

**The Evolution of Electoral Preferences:
What the Polls Reveal as the Campaign Unfolds**

Christopher Wlezien
University of Oxford

and

Robert S. Erikson
Columbia University

Nuffield College Working Paper 2003-W7
University of Oxford

Forthcoming in Allan McCutcheon, ed., *The Science of Pre-Election Polling*:
Lincoln: University of Nebraska Press.

A preliminary version of this paper was presented at the Symposium on the Science of Pre-Election Polling, Washington, D.C., 2002. Portions also were presented at the University of Essex, Nuffield College, Trinity College, Dublin, and in the seminar series of the Houston Chapter of the American Statistical Association. The research described herein builds on a larger project that 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. Special thanks are due to Bruce Carroll, Joe Howard, Jeff May, and Amy Ware for assistance with data collection and management and Charles Franklin, Don Green, Michael Marsh, and Allan McCutcheon for useful comments on the data analysis itself.

Electoral scholars are beginning to understand how electoral preferences evolve over time. Much is known about how voters make their choices on Election Day (see, e.g., Alvarez, 1997; Campbell, 2000; Gelman and King, 1993, Johnston, Brady, and Crete, 1992; Lewis-Beck, 1988) and why election outcomes turn out as they do (see, e.g., Campbell and Garand, 2000; Sanders, 1996). Of course, what we do know is imperfect.¹ Much less is known, however, about how voter preferences come into focus over the campaign (Campbell, 2000; Wlezien and Erikson, 2002). One major reason is that voter preferences during the campaign are incompletely observed. Poll readings are not available for every day of the campaign, and the readings that are available are measured with survey error. How then can we trace the evolution of electoral preferences over the timeline of the campaign? What do the polls reveal about what is happening as the campaign unfolds before us?

This challenge is the subject of the present chapter, focusing on one contest: the U.S. presidential election of 2000. The goal is to extract the underlying electoral preference over the election cycle using the full set of available polls, that is, in the presence of multiple survey organizations and designs. The approach builds on previous research in which poll results from various organizations were combined into one time series, that is, to capture the history of voter preferences during the 1996 election (Erikson and Wlezien, 1999). Using this technology, it is possible to partially disentangle preferences from poll results as the campaign actually evolves. The concluding section considers what polls during the campaign reveal about the ultimate vote on Election Day.

¹ Witness the US presidential election of 2000.

The Polls in 2000

There were an almost incredible number of pre-election polls during the 2000 presidential election cycle. For the 2000 election year itself, the pollingreport.com website contains some 524 national polls of the Bush-Gore (-Nader) division reported by different survey organizations. In each of these polls, respondents were asked about how they would vote “if the election were held today” with slight differences in question wording. The differences would appear to matter relatively little (see Lau, 1994). Of course, there are other differences in the practices of survey organizations and these are quite important, as we will see.

Figure 1 displays the basic data. It shows Gore’s percentage share of the two-party vote (ignoring Nader and Buchanan) for the complete set of 524 polls. Each observation represents the “vote” from a specific poll. 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.² Where multiple results for different universes are available for the same polling organizations and dates, data for the universe that best approximates the actual voting electorate is used, e.g., a sample of likely voters over a sample of registered voters. Wherever possible, respondents who were undecided but leaned toward one of the candidates were included in the tallies.

-- Figure 1 about here --

The data in Figure 1 offer only a very sketchy pattern. There is evidence of the (fairly) predictable convention bounces, first for Bush (about 100 days before Election Day) and then for Gore, after which Gore’s support shrinks, though there appears to be a lot a variation in poll results on particular days. What is most evident from the figure is the clustering of polls late in the election cycle, which is as we might expect.

The numbers in Figure 1 are somewhat deceiving in that there is considerable overlap in the polls—typically tracking polls—conducted by the same survey houses for the same reporting organizations. That is, many survey organizations report results on a daily basis that average results from preceding days, so that results on successive days are not independent. Where an organization operates a tracking poll and reports 3-day moving averages on each day, for example, the result reported on day t contain information from days $t-1$, $t-2$, and $t-3$. Clearly, it is necessary to remove this overlap in polls and doing so is straightforward. Using the example from above, where a survey organization operates a tracking poll and reports a 3-day moving average on each day, one simply selects poll results for every third day.

Following this procedure from the beginning of the election year to Election Day leaves 295 separate national polls. These polls allow readings for 173 separate days during 2000, 59 of which are after Labor Day. We thus have a virtual day-to-day monitoring during the general election campaign. It is important to note, however, that polls on successive days still are not truly independent: Although they do not share respondents they do share overlapping *polling periods*. That is, polls results centered on a particular day t usually include information collected on adjacent days, $t-1$ and $t+1$.³ Thus, results on neighboring days capture many of the same things. This is of consequence for any analysis of preference change.

-- Figures 2 and 3 about here --

The data in Figure 2 indicate patterned movement in the polls over time, though the portrait of preference movement remains impressionistic. Most importantly, the results for any given date differ quite considerably. Some of the noise is mere sampling error, though there are

² For surveys in the field for an even number of days, the fractional midpoint is rounded up to the following day.

³ There is a good amount of variance in the number of days surveys are in the field, ranging from 1 day snapshots to surveys that stretch up to 31 days. The mean number of days in the field is 3.57; the standard deviation is 2.39 days.

other sources of survey error as well. The daily poll-of-polls in Figure 3 reveals more distinct pattern. The observations in this figure represent Gore's share for all respondents aggregated by the mid-date of the reported polling period. We see more clearly that Gore began the year well behind Bush and gained through the spring, where his support settled at around 47 percent until the conventions. We then see the expected convention bounces, from which Gore emerged in the lead heading into the autumn. According to these polls, things were playing out as political science election forecasters expected they would, and much like 1988 (see Wlezien, 2001).⁴ But by the start of October, the parallel with 1988 had ended, with Gore's support in the polls fading until just before Election Day, when it rebounded sharply. The polls in the field at the very end of the campaign showed a dead-heat.

Disentangling Preferences

Trial-heat poll results represent a combination of true preferences and survey error. Survey error comes in many forms, the most basic of which is sampling error. All polls contain some degree of sampling error. Thus, we will observe changes from poll to poll even when the division of candidate preferences is constant and unchanging. This is well known. The problem is that we cannot actually separate sampling error from reported preferences. Sampling error is random.

All survey results also reflect design effects. These effects represent the consequences of the departure in practice from simple random sampling that results from clustering, stratifying, and the like (Groves, 1989). When studying election polls, the main source of design effects surrounds the polling universe. It is not easy to determine in advance who will vote on Election Day. All we can do is estimate the voting population. Survey organizations typically rely on

⁴ The sitting Vice-President is running, the economy and presidential approval are favorable; he is behind in the polls

“screens” to identify likely voters. How one screens for likely voters has consequences for the poll margins one reports at each point in time. It also may have implications for the variance in the polls over time.⁵ Many organizations also use some sort of weighting procedure, e.g., a particular distribution of party identification, to approximate the likely voting electorate. To the extent the organizations use variables that predict the vote on Election Day, weighting will reduce the observed sampling error. To the extent the weighting schemes differ, however, it may introduce additional error, as when organizations weight by different distributions of party identification (Wlezien and Erikson, 2001).

When combining polls from different survey organizations, house effects also are a problem. These effects represent the consequences of survey houses employing different methodologies, including survey design itself. Indeed, much of the observed difference across survey houses may reflect underlying differences in screening and weighting procedures. Results can differ across houses for other reasons, including data collection mode, interviewer training, procedures for coping with refusals, and the like (see Converse and Traugott, 1986; Lau, 1994; also see Crespi, 1988). Whatever the source, poll results can vary from day to day because polls reported on different days are conducted by different houses.

Although we cannot eliminate sampling error, we can to some extent account for design and house effects. That is, we can control for the polling universe—at least very broadly defined—as well as different survey houses (Erikson and Wlezien, 1999). The data include results from 36 different survey organizations. These organizations sampled three different polling universes during the year, specifically, adults, registered voters, and likely voters. Of

early in the year, and then gains the lead for good after the party convention.

⁵ There actually is reason to suppose that a basic sample of registered voters better may be better, that is, since it is less artificially variable (Wlezien and Erikson, 2001).

course, as noted above, these universes do not necessarily mean the same things to different organizations, particularly as relates to “likely” voters. The problem is that we do not have much information on the various likely voter screens (or weighting schemes) that different organizations use.⁶ In order to adjust for possible design and house effects given what little we do know, the results for all 295 polls are regressed on a set 35 ($N - 1$) survey house dummy variables and two polling universe dummy variables. Dummy variables also were included for each of the 173 dates with at least one survey observation. With no intercept, the coefficients for survey dates comprise the adjusted estimates of public preferences during the campaign.

-- Tables 1 and 2 about here --

The results of the analysis are shown in Table 1. Here, we can see that the general polling universe did not meaningfully affect poll results during 2000, controlling for survey house (and date). That is, the choice of “likely” versus “registered” voters did not affect the results. This is not entirely surprising, as the same was true in 1996 (see Erikson and Wlezien, 1999).⁷ Table 1 also shows that survey house did matter in 2000, though the effects are not highly reliable over the full election year. The estimated range of the house effects is over six percentage points, suggesting important differences among at least some survey houses (also see Traugott, 2001). After Labor Day, house effects are much more statistically significant and the range of effects is greater, approximately eight percentage points. Presumably these differences at least partly reflect differences in survey design, although we cannot be sure about the exact causes. Regardless, the evident differences have consequences for our portrait of preferences during 2000: Poll results will differ from day-to-day merely because different houses report on different days.

For illustrative purposes, Table 2 lists the estimated effects of selected, well-known houses

⁶ What information we have indicates that they matter quite a lot (Wlezien and Erikson, 2001).

over the full election year and after Labor Day. (Recall that our data include polls from 36 different survey organizations.) The Table displays the degree to which poll results from the different houses overstate or understate Gore's poll share compared to the median house, which was ABC.⁸ We can see that most of the organizations shown tended to overstate Gore's share relative to the median house, especially over the full election year. The important lesson is that poll results differ quite dramatically even among those houses that are within the mainstream. The differences make clear why we can't (or shouldn't) rely on a single survey house to gauge change in voters' preferences. Which one should we pick? Were we to pick CNN/Time, the results would overstate Gore's share by more than three percentage points on average compared to the base survey house. Were we to pick Voter.com, the results would understate Gore's share by over one point. By relying on data from different houses, we have a basis for comparison, one that we can explicitly take into account in our analyses. Of course, it also gives us more poll readings and larger daily samples.

Notice from our analysis of variance in Table 1 that the collective effects of the survey date easily meet conventional levels of statistical significance ($p < .001$) during both the full election year and the general election campaign after Labor Day. This is statistical evidence that underlying electoral preferences actually changed over the course of the campaign—a simple fact of obvious importance. Figure 4 displays the survey date estimates, which constitute our (house-adjusted) estimates of preferences on different days. Of course, because of the differences across survey houses, it was necessary to select a “base” house. For this analysis, the median house from the analysis of variance for the full election year was used. This seems an appropriate strategy, though it may not be quite right: the problem is that we cannot tell for sure. It may be tempting to

⁷ Again, this does not mean that the differences between the different likely voter universes did not matter.

use the house that best predicted the final outcome, though this is even more tenuous. Given that all polls contain sampling error, getting it right at the end is as much the result of good luck as good survey design.

-- Figures 4 and 5 about here --

The estimates in Figure 4 exhibit more pattern than the unadjusted polls in Figure 2, especially late in the cycle. Figure 5 displays results for the post-Labor Day period. There still is evidence of noise in these estimates the result of sampling error. That is, the polls bounce around quite a lot from day to day. There also is reason to think that the series contains other sources of error, such as design effects, that are not easily captured statistically because they vary over time (Wlezien and Erikson, 2001). This might help explain some of the sharper spikes in poll results during the early fall. The problem is that we do not know: As noted above, we do not have much information about the various screens and weighting procedures used by the many different survey organizations. We thus must accept that our series of polls, even adjusted for systematic house and design effects, is imperfect. It may be the best we can do, however.

It also may be possible to address sampling error, at least to some extent. Although we cannot fully eliminate sampling error, it is possible to adjust for it somewhat by “smoothing” the data. Donald Green and Alan Gerber have developed such a routine, called *Samplemiser*. The routine uses the Kalman filtering algorithm (see Beck, 1990) to generate new estimates from previous poll results (weighted by their sample sizes) and then corrects based on systematic variation in the results. In effect, it adjusts for the negative autocorrelation in poll results that random sampling error will produce. (For details, see Green, Gerber, and DeBoef, 1999.) Let us

⁸ Coincidentally, the median *poll* was conducted by ABC/Washington Post.

see what the Samplemiser adjustment provides.⁹

-- Figure 6 about here --

The resulting estimates and the raw house-adjusted poll-of-polls are shown in Figure 6. What we can see here is that smoothing adds only a little clarity. During the fall campaign, it makes almost no difference whatsoever. This does not mean that our series of polls does not contain sampling error; Rather, it means that the systematic portion of sampling error that we can detect is relatively small. The evident differences owing to smoothing even may be somewhat generous given that the procedure can capture the effects of shocks to true preferences that are very short-lived, and look much like the effects of random sampling error. Regardless, it appears that our house-adjusted estimates are pretty reliable.

What Polls Reveal about Preferences as the Campaign Unfolds

The foregoing analysis is entirely historical. It tells us what happened after the fact, using data that are available after the election is over. What do the polls tell us in advance, as the campaign itself unfolds? We obviously cannot step back in time. We can, however, behave as if we could. That is, we can produce estimates at each point of the cycle using information about poll results up to that date. Specifically, extending the preceding analyses, house-adjusted estimates of preferences can be generated for each day of the election cycle using polls for the particular day and all preceding days. Doing so seems straightforward. The only question is: Which house do we select as our base? This is of obvious importance, as we have seen, for the

⁹ The Samplemiser routine requires that one make a number of decisions about the nature of the underlying time series process, the variance of disturbances or shocks to preferences, and priors about a number of things. The estimates for this analysis are based on the assumption of an integrated or 'random walk' process, an unknown disturbance variance, and diffuse priors. The procedure presumes that the variance of shocks to preferences is approximately constant over time, though this may not be true (see Wlezien and Erikson, 2002). It thus is important to note that generating estimates for different periods during the year, e.g., pre- and post-Labor Day, produces virtually

decision affects the estimates of underlying electoral preferences in substantial ways. After the fact, the decision is relatively easy: Pick the median house. But what does one do during the campaign? There would seem to be (at least) two criteria for deciding: (1) the past performance of survey houses; (2) their presence in the field throughout the cycle, that is, from beginning to end. Based on these criteria, the obvious choice seems Gallup. It has the longest continuous record of pre-election polling. It also was in the field throughout the campaign and polled much more frequently than any other organization.¹⁰

To generate the house-adjusted estimates of electoral preferences at each point in time, the ANOVA model described in Table 1 is estimated for each date for which we have poll data. For instance, on day *ED-200*, the effects of survey date, survey house and general poll universe are estimated using all poll results for that day and all preceding days, i.e., *ED-200*, *ED-201*, *ED-202*, and so on, using Gallup as the base house. The date estimate for *ED-200* is recorded and the process is repeated for day *ED-199* and each successive day for which there are poll data up to Election Day. These rolling house-adjusted estimates are shown in Figure 7 along with the original house-adjusted poll readings for the last 200 days of the campaign.¹¹

-- Figures 7 and 8 about here --

In the figure we can see fairly substantial disjunctures between the two series early on, until about 100 days before Election Day. This is to be expected given that house adjustments during this period are based on relatively small numbers of cumulative surveys, and thus are not highly reliable. After that point, the two series closely track each other, with one notable

identical estimates.

¹⁰ Fully 43 of the 295 separate national polls used in this analysis are from Gallup, 50 percent more than from any other organization.

¹¹ Of course, we could change (or “update”) the base survey house as the campaign unfolds and we learn about the differences across houses. For instance, heading into the fall general election campaign, we might choose the median house based on analysis using polls from the pre-Labor Day period.

exception about 70 days out. Indeed, through the post-Labor Day period, the rolling estimates virtually parallel the fully adjusted series based on data available after the election. This is clear in Figure 8. The pattern is satisfying: Although we expect the two series to increasingly correspond as we move closer to Election Day, we do not necessarily expect such a high level of consistency throughout the fall.¹² The rolling estimates are track slightly below the original house-adjusted readings, reflecting the smaller Gallup baseline. Recall that Gallup understated Gore's vote share by 0.42% on average compared with the median house during the fall (see Table 2).

It is tempting to treat the estimates in Figures 7 and 8 as “forecasts” of sorts but they really are not. The estimates are based on polls aggregated by the mid-date of the reported polling period, e.g., the estimate for day *ED-30* is based on polls centered on that date. Of course, the data are not available at that time, when polls are still in the field. The data are only available when the polling is complete and there often is a lag in the actual release. To produce poll forecasts for each day, therefore, it is necessary to date polls by the release date itself. Unfortunately, this information is not readily available now, in the autumn of 2002. We can, however, guess at the release date. For instance, we might assume that the release date is the date after the last day the poll is in the field. While reasonable, this is a tenuous assumption. It may make more sense to assume that a poll is released in the two days after coming out of the field, that is, by “pooling” data on each pair of days. Polls ending on *ED-30* can be dated as *ED-29* and *ED-28*, and polls ending on *ED-29* can be dated as *ED-28* and *ED-27*, and so on. The estimate for each day *t* thus would be the equivalent of an average forecast using polls ending on days *t-1* and *t-2*. This seems a reasonable approach given that we are forecasting after the fact, without

¹² The pattern implies that estimated house effects do not change much after Labor Day or that combining the often

information on poll release; in the future, we can do it right. Now, to generate the estimate for each day, we simply follow the estimation algorithm used to produce the rolling estimates (in Figure 7) using the newly dated polls.

-- Figure 9 about here --

Figure 9 plots the estimate and a 95% confidence interval for each date in the post-Labor Day period. The confidence interval is based on the daily sample size and the poll margin itself.¹³ Notice that the forecasts are less variable than the fully house-adjusted estimates or the rolling estimates (see Figure 8), which is the simple result of pooling. The forecasts together with the confidence intervals reveal additional information about what we “could” see as the fall campaign evolved. First, given that the confidence drifts above the 50-50 threshold, it appears fairly certain that Gore entered the fall with a lead, which quickly disappeared. Second, it also appears that Bush really did gain the lead during the last 30 days, though his support ebbed and flowed somewhat over the period, i.e., on some days, we could not say with great (95%) confidence that he was “winning.” Third, in the days leading up to the election, we could forecast a dead heat only at the very last minute, on Election Day itself (and just barely at that). This implies one fairly obvious lesson: In a close race, stay in the field until the very end. The polls that did get it right.

What Polls Reveal about the Vote

This chapter has focused on what polls tell us about electoral *preferences* over time. Specifically, it has provided a basic method for extracting elements of survey error from poll

multiple readings on each day during the period largely cancels out the evident differences across houses.

¹³ Since poll results on each pair of successive days are pooled, the sample size used to generate the confidence interval is the pooled sample divided by 2. Note that the confidence intervals technically are affected slightly by the loss of degrees of freedom from estimating house and date effects, and we do not take these into account.

results reported by different survey organizations as the election cycle evolves. This may help us determine what the likely vote would be on any given day. In one sense, it also provides rolling forecasts of the vote on Election Day, as the foregoing discussion implies. That is, we might expect that if nothing happens between the poll date and the election, the estimate from the polls would be the final vote. This is a fairly standard approach to forecasting using the polls, certainly in media coverage of the campaign, and it also is fairly reasonable. It nevertheless relies on strong assumptions about dynamics. That is, it presumes that all changes in preferences persist.

If innovations in preferences don't persist, however, this is not right. The best forecast would not be the current estimate from the polls: It would be the equilibrium of the series (if knowable), except at the end of the race, when any shocks to preferences would not have time to decay.¹⁴ This is what most political scientist forecasters presume, either explicitly or implicitly (see Campbell and Garand, 2000). In their models, the economic and political “fundamentals” of the election are considered to be constant at least for much of the campaign.

Now, it may be that neither model is correct. Indeed, we might expect that that some shocks persist and others decay (Wlezien and Erikson, 2002). Certain events may cause permanent “bumps” and others temporary “bounces.” Yet other events may cause some of both. In theory, we can model the underlying process as the campaign evolves, much as with the rolling house adjustments. That is, for each day t , we can estimate a basic (ARIMA) time series model that allows for both persistence and decay. We then can generate a forecast (or set of forecasts) based on the results. Although this is straightforward in theory, it is much less so in practice. As we have seen, poll estimates on successive days are not truly independent—they share overlapping polling periods and thus contain much of the same information. This poses obvious

¹⁴ The degree to which this is true, of course, depends on the size of shocks and their rate of decay.

complications: We simply can't pursue standard time series analyses.

Even armed with perfect data and a good model, there are limits on our ability to forecast using pre-election poll. Most importantly, we cannot (fully) anticipate campaign effects and whether and to what extent they last. All we can offer at any point in time are "mean" estimates of sorts. And, of course, no matter how well we design our surveys, comprehend the underlying dynamics of voter preferences, and anticipate the campaign itself, we ultimately can't predict whose votes will and won't count on Election Day.

References

- Alvarez, R. M. (1997). *Information and Elections*. Ann Arbor: University of Michigan Press.
- Beck, Nathaniel. 1990. 'Estimating Dynamic Models Using Kalman Filtering.' *Political Analysis* 1: 121-156.
- Campbell, A., Converse, P.E., Miller, W.E., & Stokes, D.E. (1960). *The American Voter*. New York: Wiley.
- Campbell, J.E., (2000). *The American Campaign: U.S. Presidential Campaigns and the National Vote*. College Station, Texas: Texas A&M University Press.
- , & Garand, J.C. (2000). *Before the Vote: Forecasting American National Elections*. Thousand Oaks, Calif.: Sage Publications.
- Converse, P.E., & Traugott, M.W. (1986). 'Assessing the Accuracy of Polls and Surveys.' *Science* 234:1094-1098.
- Crespi, I. (1988). *Pre-Election Polling: Sources of Accuracy and Error*. New York: Russell Sage.
- Erikson, R.S., & Wlezien, C. (1999). 'Presidential Polls as a Time Series: The Case of 1996.' *Public Opinion Quarterly* 63:163-177.
- (1996). 'Of Time and Presidential Election Forecasts.' *PS: Political Science and Politics* 29:37-39.
- Erikson, Robert S., and Lee Sigelman. 1996. 'Poll-Based Forecasts of the House Vote in Presidential Election Years.' *American Politics Quarterly* 24:52-531.
- , 1995. 'Poll-Based Forecasts of Midterm Congressional Elections: Do the Pollsters Get

- it Right?' *Public Opinion Quarterly* 59:589-605.
- Gelman, A., & King, G. (1993). 'Why are American Presidential Election Polls so Variable When Votes are so Predictable?' *British Journal of Political Science* 23: 409-519.
- Green, Donald P., Alan S. Gerber, and Suzanna L. DeBoef. 1999. 'Tracking Opinion over Time.' *Public Opinion Quarterly* 63: 178-192.
- Groves, R.M. (1989). *Survey Errors and Survey Costs*. New York: Wiley.
- Heise, D.R. (1969). 'Separating Reliability and Stability in Test-Retest Correlations.' *American Sociological Review* 34:93-101.
- Holbrook, T. (1996). *Do Campaigns Matter?* Thousand Oaks, California: Sage.
- Johnston, R., Blais, A., Brady, H.E., & Crete, J. (1992). *Letting the People Decide: Dynamics of a Canadian Election*. Kingston, Canada: McGill-Queen's Press.
- Lau, R. (1994). 'An Analysis of the Accuracy of 'Trial Heat' Polls During the 1992 Presidential Election.' *Public Opinion Quarterly* 58:2-20.
- Lazarsfeld, P. F., Berelson, B.R., & Gaudet, H. (1944). *The People's Choice*. New York: Columbia University Press.
- Lewis-Beck, Michael S. (1988). *Economics and Elections*. Ann Arbor: University of Michigan Press.
- Sanders, David. (1996). 'Economic Performance, Management Competence and the Outcome Of the Next General Election.' *Political Studies* 44:203-231.
- Traugott, Michael W. 2001. 'Assessing Poll Performance in the 2000 Campaign.' *Public Opinion Quarterly* 63:389-419.
- Wlezien, C. (2001). 'On Forecasting the Presidential Vote.' *PS: Political Science*

And Politics 34:25-31.

----- (2000). 'An Essay on "Combined" Time Series Processes.' *Electoral Studies* 19:77-93.

Wlezien, C., and Erikson, R.S. (2002). 'The Timeline of Presidential Election Campaigns.'

Journal of Politics 64:969-993.

----- (2001). 'Campaign Effects in Theory and Practice.' *American Politics Research*

29:419-437.

Zaller, J. (2002). 'Assessing the Statistical Power of Election Studies to Detect

Communication Effects in Political Campaigns.' *Electoral Studies* 21:297-329.

Figure 1: All Trial-Heat Presidential Polls by Date, 2000

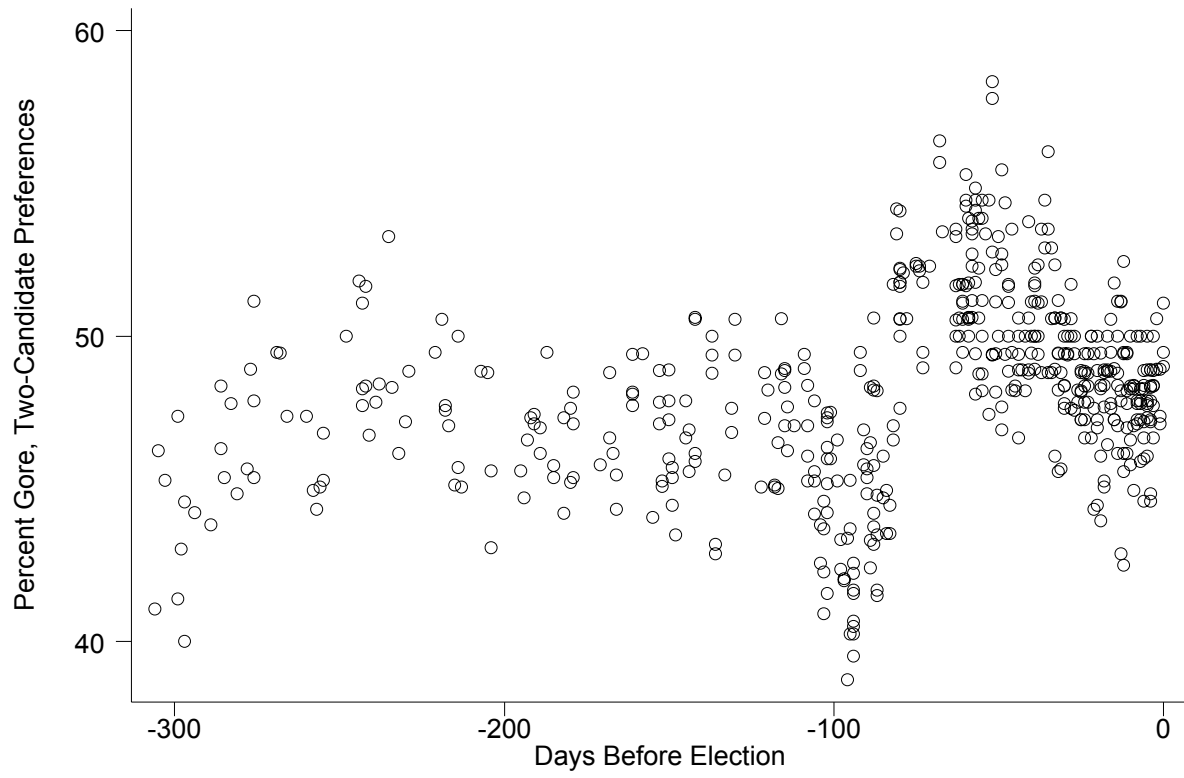


Figure 2: All Separate Trial-Heat Presidential Polls by Date, 2000




Figure 3: Trial-Heat Presidential Polls Aggregated by Date, 2000

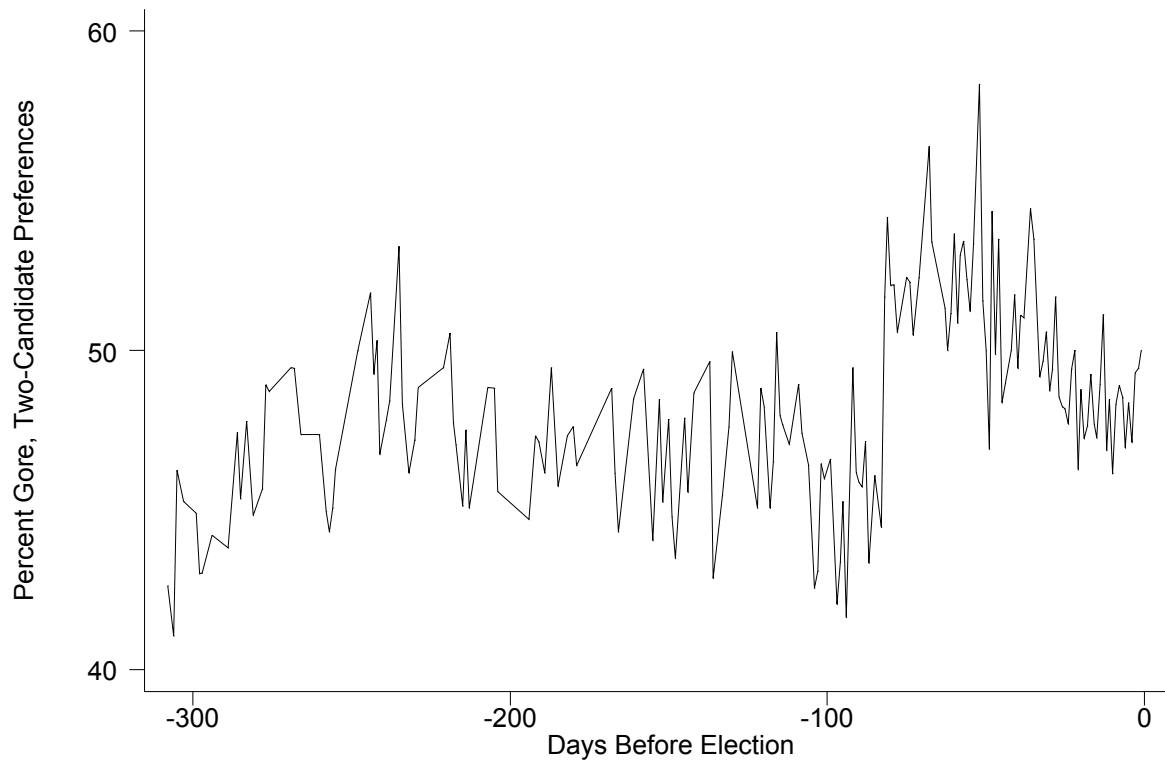


Table 1: An Analysis of General Survey Design and House Effects on
 Presidential Election Polls, 2000

Variable	Election Year	After Labor Day
Poll Universe	0.12 (0.89)	0.08 (0.78)
Survey House	1.44 (0.09)	2.15 (0.01)
Survey Date	3.88 (0.00)	4.31 (0.00)
<i>R</i> -squared	0.90	0.88
Adjusted <i>R</i> -squared	0.68	0.68
Mean Squared Error	2.79	1.74
Number of polls	295	135
Number of respondents	267,974	130,024

Note: The numbers corresponding to the variables are F-statistics. The numbers in parentheses are two-tailed *p*-values.

Table 2: Estimated Effects of Selected Survey Houses on
Trial-Heat Poll Results, 2000 (in percents)

Polling Organization	Election Year	After Labor Day
Yankelovich/CNN/Time	3.01	3.10
Fox/Opinion Dynamics	2.07	2.62
Pew Research Center	1.26	0.57
CBS	1.18	1.46
Newsweek	1.15	1.76
NBC/Wall Street Journal	0.78	-.12
ABC/Washington Post ^a	0.71	-.19
Reuters/MSNBC	0.69	0.70
ICR ^b	0.64	1.14
ABC ^c	0.00	0.00
CNN/USA Today/Gallup	-.53	-.42
Voter.com/Battleground	-1.12	-1.59

^a Prior to separate ABC and Washington Post tracking poll reports.

^b International Communications Research.

^c Based only on tracking poll reports.

Source: Pollingreport.com.

Figure 4: House-Adjusted Trial-Heat Presidential Polls by Date, 2000

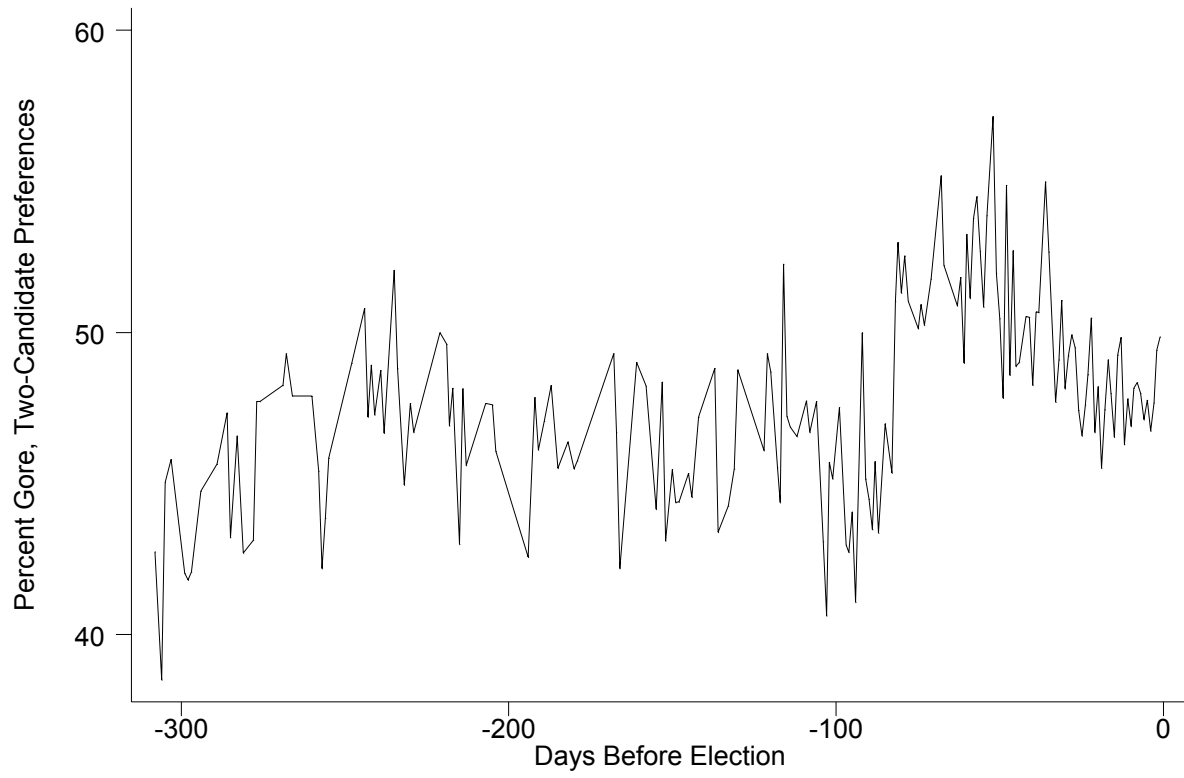


Figure 5: House-Adjusted Trial-Heat Presidential Polls by Date,
Labor Day to Election Day, 2000

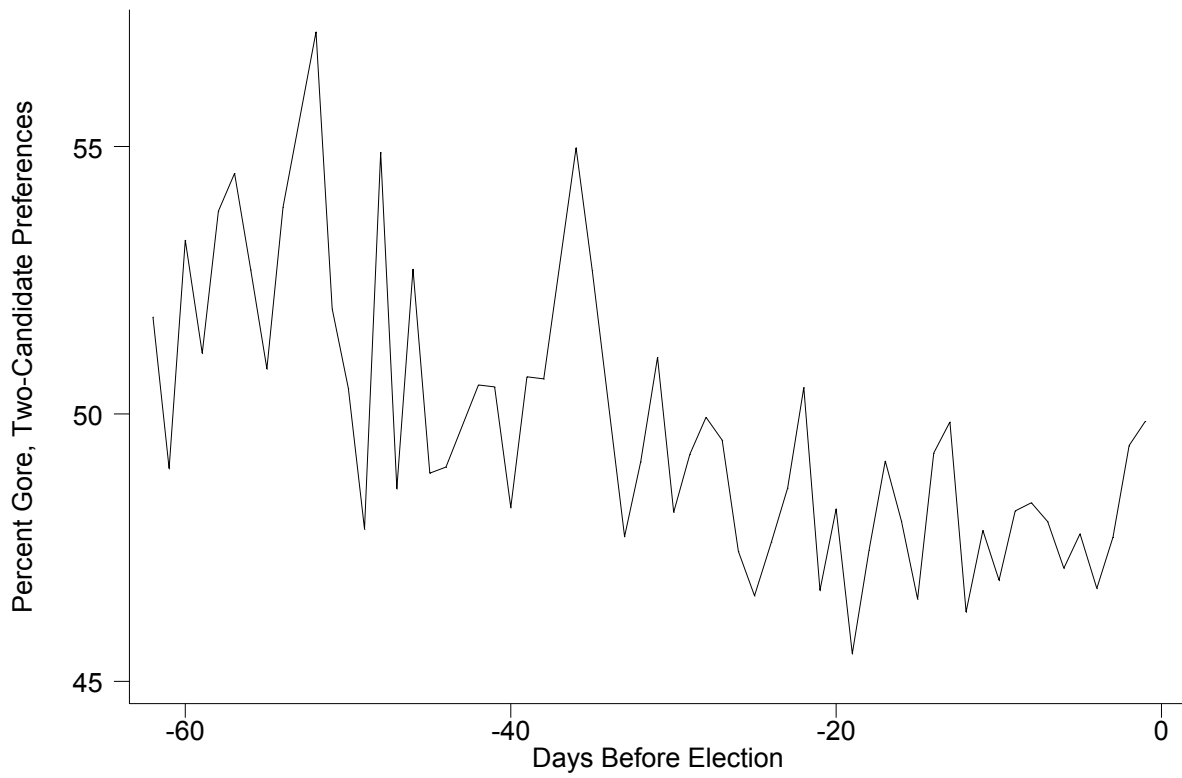


Figure 6: House-Adjusted Poll Readings and Sampliser Estimates, 2000

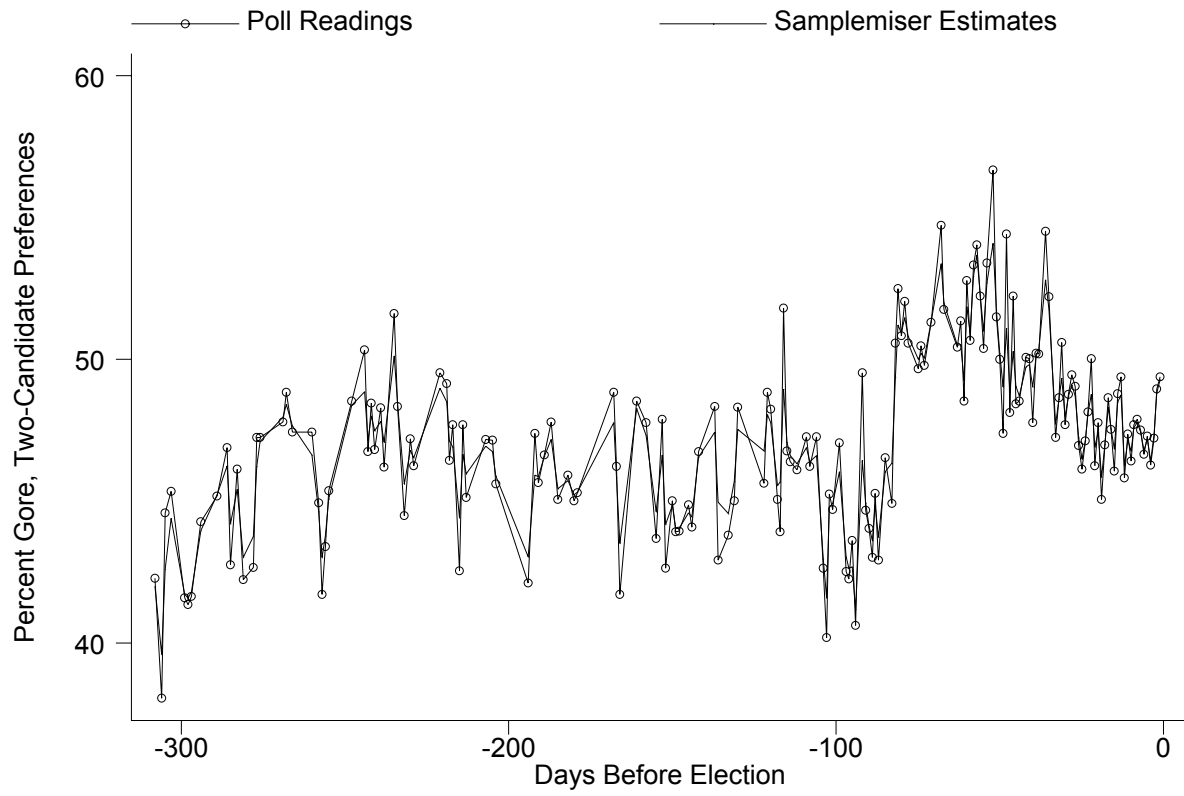


Figure 7: House-Adjusted Poll Readings and Rolling House-Adjusted Estimates, 2000

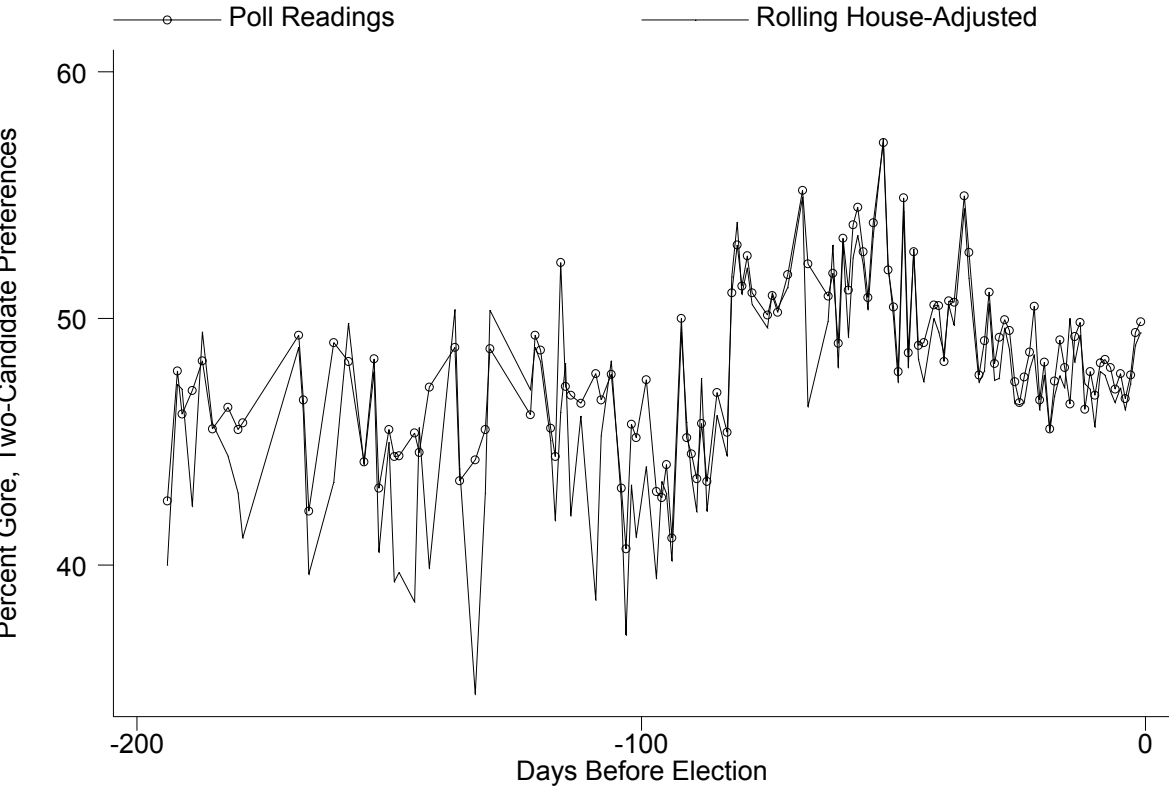


Figure 8: House-Adjusted Poll Readings and Rolling House-Adjusted Estimates, Labor Day to Election Day, 2000

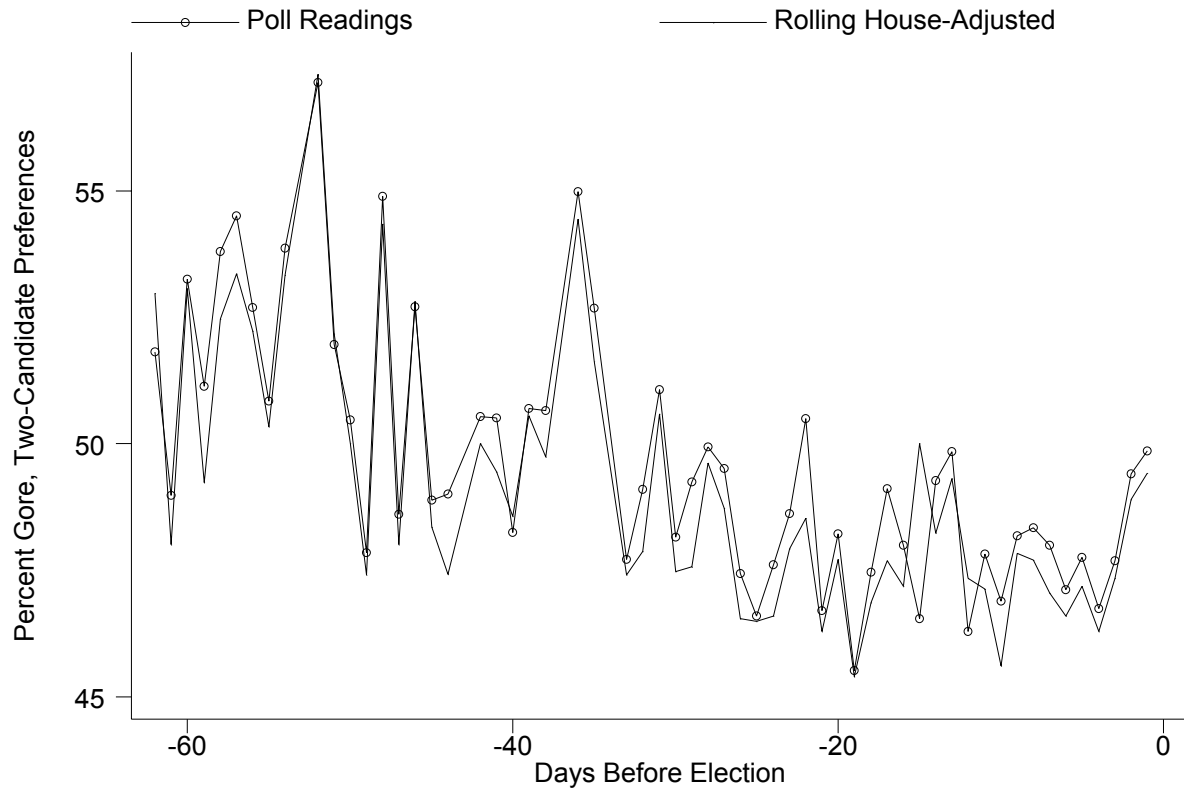


Figure 9: Pooled Rolling House-Adjusted “Forecasts” (with Confidence Intervals), Labor Day to Election Day, 2000

