 2. As Figure 3 shows, the 2000

Figure 3

Moderate Change in 2000

Questions:

Variations on: If the presidential election were being held today, would you vote for... or...?

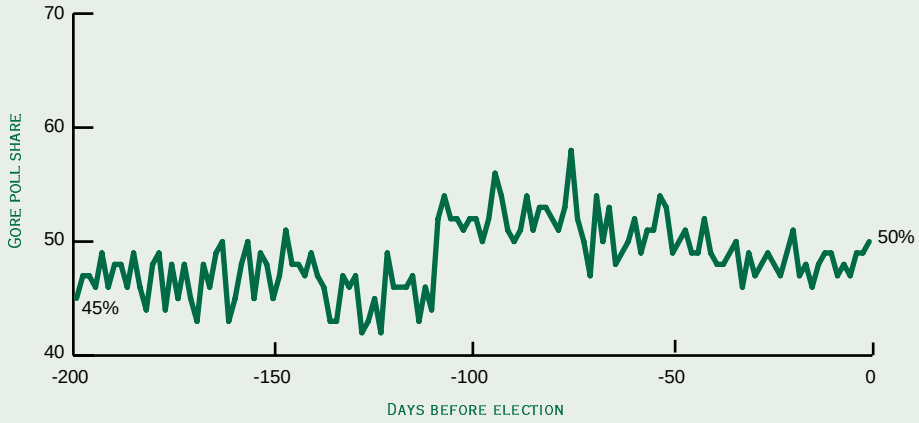

election was somewhere in between. Put simply, campaign dynamics differed meaningfully across elections.

Of course, as we already have discussed, these poll results combined real movement in preferences with survey error. Although we could not fully disentangle survey error from reported preferences, we nevertheless could ask how much of the observed statistical variance was real.

We provided a basic answer to this question by comparing the variance we observed with what we would have expected to observe, given the sample sizes and simple random sampling, if electoral preferences were, in fact, constant over the campaign. The surplus indicated the true variance of preferences, and it was relatively easy to compute.

These calculations indicated that the true variance differed significantly from year to year. In 1992, sampling error accounted for only 5% of the variance in the polls during the 200 days prior to the election. In 1996 and 2000, the same amount of sampling error accounted for more than 30% of the observed variance. In some years, up to half was sampling error.

Nevertheless, most of the variance in poll results over each long campaign was "real." On average, over the 15 presidential elections from 1944 to 2000, the ratio of true variance to error variance was about four to one. We know this because the variance we observed was about four times what we would have expected from sampling error alone, that is, if survey respondents made choices by flipping coins.

Most of the variance in preferences was concentrated in the period leading up to the fall general election campaign, however. After Labor Day, the unofficial kickoff of the general elec-

Figure 4

Outcome Comes Into Focus as Election Cycle Unfolds

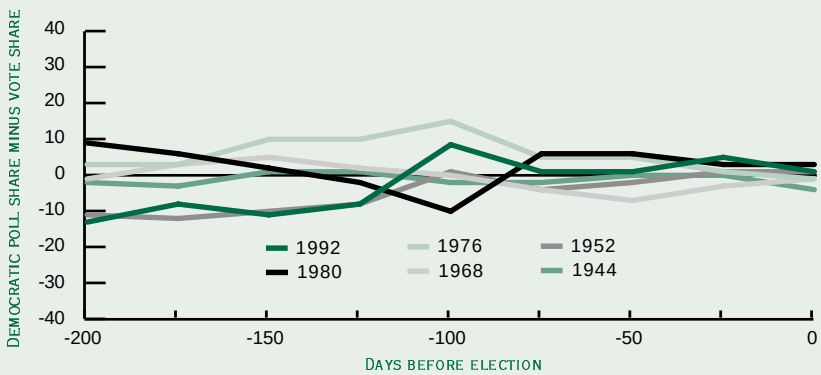

Days till election	DEMOCRATIC POLL SHARE MINUS VOTE SHARE					
	1992	1980	1976	1968	1952	1944
-200	-13	9	3	-1	-11	-2
-175	-8	6	3	3	-12	-3
-150	-11	2	10	5	-10	1
-125	-8	-2	10	2	-8	1
-100	8	-10	15	0	1	-2
-75	1	6	5	-4	-4	-2
-50	1	6	5	-7	-2	0
-25	5	3	1	-3	1	0
0	1	3	-1	-1	1	-4

Note: The Democratic poll share minus the actual vote share is shown for selected days of the election cycle.

tion campaign, more than half of the poll variance was due to sampling error alone; less than half was real. (And keep in mind that this estimate is based only on allowing for sampling error, without also taking into account house and design effects.)

Now, let us shift our focus and consider the pattern over time, across the 15 elections. Figure 4 shows selected readings of the poll of polls over 25-day intervals for six election years: 1944, 1952, 1968, 1976, 1980, and 1992. (These were chosen for expository purposes. Data for the full set of elections can be found on the Nuffield College website at http://www.nuff.ox.ac.uk/Politics/papers/2002/w27/wlezien.pdf.)

To provide comparability across years, we subtracted out the actual Democrat share of the two-party vote; thus, the numbers reflect the degree to which poll results differed from the final vote. For days after the last poll in particular election years, we assumed the numbers from the last pre-election poll. For each of the many days without polls centered on those dates, we interpolated from the most recent date with polls and the next date with polls.

The data in Figure 4 reveal that election outcomes come into focus as the election cycle evolves. At the beginning of our timeline, 200 days before the election, reported poll results differ from the final vote quite a lot—by about 6.4 points on average for the six elections shown (6.5 points for the full set of elections).

The differences remain remarkably stable over the next 100 days, and then diminish dramatically thereafter. At the 100-day mark, about the time of the national party conventions, the average difference between the polls and the ultimate vote is 6.1 points. By the very end of the campaign, the average difference is a mere 1.7 points (2.2 points for all 15 elections; and even this estimate is inflated, because the final pre-election polls often end well in advance of the election, and we simply carried forward the results.)

Clearly, the polls tell us more and more about the outcome as the campaign unfolds. This is not especially surprising. What may be surprising, though, is that much of the improvement in predictability occurs during the general election campaign after Labor Day. Indeed, the relatively small real movement in preferences during the fall appears to matter most.

Another interesting pattern in the polls over time is displayed in Figure 5. It shows the poll lead for the ultimate winner in each election over the last 200 days of the campaign. We can see that the early polls tell us something about the Election Day result—at 200 days out, the winner had the poll lead in four of the six elections shown (10 of the full 15 elections). Important sorting continued through the summer, and by 100 days out, the winner had the lead in all six elections in the figure and 12 of the 15 elections since 1944; by Labor Day, this was true in all but one year, 1948. (For the purposes of our analysis, we included Al Gore among the winners, as he won the popular vote in 2000.)

Once in the lead, however, the winner's margin tended to shrink. Leads from the Labor Day period, for instance, eventually were halved by Election Day. Even the final pre-election polls tended to distort the lead, though this partly reflects the lack of late polls in a number of years, as noted above. These results imply an underdog effect, where

Figure 5

Winner's Lead Shrinks as Election Day Nears

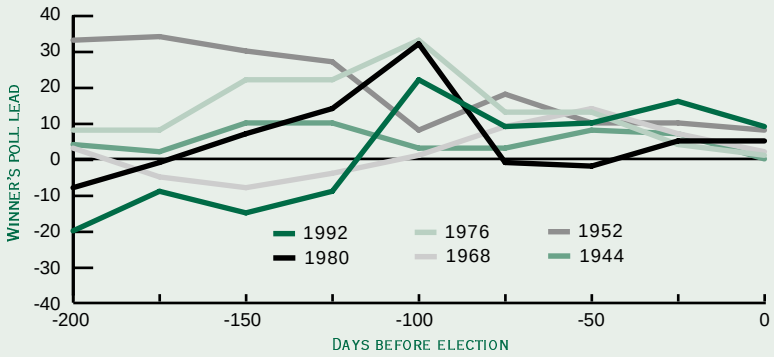

Days till election	WINNER'S POLL LEAD 1992	1980	1976	1968	1952	1944
-200	-20	-8	8	3	33	4
-175	-9	-1	8	-5	34	2
-150	-15	7	22	-8	30	10
-125	-9	14	22	-4	27	10
-100	22	32	33	1	8	3
-75	9	-1	13	9	18	3
-50	10	-2	13	14	15	8
-25	16	5	4	7	10	7
0	9	5	1	2	8	0

Note: The ultimate winner's lead in the polls is shown for selected days of the election cycle.

the projected loser gains support as the campaign persists. With such an effect, the drift of the polls is as if one candidate emerges as the favorite after the conventions and then watches the lead shrink. (A similar pattern also is evident in congressional polls.)

The basic patterns of poll movement constitute a puzzle of sorts. That the polls increase in accuracy leading up to Election Day indicates that something happens to change voter preferences, and in meaningful ways. Indeed, it appears that election campaigns really do matter.

We start with the knowledge that because the electorate's preferences do change, campaign events (broadly defined) must be exerting some sort of impact. The question then is whether these shocks from campaign events take the form of temporary "bounces" or permanent "bumps." Simply put, do the effects decay or else last?

If campaign effects are bounces, they dissipate over time. Preferences tend to revert to an "equilibrium" that is set early in each particular election year. The final outcome is the simple sum of this equilibrium plus the effects of very late events that do not fully dissipate before Election Day.

If campaign effects are bumps, conversely, they last to affect the outcome. In effect, the equilibrium drifts over time. The election outcome is the sum of all the bumps—often small in size—that happen during the campaign.

The answer may be that campaign events produce both bounces and bumps. It may be that some effects dissipate and others last. It may be that the effects combine both bounces and bumps. Statistically, it is the bumps and not the bounces that matter in the long run. They cumulate over time.

We see the evidence of permanent bumps in the fact that the polls are increasingly accurate over the fall general election campaign. If this were not so—if the effects of events dissipated—the accuracy of polls would vary little except at the very end of the campaign, reflecting the effects of late events. Something clearly happens during the fall to change voters' preferences.

Beforehand, at least prior to the conventions, we see a very different pattern. The polls during this period do not vary much in their accuracy; indeed, it is as if they bounce around an equilibrium that is constant for the particular election. The conventions have important, often realigning, effects.

The fall campaign then generates change as the accumulation of seemingly small bumps for one candidate or the other. We do not know what exactly causes preferences to change during this period. We also cannot predict it in advance.

Now, what explains the shrinking margins? Recall that trial-heat surveys show considerable movement early in the campaign, often with one candidate surging to a large lead. As the campaign progresses, the electorate's net vote intention hardens, typically moving in the direction of a tightening outcome.

This result is as if people make tentative choices early, based on the political news of the moment. In other words, early in the campaign, survey respondents act in a relatively nonpartisan or independent manner. Responding to the prevailing news about the candidates, the early campaign electorate drifts toward the early favorite, much in the manner of "independent" voters generally.

As the campaign evolves, much of these effects dissipate and preferences polarize, with a widening attitudinal gulf between the supporters of the two major-party candidates. It may be that the campaign activates voters' predispositions, causing them to gravitate toward their partisan "equilibrium" or some broader underlying preference, or it just may be that individuals react differently to the events of the campaign.

As we have discussed in an article that appeared in the September 2001 issue of *American Politics Research*, the polarization of underlying preferences over the campaign will produce a predictable decline in poll margins, regardless of its particular underpinnings. This, of course, is exactly what we observe.

In short, shifting poll results often represent chance variation due to survey error. We nevertheless can see beneath the surface that the electorate's preferences change over the course of a campaign. Early on, the likely winner holds a large initial lead. As the campaign unfolds, the race typically tightens and becomes more stable as preferences harden.

The polls also increase in accuracy leading up to Election Day. These patterns tell us that election campaigns do matter and that the general election campaign matters most of all. They do not tell us *how* campaigns actually matter. What events had effects? Which ones lasted and which ones decayed? We simply do not know. This mystery remains unsolved.

The authors wish to thank Joseph Bafumi, Christopher Carman, Bruce Carroll, Jenny Hill, Joe Howard, Jeff May, Adrian Shepherd, and Amy Ware for assistance with data collection over the years, and Steve Fisher for comments. The larger project has been supported by grants from the National Science Foundation (SBR-9731308 and SBR-0112856) and the Institute for Social and Economic Research at Columbia University.