Dynamic Bayesian Forecasting of Presidential Elections in the States

Drew A. Linzer
dlinzer@emory.edu
Department of Political Science
Emory University

Abstract

I present a dynamic Bayesian forecasting model that enables early and accurate prediction of U.S. presidential election outcomes at the state level. The method systematically combines information from historical forecasting models in real time with results from the large number of state-level opinion surveys that are released publicly during the campaign. The result is a set of forecasts that are initially as good as the historical model, then gradually increase in accuracy as Election Day nears. I employ a hierarchical specification to overcome the limitation that not every state is polled on every day, allowing the model to borrow strength both across states and, through the use of random-walk priors, across time. The model also filters away day-to-day variation in the polls due to sampling error and idiosyncratic "campaign effects," which enables daily tracking of voter preferences towards the presidential candidates at the state and national levels. Simulation techniques are used to estimate the candidates' probability of winning each state and, consequently, a majority of votes in the Electoral College. I apply the model to pre-election polls from the 2008 presidential campaign and demonstrate that the victory of Barack Obama was never realistically in doubt. The model is currently ready to be deployed for forecasting the outcome of the 2012 presidential election.

Acknowledgements: I am grateful to Cliff Carrubba, Tom Clark, and Justin Esarey for thoughtful reactions to earlier versions of this paper. Nigel Lo provided helpful research assistance. A special debt is owed to my colleague Alan Abramowitz, who could not have been more generous with his time, his insight, and most importantly, his data.

1 Introduction

Every four years, American political pundits and analysts devote seemingly endless hours to dissecting the presidential election campaign and trying to forecast the winner. These efforts increasingly rely upon the interpretation of quantitative historical data and the results of pre-election (or *trial-heat*) public opinion polls asking voters their preferred candidate for president. The 2008 presidential campaign in particular witnessed a remarkable increase in the number of pre-election polls conducted at the *state* level, where presidential elections are ultimately decided. By Election Day, over 1,700 distinct state-level surveys had been published by media organizations and private polling firms, in every U.S. state, totaling over one million individual interviews (*Pollster.com*, 2008). In the most competitive "swing" states, new polls were released almost daily as Election Day neared. The widespread availability of these survey findings critically shaped both how the campaign was reported in the news media as well as how the presidential candidates were perceived by voters (Pew Research Center, 2008; Becker, 2008; Traugott and Lavrakas, 2008).

State-level pre-election survey data represent a rich—if extremely noisy—new source of information for both forecasting election outcomes and tracking the evolution of voter preferences during the campaign. Interest in measuring and predicting these outcomes is not limited to those in the media whose job it is to explain campaign trends to the public (Broh, 1980; Stovall and Solomon, 1984; Iyengar, 1991; Rosenstiel, 2005). Campaign strategists who make decisions about the allocation of hundreds of millions of dollars worth of advertising and manpower need to be able to ascertain candidates' relative positioning in the electorate, and their potential to carry various states on the way to winning the presidency (Center for Responsive Politics, 2008; Balz and Johnson, 2009; Jamieson, 2009). In addition, academic researchers have long been interested in the factors that predict presidential election outcomes (e.g., Lewis-Beck and Rice, 1992; Campbell and Garand, 2000), the forecasting value of historical models versus pre-election public opinion polls (Brown and Chappell, 1999; Holbrook and DeSart, 1999), the earliness with which accurate forecasts can be made (Gelman and King, 1993; Erikson and Wlezien, 1996; Wlezien and Erikson, 1996), the dynamics behind public opinion during the campaign (Campbell, Cherry and Wink, 1992; Romer et al., 2006; Erikson and Wlezien, 1999; Wlezien and Erikson, 2002; Panagopoulos, 2009*a*), and the extent to which campaigns affect the eventual result (Shaw, 1999; Hillygus and Jackman, 2003; Arceneaux, 2006; Vavreck, 2009). The aim of this paper is to produce quantities of interest to each of these constituencies: state- *and* national-level estimates of not only the *current* preferences of voters at every point in the campaign; but also *forecasts* of presidential candidates' vote shares and probabilities of victory on Election Day.

I introduce a dynamic Bayesian forecasting model that unifies the regression-based historical forecasting approach developed in political science and economics with the poll-tracking capabilities made feasible by the recent upsurge in state-level opinion polling. Existing historical models are designed to predict presidential candidates' popular vote shares, at a single point in time—usually two to three months in advance of an election—from structural "fundamentals" such as the popularity of the incumbent president, whether the incumbent is running for re-election, levels of economic growth, changes in unemployment rates, and so forth (e.g., Bartels and Zaller, 2001; Abramowitz, 2008; Campbell, 2008*b*; Erikson and Wlezien, 2008; Fair, 2009). In line with theories of retrospective voting, voters tend to punish incumbent party candidates when times are bad, and reward them when economic and social conditions are more favorable (Kinder and Kiewiet, 1991; Nadeau and Lewis-Beck, 2001; Duch and Stevenson, 2008).

Although predictions from structural models can be surprisingly accurate, they are also subject to a large amount of uncertainty. Most historical forecasts are based on data from just 10 to 15 past elections, and many only generate national-level estimates of candidates' vote shares. Unless the fundamentals clearly favor one candidate over the other, it is difficult for structural models to confidently predict the election *winner*. Moreover, in the event that an early forecast is in error, structural models contain no mechanism for updating predictions once new information becomes available closer to Election Day. In 2008, for example, Democrat Barack Obama won the presidency with 53.7% of the major-party vote—a sizeable margin, by historical standards. Yet many published forecasts were unsure of an Obama victory. Two months before the election, Erikson and Wlezien (2008) gave Obama a 72% chance of winning. Lewis-Beck and Tien (2008) judged the race to be a toss-up. Campbell (2008*b*) predicted that Republican John McCain would win, with 83% probability. In closer elections, the problem is amplified: political scientists completely failed to predict the victory of Republican George W. Bush in 2000 (Campbell, 2001).

Pre-election polls provide contextual information that can be used to correct potential errors in historical forecasts, increasing both their accuracy and their precision. Polls conducted just before an election generate estimates that are very close to the eventual result, on average (Traugott, 2001, 2005; Pickup and Johnston, 2008; Panagopoulos, 2009*b*). Earlier in the campaign, polls are less effective for forecasting (e.g., Campbell and Wink, 1990; Gelman and King, 1993; Campbell, 1996), but remain useful for detecting trends in voter preferences. This presents certain practical challenges, however. First, not every state is polled on every day, leading to large gaps in the time series; especially in less-competitive states and earlier in the campaign. Second, measured preferences fluctuate greatly from poll to poll due simply to random sampling variability. Such swings have been prone to misinterpretation as representing "real" changes in attitudes. Some amount of multi-survey aggregation and smoothing is therefore necessary to reveal any underlying trends (Erikson and Wlezien, 1999; Jackman, 2005; Wlezien and Erikson, 2007).

The integrated modeling framework that I describe will enable researchers to refine and update structural state-level election forecasts in real time, using the results of every newly available state-level opinion poll. Older polls that contribute less to the forecast are used to estimate past trends in state-level opinion. To handle the uneven spacing of pre-election polls, the model borrows strength hierarchically across both states and days of the campaign. It also detects and accounts for "campaign effects" due to party conventions or major news events that influence mass opinion in the short term, but may or may not be related to the election outcome (Finkel, 1993; Holbrook, 1994; Shaw, 1999; Wlezien and Erikson, 2001).

The result is a set of election forecasts that are produced early in the campaign, become increasingly accurate as Election Day nears, yet remain relatively stable over time. Because these forecasts depend on reported levels of support for each candidate in the trial-heat polls, my model also yields daily estimates of "current" opinion in each state at any point during the campaign. In sum, I estimate not only where opinion has been, but also where it is going, with associated measures of uncertainty. The model further generates logically valid estimates of the *probabilities* that either candidate will win each U.S. state and the Electoral College vote as a whole, as a function of the available polling data, the prior confidence in the predictions of the historical model, and the proximity to Election Day.

I apply the model to the problem of forecasting the outcome of the 2008 U.S. presidential election, using the benefit of hindsight to evaluate model performance. I first analyze the entire trend of voter preferences in all fifty states through Election Day. I then return to multiple points prior to the election to investigate the predictions of the model, had it been used. Contrary to much of the media speculation at the time, Obama's victory was highly predictable far in advance, and, based on all available information, never realistically in doubt.

2 Research background

Presidential elections in the United States are decided at the state level, through the institution of the Electoral College. Within each state, candidates are awarded electoral votes on a winner-take-all basis, with the number of electoral votes per state equal to a state's total number of federal Representatives and Senators. (There are minor exceptions to this rule in Maine and Nebraska.) The candidate receiving a majority of electoral votes wins the election. In recent elections, election outcomes in a majority of states have not been competitive. In these "safe" states, the winning candidate is largely predetermined, even if the exact percentage of the vote that each candidate will receive remains unknown. The division of the country into Republican "red states" and Democratic "blue states" has been much remarked upon (e.g., Farhi, 2004; Dickerson, 2008). Most observers consider 30 to 35 of the fifty states to be safe, with each side containing a similar number of electoral votes.

Presidential elections are, as a result, effectively won or lost in a smaller number of pivotal "swing" or "battleground" states. Florida and Ohio stand out as the most prominent recent examples. Outcomes in these states are, by their very nature, both more important—and more difficult—to predict in advance. It is especially in the swing states where the potential value of pre-election polling to forecasting and opinion tracking is the greatest.

2.1 Pre-election polling data

Pre-election polls are typically conducted as random samples of registered or "likely" voters who are asked their current preferences among the presidential candidates. The wording of the 2008 Washington Post-ABC News tracking poll, for example, read "If the 2008 presidential election were being held today and the candidates were Barack Obama and Joe Biden, the Democrats, and John McCain and Sarah Palin, the Republicans, for whom would you vote?" Pollsters tabulate the answers to this question and report the percentages of voters providing each response, accompanied by a margin of sampling error that usually corresponds to a 95% level of confidence. Most polls also record the percentage of voters who are undecided; others tally support for non-major party candidates as well.

During the 2008 U.S. presidential campaign, survey researchers and media organizations released the results of 1,731 state-level public opinion polls asking voters their current choice for president (*Pollster.com*, 2008). I will use these data to illustrate and validate the model I describe, but the model could easily be applied to any previous (or future) campaign for which a series of state-level opinion polls are available. The quantity of interest here will be the Obama share of the major party vote, measured in surveys as the proportion of respondents favoring Obama out of the total number supporting either Obama or McCain.

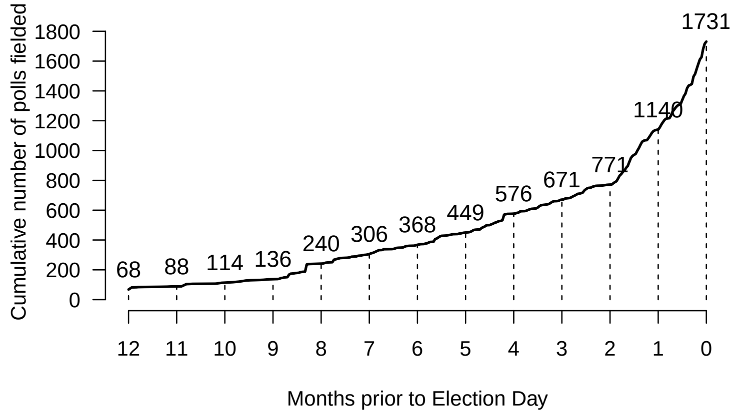

Figure 1: *Cumulative number of state-level presidential pre-election polls fielded in advance of the 2008 election. Source:* Pollster.com *(2008).*

Polls conducted during the primary season, before the nominations of Obama and McCain were assured, asked only about a hypothetical match-up between the two candidates.

More than 150 distinct entities published state-level polls in 2008, but a much smaller number were responsible for the majority of the polling. (The exact total is somewhat approximate because many polls were conducted in association between private polling firms and news companies.) The most active survey firm was Rasmussen Reports, which accounted for 322 polls as an independent unit, and another 47 working with FOX News. The seven largest firms—Rasmussen, SurveyUSA, the Quinnipiac University Poll, Research 2000, Zogby, American Research Group, and PPP—were responsible for 1,106 of the published polls. The median pre-election survey contained 600 respondents. On average, 91% of those polled reported a preference for Obama or McCain; unsurprisingly, the proportion of undecided voters was larger in early polls and decreased closer to Election Day. As most polls spend multiple days in the field to complete the sample, I will consider each poll as having "occurred" on the final day in which interviews were conducted.

Towards the end of the campaign, the rate of pre-election polling accelerated, with more than half of all surveys being fielded in the final two months before the election (Figure 1). There were also more polls fielded in states that were expected to be closely competitive: Florida and Ohio were each surveyed 113 times, Pennsylvania was surveyed 101 times, and

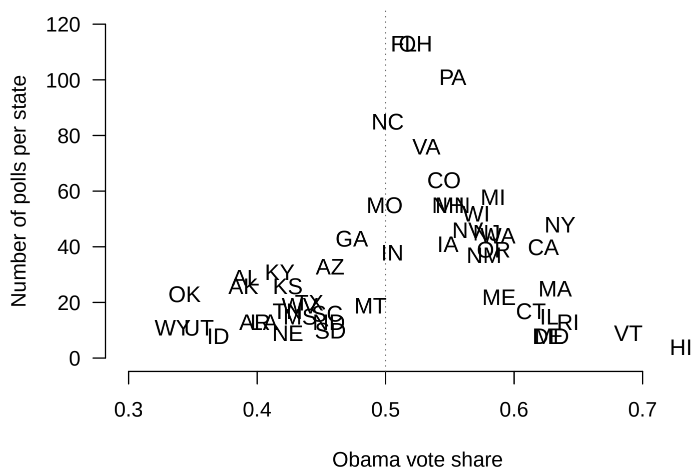

Figure 2: *More pre-election presidential polls were fielded in states that were expected to be competitive in 2008, as indicated by the closeness of Obama's eventual vote share to 50%.*

another 85 polls were conducted in North Carolina (Figure 2). On the low end, fewer than ten polls were conducted in Hawaii, Delaware, Maryland, and Vermont; all safe Democratic states—as well as in Idaho and Nebraska; both safe Republican states. States such as Missouri, Indiana, Georgia, and especially Montana are among the most interesting from a forecasting perspective because the outcomes in those states were very close despite being polled relatively infrequently.

2.2 Characteristics of trial-heat survey results

In any given poll, a large number of factors will contribute to differences between survey estimates of voters' preferences and the actual election outcome. Sampling variability is the largest single source of error, accounting for half or more of the total variation in survey-based estimates during the campaign (Erikson and Wlezien, 1999; Wlezien and Erikson, 2002). But even very close to Election Day, additional discrepancies may arise. Accurate forecasts rely on unbiased samples and truthful respondents. Other errors due to so-called *house effects* reflect systematic differences between survey organizations in screening methodology, survey design, question wording, interviewer quality, and weighting schemes, among other factors (Deming, 1944; McDermott and Frankovic, 2003; Wlezien and Erikson, 2007). Fortunately, these errors oftentimes—although certainly not always—cancel out by averaging across multiple concurrent polls (Traugott and Wlezien, 2009). Without observing the election outcome, it is nearly impossible to estimate the "average" bias due to house effects (Jackman, 2005). Most survey analysts are therefore comfortable with the simplifying assumption that the mean house effect is zero; implying that the results of multiple polls of the same population at around the same time should vary around the true, underlying population proportion.

Farther away from Election Day, polls are less effective for predicting the eventual result. This is because voters' reported preferences change during a campaign in response to temporally (and temporarily) salient campaign events, and as voters learn more about the candidates and pay greater attention to the race as the election nears (Finkel, 1993; Gelman and King, 1993; Stevenson and Vavreck, 2000). For voters who are undecided or who have devoted minimal effort to evaluating the candidates, the mere presence of intense and consistently favorable media coverage of one of the candidates—as occurs during the party conventions, for example—can sway individuals to report preferences that differ from their eventual vote choice (Zaller, 1992). Research into campaign dynamics suggests that many voters wait until the very end of the campaign to ultimately make up their mind (Rosenstone, 1983; Gelman and King, 1993; Arceneaux, 2006).

Taken together, these characteristics of pre-election polls motivate the specification of a statistical model that combines state-level election forecasting with the tracking of voter preferences during a campaign. Far removed from Election Day, forecasts should rely primarily on a historical, "structural" model. By Election Day, those forecasts should transition to being based upon the available polling data. The model should account for—and smooth away—sampling variability, with polls containing larger samples contributing more information to the estimator. In states where polling is infrequent, we would like the model to estimate the trend in opinion between consecutive surveys, so that forecasts can be made for every state on every day of the campaign, regardless of whether a survey was conducted on that day. At the same time, the model should distinguish the projected election outcome from the "current" proportion favoring either candidate. Finally, the specification of the model should separate short-term, national campaign effects (as due to the party conventions) from underlying, long-term dynamics in state-level mass opinion. Short-term effects will appear in the data as common daily trends in the polls across all fifty states. Discerning such trends requires a multivariate time series and would not be possible using only a single sequence of national-level polls, for example.

Existing approaches to integrating polling data and historical forecasting models contain certain of these elements. Campbell (2008*b*) and Erikson and Wlezien (2008) generate static election forecasts from a multiple regression model fitted to past election data, including trial-heat poll results as an independent variable. The primary limitation of this method for forecasting future elections is that, as emphasized by Holbrook and DeSart (1999), the regression weights estimated for the opinion variable are subject to uncertainty (sample sizes are typically small) and may have changed since previous elections. A second strategy uses contemporary survey data to update historical model-based forecasts in a Bayesian manner (e.g., Brown and Chappell, 1999; Strauss, 2007; Lock and Gelman, 2010), yet none of these methods are general enough to utilize data from all available state-level opinion polls in real time. Bayesian techniques for estimating trends in voter preferences using pre-election polls either do not produce forecasts until very late in the campaign (Christensen and Florence, 2008), or require that the election outcome is already known (Jackman, 2005), making forecasting impossible.

3 A dynamic Bayesian forecasting model

To begin, denote as h_i a forecast of the election outcome in states $i = 1 \ldots 50$, based upon a historical model that produces predictions far in advance of Election Day. There are a variety of approaches to generating these baseline forecasts (e.g., Rosenstone, 1983; Holbrook, 1991; Campbell, 1992; Lock and Gelman, 2010). Since no definitive model exists, the precise choice of how to estimate h_i—as well as how much prior certainty to place in those estimates—is left to the analyst. Values chosen for h_i should be theoretically well-motivated, however, as they will be used to specify an informative prior distribution for the estimate of the state-level election outcome. The performance of the forecasting component of the model will depend on h_i most heavily early in the campaign. Closer to the election, the choice of h_i matters less because the forecast gives greater weight to the newly available polling data.

3.1 Specification

As the campaign progresses, increasing numbers of pre-election polls are released. Let $j = 1 \ldots J$ index days of the campaign, so that $j = 1$ corresponds to the first day of polling and $j = J$ is Election Day. The model can be fitted on any day of the campaign, using as many polls are presently available. The J days prior to Election Day need not include the dates of every single pre-election poll if, for example, an investigator wishes to disregard polls conducted far ahead of the election. On day j of the campaign, let K_j denote the total number of state-level polls that have been published until that point. Indexing polls as $k = 1 \ldots K_j$, I denote the sample size of the kth survey as n_k, of whom y_k respondents indicate support for the Democratic candidate. I restrict n_k to include only those respondents who indicate a preference for one of the two major party candidates. In survey k, assuming a random sample,

$$y_k \sim \text{Binomial}(\pi_{i[k]j[k]}, n_k), \tag{1}$$

where π_{ij} represents the true, but unobserved, proportion of voters in state i *who would tell pollsters* that they intend to vote for the Democrat on day j. Notationally, $i[k]$ and $j[k]$ represent the state and day of poll k, respectively. Polls with larger sample sizes n_k contribute more information about π_{ij}.

The proportion π_{ij} is a function of two components: a state-level effect β_{ij} that captures the long-term dynamics of voter preferences particular to state i, and a national-level effect δ_j that detects idiosyncratic departures from β_{ij} due to "campaign effects" common to all fifty states. I place both β_{ij} and δ_j on the logit scale, as π_{ij} is bounded by zero and one:

$$\pi_{ij} = \text{logit}^{-1}(\beta_{ij} + \delta_j). \tag{2}$$

Substantively, the day-to-day trend in δ_j contributes to the results of pre-election polls, but is a source of error when trying to forecast each state's Election Day outcome. To the extent that δ_j diverges from zero, we would like to filter away its effects on π_{ij}. This leaves behind the state-level quantity β_{ij} which forms the basis for the Election Day forecast, $\text{logit}^{-1}(\beta_{iJ})$.

The presence in the model of the national-level effect δ_j is crucial to estimating the temporal dynamics of voter preferences within each state. It is the existence of these common trends in state-level opinion that enables the model to borrow strength across states, to estimate π_{ij} during intervals when polling data are scarce. Through δ_j, the model accounts for daily cross-state covariation in voter opinion. If and when opinion across states does not trend together, this will also be detectable by the model.

Across time, estimates of β_{ij} and δ_j are connected by two distinct Bayesian random walk processes. When forecasting weeks or months ahead of the election, there will be a gap in the polling data between the last published survey on day j, and Election Day J. To bridge this interval, the β_{ij} are assumed to follow a reverse random walk, "beginning" on Election Day. The idea is similar to Strauss (2007). State-level historical forecasts h_i are incorporated into the model through an informative Normal prior distribution over the β_{iJ},

$$\beta_{iJ} \sim N(\text{logit}(h_i), s_i^2). \tag{3}$$

The variances s_i^2 are set in advance by the analyst, to reflect prior uncertainty about the accuracy of the respective h_i. Depending upon the procedure used to generate h_i, it may be possible to estimate s_i^2. Larger values of s_i^2 indicate greater uncertainty in h_i, which upweights the influence of the polling data. Smaller values of s_i^2 place greater confidence in the historical forecast and make model forecasts less sensitive to new polling data. Moving back in time, each previous day's estimate of β_{ij} is given the prior distribution

$$\beta_{ij} \sim N(\beta_{i(j+1)}, \sigma_\beta^2), \tag{4}$$

where the estimated variance parameter σ_β^2 captures the rate of daily change in β_{ij}. σ_β is assigned a uniform prior distribution. As older polls are superseded by newer information, they contribute less to the forecast, but they leave behind the historical trend in β_{ij} from days 1 to j of the campaign.

The national effects δ_j are also modeled as a reverse random walk process,

$$\delta_j \sim N(\delta_{(j+1)}, \sigma_\delta^2), \tag{5}$$

with $\delta_J = 0$ fixed by definition. From Equations 2 and 3, the Election Day forecast is therefore a compromise between the recent poll results and the predictions of the structural model. This assumption also identifies the model by anchoring the scales of β_{ij} and δ_j on day J. The estimated variance parameter σ_δ^2 captures the rate of daily change in δ_j; I assume a uniform prior over σ_δ. Note that estimates of δ_j can only be meaningfully interpreted until the most recent day on which a survey was administered, which may be far in advance of the election. Past estimates of δ_j provide a historical record of the magnitude and direction of national campaign effects—quantities of considerable substantive interest. Combined with earlier estimates of β_{ij}, the model produces a smoothed trend of "current" state-level preferences π_{ij} over the entire length of the campaign. This series is both important to analysts and useful for posterior model checking of proper fit of the model to the data.

In cases where the result of the election is known, as when researching trends in voter preferences from past elections, h_i can be set equal to the outcome in state i. We would then fix $\beta_{iJ} = \text{logit}(h_i)$ instead of specifying the prior distribution in Equation 3, since forecasting (based upon estimating β_{iJ}) is no longer of interest.

3.2 Estimation

Given K_j state-level pre-election polls, and fifty historical forecasts h_i with measures of prior uncertainty s_i^2, the Bayesian model may be estimated using a MCMC sampling procedure. I implement the estimator in the WinBUGS and R software packages (Lunn et al., 2000; Sturtz, Ligges and Gelman, 2005; R Development Core Team, 2011). This produces a rich (and large) set of parameter estimates: the average preferences of voters, π_{ij}, as they would be reported to pollsters in each state at each day in the campaign prior to day j, the trend in national-level campaign effects δ_j, the filtered state-level vote preferences β_{ij}, and the state-level Election Day forecasts, $\text{logit}^{-1}(\beta_{iJ})$. Measures of uncertainty for each of these estimated quantities are based on the spread of the simulated posterior draws of each parameter.

Following estimation, the posterior probability that the Democratic candidate will win the election in state i, $\Pr(\text{logit}^{-1}(\beta_{iJ}) > 0.5)$, is calculated as the proportion of posterior draws of β_{iJ} that are greater than 0. To calculate the probability that the Democratic candidate wins the presidency, I select the fifty posterior draws of β_{iJ} produced in a single iteration of the sampling algorithm. I then tally the total number of electoral votes in states where the Democratic candidate is forecast to receive more than 50% of the two-party vote, and add the three electoral votes of the District of Columbia, which is reliably Democratic. Repeating this calculation across multiple sampling iterations produces a distribution of predicted electoral vote outcomes. The proportion of these simulations in which the Democratic candidate receives an absolute majority—270 or more—of the 538 electoral votes is taken as the Democratic candidate's probability of victory.

One limitation is that for forecasts being made far in advance of Election Day, the model becomes slow to converge due to the lack of available polling data. A slight modification to the specification of the β_{ij} parameters makes the problem tractable and accelerates MCMC convergence. Rather than let β_{ij} vary by day, I divide the J days of the campaign into J/W short spans or "windows" of W days apiece. In Equation 2, I replace β_{ij} with $\beta_{it[j]}$, which denotes the value of β in state i for the time period $t = 1 \dots J/W$ containing day j. Parameters

δ_j are still estimated for each of the J days of the campaign. The election forecast becomes $\text{logit}^{-1}(\beta_{it[J]})$, and this parameter receives the prior distribution in Equation 3. In practice, values of W equal to just three to five days can significantly improve the estimation process, without substantively altering the election forecast. This simplification works because while δ_j fluctuates quite a bit on a day-to-day basis, β_{ij} changes far more gradually over time.

4 Application: The 2008 U.S. presidential election

The 2008 U.S. presidential election was widely predicted to result in a victory for the Democratic candidate, Barack Obama (Campbell, 2008*a*). Republican John McCain suffered from two major drags on his candidacy: an extremely low approval rating for the incumbent Republican president, George W. Bush; and a weak economy, whether measured in terms of GDP growth, consumer satisfaction, unemployment rates, or other factors. Yet as a candidate, Obama consistently lagged behind expectations in national pre-election polls—even falling behind McCain for a brief period after the Republican National Convention in early September (*Pollster.com*, 2008). News reports quoted worried Democrats suddenly wondering if Obama would lose after all (Kuhn and Nichols, 2008). Contributing to the uncertainty were the lingering effects of the unusually long Democratic primary battle between Obama and then-Senator Hillary Clinton; as well as questions about what effect Obama's race might have on the willingness of white voters to support him in the November election.

I investigate a series of issues surrounding the election. By how much did voter preferences change over the course of the campaign? What were the short-term effects of particular campaign events on reported voter preferences? Which were the actual "swing" states and how soon was this knowable? Finally, how early, and with what precision, was the election outcome predictable from a combination of structural factors and pre-election polls?

4.1 The historical forecast

Forecasting presidential elections using a structural model imposes a tradeoff between earliness and accuracy. I produce both an inaccurate "early" forecast and an accurate "late" forecast based on the Abramowitz (2008) "Time-for-Change" model. The late forecast, available about two months prior to Election Day, predicts the national-level vote share for the incumbent party candidate from three variables: the annualized growth rate of second quarter GDP, the June approval rating of the incumbent president, and a dummy for whether the incumbent party has been in office for two or more terms. For the earlier forecast, available up to six months in advance, I use a variation of the model fitted to changes in *first* quarter GDP, and the March presidential approval rating. Based on the results of 15 previous presidential elections, the early forecast would have predicted Obama to receive 56.8% of the national major-party vote in 2008. The more proximate late forecast predicted that Obama would receive 54.3%. Both of these forecasts overestimated Obama's actual vote share of 53.7%.

To translate national forecasts to the state level, I exploit a twenty-year pattern in presidential election outcomes. Since 1980, state-level Democratic vote shares have tended to rise and fall in national waves, by similar amounts across states. In 2004, Democrat John Kerry received 48.8% of the two-party presidential vote. Assuming the same trend would continue in 2008 (as it did), the Time-for-Change model predicts an average state-level gain by Obama of 8% in the early forecast and 5.5% in the late forecast.

I calculate h_i by first adding to Kerry's 2004 state-level vote shares either 5.5% or 8% depending upon the timing of the forecast. I then adjust the forecast by a further 6% in candidates' home states, adding in for Hawaii (Obama) and Texas (Bush in 2004), and subtracting away for Arizona (McCain) and Massachusetts (Kerry in 2004). This correction was estimated from past elections by Holbrook (1991) and Campbell (1992), and also employed by Lock and Gelman (2010). For the late forecast, I set $s_i^2 = 0.01$ on the logit scale, which places 95% of the prior probability in an approximately $\pm 5\%$ range around each historical forecast. To reflect greater uncertainty in the early forecast, I set $s_i^2 = 0.04$, which places

95% of the prior probability in an approximately ±10% range around each h_i. Both values are consistent with the amount of state-level variability estimated by Lock and Gelman (2010).

To test the sensitivity of the broader model to the choice of h_i and s_i^2, I also examine an alternative, more general approach to forecasting based on the idea of a Democratic "normal vote." For this, I set h_i equal to the mean Democratic vote share in each state over the previous four presidential elections: 1992 and 1996, which were won by a Democrat, and 2000 and 2004, which were won by a Republican. In 2008, this would have systematically under-predicted Obama's performance by 1.8%, on average. I again let $s_i^2 = 0.04$.

4.2 Election Day estimates

I fit the model to the 1,363 state-level pre-election polls fielded up to six months prior to Election Day, basing h_i on the final Time-for-Change forecast. The model is estimated using a three-day window ($W = 3$) for parameters $\beta_{it[j]}$. I run three sequences of 50,000 MCMC iterations, discarding the first half as burn-in and thinning to keep every 75th draw, saving 1,000 posterior samples. Convergence of the MCMC algorithm is assessed by values of $\hat{R} \approx 1$ and visual confirmation that the three sequences are completely mixed (Gelman et al., 2004).

To evaluate the fit of the model, I compare the observed pre-election poll results to the estimates of π_{ij} in the eight states with the closest election outcomes (Figure 3). In states where large numbers of polls were fielded (e.g., Florida), π_{ij} appears to accurately capture the underlying trend in voters' reported preferences during the campaign. In states with fewer polls, the trend in π_{ij} still matches the survey data, but with a more noticeable effect of the common δ_j term during gaps in the polling. Final estimates of π_{iJ} are very close to the actual election outcome, which validates the accuracy of late pre-election polls.

In a typical state, the variance in the poll results was three to five times greater than the variance in the estimated π_{ij}. This suggests that the combined error in pre-election polls due to sampling variability and house effects may be even greater than what was estimated

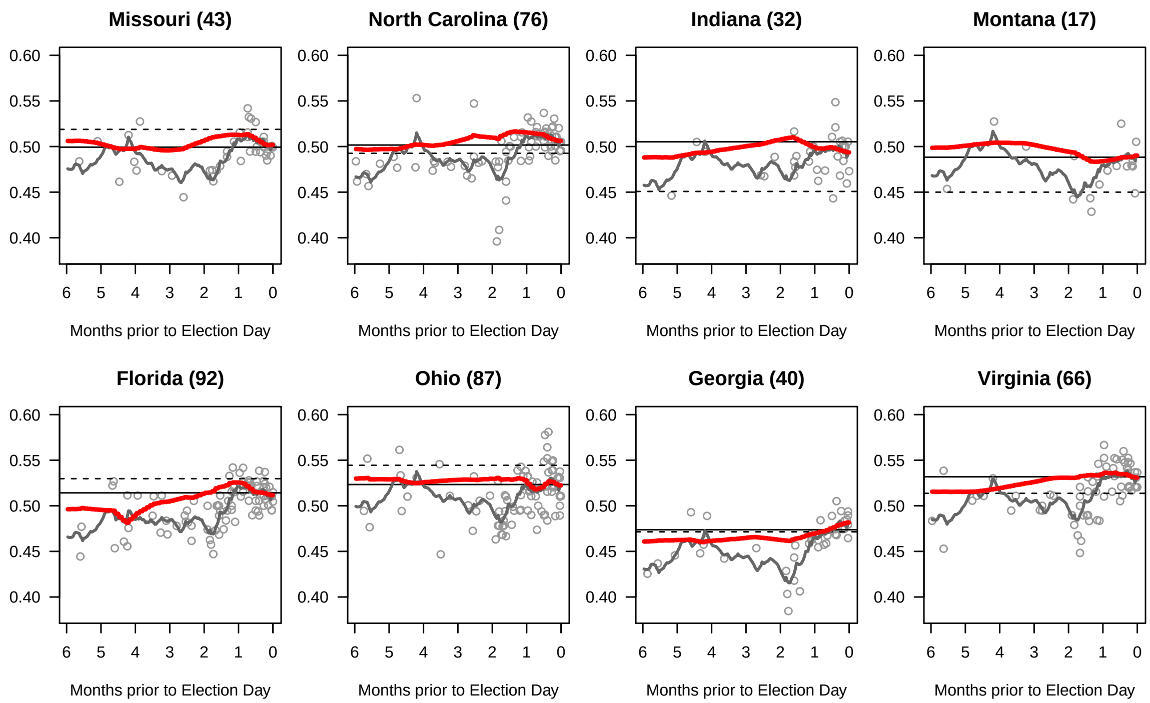

Figure 3: *Tracking the 2008 presidential polls in the eight closest states. The vertical axis is the percent supporting Obama. Points denote observed poll results; horizontal lines indicate the actual election outcome (solid) and the structural model prediction (dashed). The jagged gray line is the state-level daily estimate of voter preference for Obama,* π_{ij}. *The red line is the filtered within-state trend of logit*$^{-1}(\beta_{it[j]})$. *The number of polls fielded in each state within six months of the election appears in parentheses.*

by Erikson and Wlezien (1999) and Wlezien and Erikson (2002). Across all fifty states, estimates of π_{ij} varied by 4% to 10% over the final six months of the campaign, but 98% of the day-to-day changes in π_{ij} were by less than 0.5%. The opinions that voters held about the candidates changed very gradually over time, compared to the large swings in the polls.

Results of the model further suggest that most of the temporal variation in state-level opinion was due to national-level campaign effects. Once the national effect of δ_j is filtered away, the state effect $\beta_{it[j]}$ is highly stable over time (Figure 4). Where trends in $\beta_{it[j]}$ do occur, they tend to happen in one sustained direction. The consistently negative values of δ_j reflect the fact that Obama ran "behind" his eventual election performance in most states for most of the campaign. Interestingly, some of the most significant short-term movements in δ_j coincide with major campaign events. Hillary Clinton's concession of the Democratic

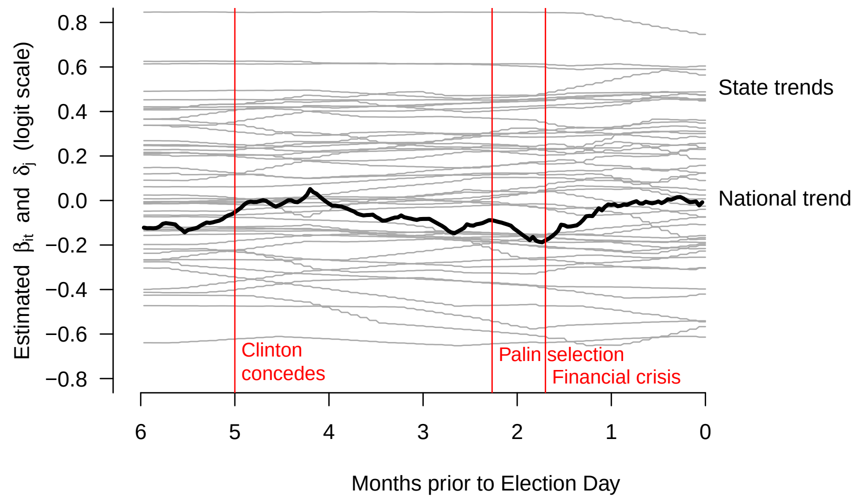

Figure 4: *State-level trends in $\beta_{it[j]}$ (gray) and the national trend in δ_j (black) during the final six months of the 2008 U.S. presidential election.*

nomination and endorsement of Barack Obama in early June occurs in the middle of a runup in support for Obama. The selection by John McCain of Sarah Palin as his vice presidential nominee on August 29 begins a downward slide for Obama. And the start of the financial crisis in mid-September marks the beginning of Obama's final push towards victory. This effect in particular was already evident in the trends in π_{ij} in Figure 3.

On Election Day, the model produces highly accurate forecasts of the state-level election outcomes (Figure 5). As imperfect predictions may arise from flaws in the model, the underlying survey data, or both, this result is reassuring. Compared to the actual result, predictions from the structural model alone—that is, without using any information from pre-election polls—had a mean absolute deviation (MAD) of 2.6%, and incorrectly predicted Arkansas, Indiana, Missouri, and North Carolina. Once the pre-election polls are incorporated, the MAD drops to 1.4%. The complete model correctly predicts Arkansas and North Carolina, and only misses outcomes in Indiana by 1.2% and Missouri by 0.3%. These two states were essentially toss-ups, however. In Missouri, the model estimated a 66% chance of victory for Obama; he lost with 49.9% of the vote. In Indiana, the model estimated a 20% chance of victory for Obama; he won with 50.5% of the two-party vote. The next

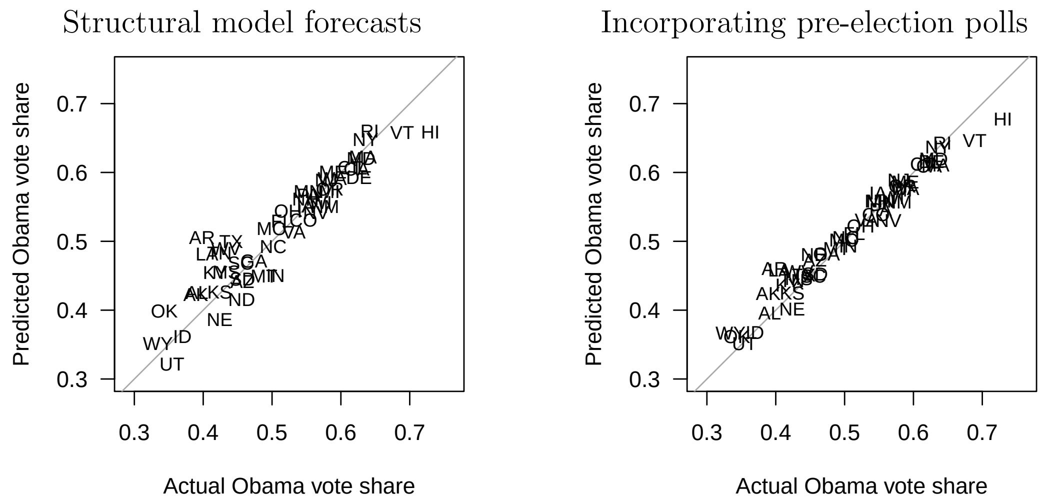

Figure 5: *Comparison of state-level forecasts to presidential election results. Left: forecasts based only on the structural Time-for-Change model with home-state effects. Right: forecasts incorporating pre-election polling through Election Day.*

most uncertain states were North Carolina (91% chance that Obama would win; he did) and Montana (9% chance that Obama would win; he lost).

With so little uncertainty in the state-level predictions, the posterior distribution of Obama's electoral vote tally indicates that he would win the presidency with 100% probability (Figure 6). In only one simulation out of a thousand did Obama even receive as few as 311 electoral votes. The modal number of electoral votes predicted for Obama was 364. In actuality he won 365, which (in a historical anomaly) included the single electoral vote of Nebraska's 2nd Congressional District. By comparison, the range of predicted electoral votes generated by Lock and Gelman (2010), who updated a prior structural forecast made close to Election Day using only *one set* of state-level pre-election polls conducted nine months before the election, was between 250 and 450 electoral votes. Their simulations gave Obama a 99.9% chance of victory, but this high level of certainty was only achievable because Obama won by a relatively large margin in 2008. My results indicate that by updating continually during the election, outcomes can be forecasted with much greater precision. Naturally, this will matter more when the election is close.

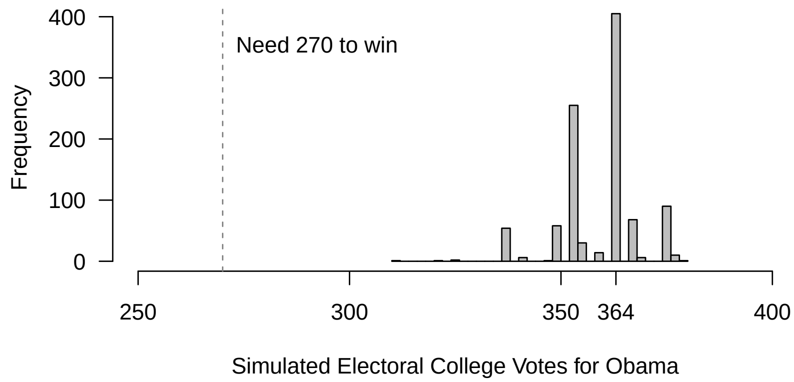

Figure 6: *Posterior distribution of simulated Electoral College votes for Obama, on Election Day. The uneven spacing occurs because only a small number of states had nontrivial posterior probabilities that Obama would either win (100%) or lose (0%).*

Using a less-accurate structural forecast makes almost no difference to the model estimates once a complete set of pre-election polls has been observed through Election Day. Estimates of δ_j and $\beta_{it[j]}$ exhibit the same historical trends as in Figure 3 whether the structural forecast is based on the Time-for-Change model (shown) or the "normal vote" model with weaker prior beliefs. The "normal vote" forecast mispredicts nine states and has a MAD of 4.2%. But with the polls incorporated, the MAD becomes 1.6% and only Indiana is forecasted incorrectly. The probability of an Obama victory remains 100%.

4.3 Forecasting during the campaign

How accurately can the outcome of a presidential election be forecasted at earlier points in the campaign? Suppose that eight weeks remain before Election Day. To illustrate how the model forecasts ahead in time, I fit the model to all state-level opinion polls fielded between six months prior to the election and the current day. On day $j = J - 56$, the proportion supporting supporting Obama in each state, π_{ij}, is estimated as a function of national effects δ_j and state effects $\beta_{it[j]}$. The quantity of interest, however, is not the current $\text{logit}^{-1}(\beta_{it[j]})$, but rather the Election Day forecast $\text{logit}^{-1}(\beta_{it[J]})$. The model therefore estimates a forward trend in $\beta_{it[j]}$ that shrinks via the reverse random-walk process towards the prior distribution of $\beta_{it[J]}$ (Figure 7). The pace of daily change in $\beta_{it[j]}$ is controlled by σ^2_β. In this manner, the

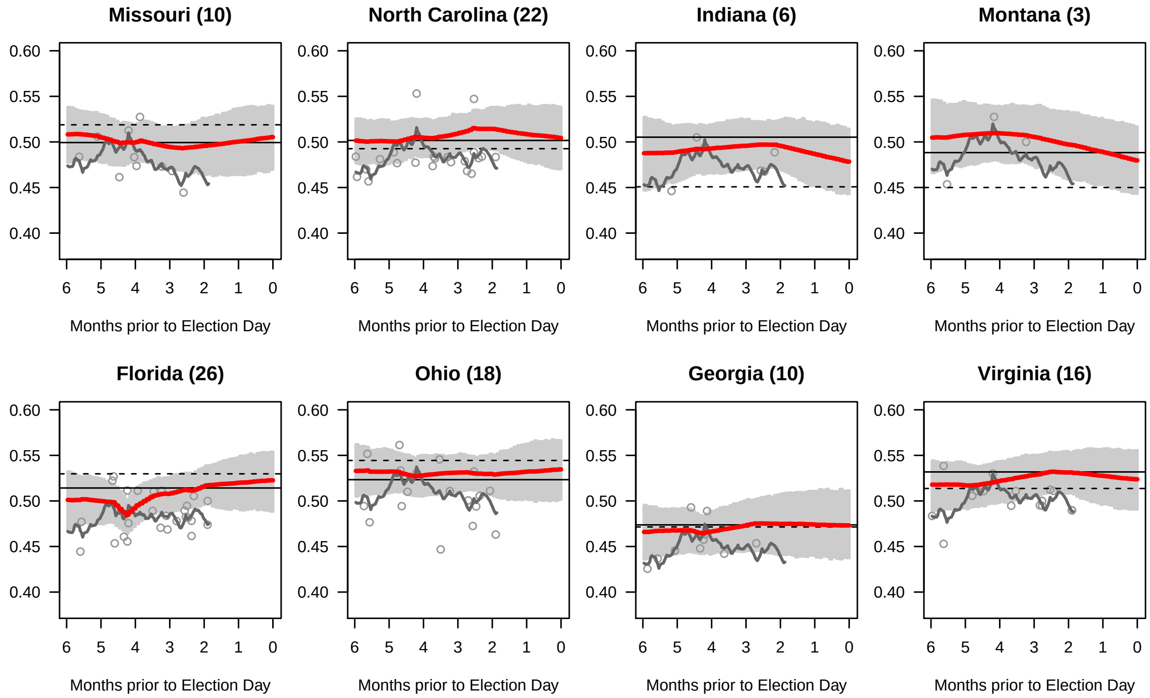

Figure 7: *The 2008 presidential election forecast, with eight weeks remaining in the campaign. As in Figure 3, the vertical axis is the percent supporting Obama, points denote poll results, and horizontal lines indicate the actual election outcome (solid) and the structural model prediction (dashed). The red line is* $\text{logit}^{-1}(\beta_{it[j]})$*, which trends forward to the Election Day forecast,* $\text{logit}^{-1}(\beta_{it[J]})$*; shaded areas represent 95% posterior credible intervals. Gray lines show current opinion,* π_{ij}*, through the final day of polling. Obama's standing in the polls at the eight-week point was well below both his forecasted* and *actual vote shares.*

model forecasts are jointly determined by the prior structural forecast and the most recent filtered estimate of within-state opinion. Nearer to the election, the $\beta_{it[j]}$ on the most recent day of polling have less time to "revert" back to the structural forecasts, which effectively upweights the information in the polls closest to Election Day, and increases the precision of the Election Day forecasts. The probability that Obama will win each state as of day j is calculated from the posterior distribution of $\text{logit}^{-1}(\beta_{it[J]})$.

I now simulate the entire process of forecasting the 2008 election in real time. Beginning four months prior to Election Day, I advance through the campaign in weekly increments, updating the state-level forecasts with all newly available pre-election polling data. I initially base structural forecasts h_i on the "early" Time-for-Change model, but with eight weeks remaining in the campaign, I switch to the "late" Time-for-Change forecast. As a robustness

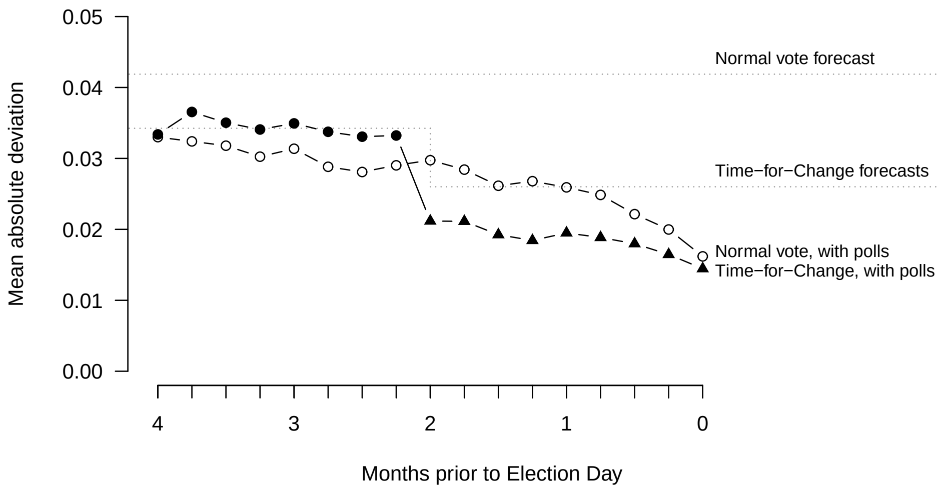

Figure 8: *The mean absolute deviation between state-level election forecasts and the actual outcomes decreases as more pre-election polls become available. Dotted lines denote the MAD of structural forecasts without incorporating polling data. Points indicate the MAD from combining structural forecasts with all available pre-election polls. With two months remaining, estimates based on the Time-for-Change model switch from the "early" (●) to the "late" (▲) forecast.*

check, I also re-run the simulation using the "normal vote" structural forecast. This forecast underestimates Obama's vote share on average, but more significantly, it is less correlated with the actual election outcome, which gives it a larger MAD than either Time-for-Change-based estimate. In each case, I only use polls published up to six months before the election.

The results indicate that by incorporating information from successive weeks of pre-election polls, candidates' state-level vote shares can be predicted with increasingly greater accuracy (Figure 8). The largest improvements to the forecasts occur in the final month of the campaign, when the polls become most informative about the election outcome (e.g., Campbell, 1996). Yet even polls conducted *four to six months* before the election reduce the MAD of the "normal vote" forecast by 1%. In contrast, the early polls have little consistent effect on the MAD of the initial Time-for-Change-based forecast. But with two months remaining in the campaign, combining the polls with the *late* Time-for-Change forecast results in a substantial improvement in accuracy. The naive "normal vote" forecast, once combined with the polls, produces estimates that are as accurate as the late Time-for-Change

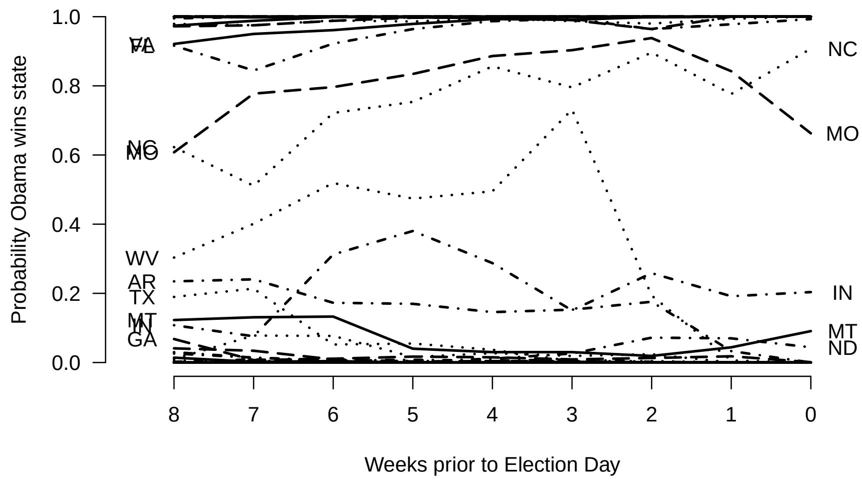

Figure 9: *Weekly estimates of the posterior probability that Obama wins each state, starting eight weeks before Election Day. Forecasts are based on the late Time-for-Change model.*

forecast by the time six weeks remain. In the final two weeks of the campaign, both sets of forecasts converge to a MAD of approximately 1.5%, as noted above.

The forecast MAD measures the overall accuracy of the model across all fifty states, but it is the smaller set of "swing" states in which accurate forecasting matters most. In Figure 9, I plot the posterior probability that Obama wins each state, as estimated over the final two months of the campaign, by combining the polls with the late Time-for-Change forecast. The key feature of this plot is the stability of the model predictions. Large states that proved pivotal to Obama's victory—including Florida, Pennsylvania, Ohio, and Virginia—were already nearly certain to be won by Obama with two months remaining. The battleground states of Missouri, North Carolina, and Indiana were likewise identifiable far in advance of the election. Only West Virginia demonstrates any significant fluctuation, peaking at a 73% chance of Obama victory three weeks prior to Election Day, then decreasing to zero.

It follows that Obama's expected share of the electoral vote was consistently above the 270 needed to win the presidency. Under the Time-for-Change forecast, Obama was never simulated to receive fewer than 270 electoral votes, even as far as four months before Election Day (Figure 10). Moreover, the precision of the Election Day forecasts in Figure 6 was achievable as soon as the late Time-for-Change forecast became available at the two-month

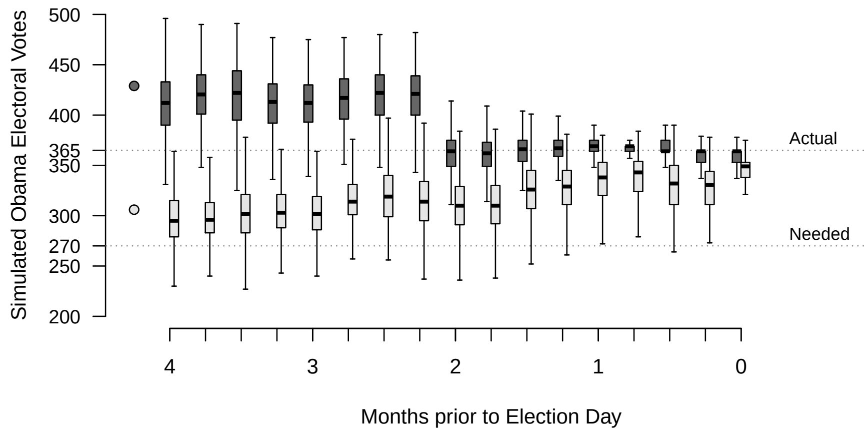

Figure 10: *Boxplots comparing simulated, forecasted Obama electoral votes from the Time-for-Change (dark) and normal vote (light) models, over the final four months of the 2008 presidential campaign. Points at left denote the predictions of the forecasting models without incorporating polls. A majority is 270 electoral votes.*

point. The normal vote-based forecasts, despite underestimating Obama's state-level vote share, also led to highly certain predictions of an Obama victory. Four months in advance of the election, the model predicted an 85% chance that Obama would win. This increased to 95% with two months remaining, and 100% during the final four weeks. Combining a well-motivated structural forecast with information from large numbers of pre-election polls thus generates early and accurate predictions of the presidential election outcome.

5 Discussion

The trend towards increased pre-election polling—especially at the state level—appears likely to continue in the 2012 presidential campaign, and beyond. Public opinion polls have become integral to political reporting, and interest in following the "state of the race" only seems to grow each year. For analysts, the availability of these survey data creates new opportunities for statistical models that can apply theories of mass opinion formation and voter behavior to produce better estimates of voter preferences during the campaign, as well as forecasts of the outcome on Election Day.

This paper has presented a dynamic Bayesian statistical procedure for processing and interpreting the results of state-level pre-election opinion polls in real time. Applied to the 2008 presidential election, the model generated a nearly perfect prediction of which states would be won by Barack Obama and John McCain, and by how much; and estimated with certainty that Obama would win the presidency. The model also produced daily estimates of state-level opinion during the campaign and reliably predicted which states would be most competitive on Election Day. It is ready to be deployed for the 2012 election as soon as the first structural forecasts become available.

The results of my analysis highlight a number of important lessons about presidential campaigns and elections. First, presidential election forecasts can, and should, be made at the state level. State-level outcomes can be predicted accurately and reliably by combining readily available historical and public opinion data. Furthermore, these forecasts need not be overly sensitive to short-term fluctuations in voter preferences during the campaign. Most of the variation in pre-election polls is due to sampling variability. But even after averaging this away, much of the remaining day-to-day variation in state-level opinion is attributable to idiosyncratic, national-level campaign effects. By smoothing and filtering the pre-election polling data, it is possible to produce election forecasts that converge towards the outcome in a gradual and stable manner. During the campaign, any report suggesting that voter preferences have changed by more than a few tenths of a percent on a daily basis should be treated with extreme suspicion.

There nevertheless remain inherent limitations to what can be learned from state-level public opinion data—no matter how many surveys are released in the next election cycle. With current numbers of polls, it is relatively easy to forecast the outcomes of state-level presidential elections on the eve of the election, as I have shown. The challenge remains to produce accurate forecasts many months in advance. My solution seeks to combine the best features of structural forecasts and pre-election polls, downweighting the historical forecasts over time in favor of the information contained in more recent survey data. But even so, the biggest forecasting improvements only occur one or two months in advance of the election. This is not because there is not *enough* polling data, but because the polling data themselves are noisy and, far before Election Day, subject to inaccuracies. Future research into presidential campaign dynamics may yet discover new ways to extract meaning from those early polls.

References

Abramowitz, Alan I. 2008. "Forecasting the 2008 Presidential Election with the Time-for-Change Model." *PS: Political Science & Politics* 41(4):691–695.

Arceneaux, Kevin. 2006. "Do Campaigns Help Voters Learn? A Cross-National Analysis." *British Journal of Political Science* 36(1):159–173.

Balz, Daniel J. and Haynes Johnson. 2009. *The Battle for America 2008: The Story of an Extraordinary Election*. New York: Viking.

Bartels, Larry M. and John Zaller. 2001. "Presidential Vote Models: A Recount." *Political Science & Politics* 34(1):9–20.

Becker, Bernie. 2008. "Competing Web Sites Track Election Polls." *The New York Times* . October 28.

Broh, C. Anthony. 1980. "Horse-Race Journalism: Reporting the Polls in the 1976 Presidential Election." *Public Opinion Quarterly* 44(4):514–529.

Brown, Lloyd B. and Henry W. Jr. Chappell. 1999. "Forecasting presidential elections using history and polls." *International Journal of Forecasting* 15(2):127–135.

Campbell, James E. 1992. "Forecasting the Presidential Vote in the States." *American Journal of Political Science* 36(2):386–407.

Campbell, James E. 1996. "Polls and Votes: The Trial-Heat Presidential Election Forecasting Model, Certainty, and Political Campaigns." *American Politics Research* 24(4):408–433.

Campbell, James E. 2001. "The Referendum that Didn't Happen: The Forecasts of the 2000 Presidential Election." *Political Science & Politics* 34(1):33–38.

Campbell, James E. 2008*a*. "Editor's Introduction: Forecasting the 2008 National Elections." *PS: Political Science & Politics* 41(4):679–682.

Campbell, James E. 2008*b*. "The Trial-Heat Forecast of the 2008 Presidential Vote: Performance and Value Considerations in an Open-Seat Election." *PS: Political Science & Politics* 41(4):697–701.

Campbell, James E. and James C. Garand, eds. 2000. *Before the Vote: Forecasting American National Elections*. Thousand Oaks, CA: Sage Publications.

Campbell, James E. and Kenneth A. Wink. 1990. "Trial-Heat Forecasts of the Presidential Vote." *American Politics Research* 18(3):251–269.

Campbell, James E., Lynna L. Cherry and Kenneth A. Wink. 1992. "The Convention Bump." *American Politics Research* 20(3):287–307.

Center for Responsive Politics. 2008. *U.S. Election Will Cost $5.3 Billion, Center for Responsive Politics Predicts*. http://www.opensecrets.org/news/2008/10/us-election-will-cost-53-billi.html.

Christensen, William F. and Lindsay W. Florence. 2008. "Predicting Presidential and Other Multistage Election Outcomes Using State-Level Pre-Election Polls." *The American Statistician* 62(1):1–10.

Deming, W. Edwards. 1944. "On Errors in Surveys." *American Sociological Review* 9(4):359–369.

Dickerson, John. 2008. "So You Think You're a Swing Voter? Think again: It depends on whether you live in a swing state." *Slate* . September 30. http://www.slate.com/id/2201071.

Duch, Raymond M. and Randolph T. Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge: Cambridge University Press.

Erikson, Robert S. and Christopher Wlezien. 1996. "Of Time and Presidential Election Forecasts." *PS: Political Science and Politics* 29(1):37–39.

Erikson, Robert S. and Christopher Wlezien. 1999. "Presidential Polls as a Time Series: The Case of 1996." *Public Opinion Quarterly* 63(2):163–177.

Erikson, Robert S. and Christopher Wlezien. 2008. "Leading Economic Indicators, the Polls, and the Presidential Vote." *PS: Political Science & Politics* 41(4):703–707.

Fair, Ray C. 2009. "Presidential and Congressional Vote-Share Equations." *American Journal of Political Science* 53(1):55–72.

Farhi, Paul. 2004. "Elephants Are Red, Donkeys Are Blue; Color Is Sweet, So Their States We Hue." *The Washington Post* p. C01. November 2.

Finkel, Steven E. 1993. "Reexamining the 'Minimal Effects' Model in Recent Presidential Campaigns." *The Journal of Politics* 55(1):1–21.

Gelman, Andrew and Gary King. 1993. "Why Are American Presidential Election Campaign Polls So Variable When Votes are So Predictable?" *British Journal of Political Science* 23(1):409–451.

Gelman, Andrew, John B. Carlin, Hal S. Stern and Donald B. Rubin. 2004. *Bayesian Data Analysis; Second Edition*. Chapman & Hall/CRC.

Hillygus, D. Sunshine and Simon Jackman. 2003. "Voter Decision Making in Election 2000: Campaign Effects, Partisan Activation, and the Clinton Legacy." *American Journal of Political Science* 47(4):583–596.

Holbrook, Thomas M. 1991. "Presidential Elections in Space and Time." *American Journal of Political Science* 35(1):91–109.

Holbrook, Thomas M. 1994. "Campaigns, National Conditions, and U.S. Presidential Elections." *American Journal of Political Science* 38(4):973–998.

Holbrook, Thomas M. and Jay A. DeSart. 1999. "Using state polls to forecast presidential election outcomes in the American states." *International Journal of Forecasting* 15:137–142.

Iyengar, Shanto. 1991. *Is anyone responsible? How television frames political issues*. Chicago: University of Chicago Press.

Jackman, Simon. 2005. "Pooling the Polls Over an Election Campaign." *Australian Journal of Political Science* 40(4):499–517.

Jamieson, Kathleen Hall, ed. 2009. *Electing the President, 2008: The Insiders' View*. Philadelphia: University of Pennsylvania Press.

Kinder, Donald R. and D. Roderick Kiewiet. 1991. "Sociotropic Politics: The American Case." *British Journal of Political Science* 11(2):129–161.

Kuhn, David Paul and Bill Nichols. 2008. "Autumn Angst: Dems fret about Obama." *Politico* . September 10. http://www.politico.com/news/stories/0908/13357.html.

Lewis-Beck, Michael S. and Charles Tien. 2008. "The Job of President and the Jobs Model Forecast: Obama for '08?" *PS: Political Science & Politics* 41(4):687–690.

Lewis-Beck, Michael S. and Tom W. Rice. 1992. *Forecasting Elections*. Washington, DC: CQ Press.

Lock, Kari and Andrew Gelman. 2010. "Bayesian Combination of State Polls and Election Forecasts." *Political Analysis* 18(3):337–348.

Lunn, D.J., A. Thomas, N. Best and D. Spiegelhalter. 2000. "WinBUGS–a Bayesian modelling framework: concepts, structure, and extensibility." *Statistics and Computing* 10:325–337.

McDermott, Monika L. and Kathleen A. Frankovic. 2003. "Review: Horserace Polling and Survey Method Effects: An Analysis of the 2000 Campaign." *Public Opinion Quarterly* 67(2):244–264.

Nadeau, Richard and Michael S. Lewis-Beck. 2001. "National Economic Voting in U.S. Presidential Elections." *The Journal of Politics* 63(1):159–181.

Panagopoulos, Costas. 2009*a*. "Campaign Dynamics in Battleground and Nonbattleground States." *Public Opinion Quarterly* 73(1):119–129.

Panagopoulos, Costas. 2009*b*. "Polls and Elections: Preelection Poll Accuracy in the 2008 General Elections." *Presidential Studies Quarterly* 39(4):896–907.

Pew Research Center. 2008. "Winning the Media Campaign: How the Press Reported the 2008 Presidential General Election." *Project for Excellence in Journalism* . October 22. http://www.journalism.org/node/13307.

Pickup, Mark and Richard Johnston. 2008. "Campaign trial heats as election forecasts: Measurement error and bias in 2004 presidential campaign polls." *International Journal of Forecasting* 24(2):272–284.

Pollster.com. 2008. *The Polls: The 2008 Presidential Election*. http://www.pollster.com/polls/2008president.

R Development Core Team. 2011. *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0. http://www.R-project.org.

Romer, Daniel, Kate Kenski, Kenneth Winneg, Christopher Adasiewicz and Kathleen Hall Jamieson. 2006. *Capturing Campaign Dynamics, 2000 and 2004: The National Annenberg Election Survey*. Philadelphia: University of Pennsylvania Press.

Rosenstiel, Tom. 2005. "Political Polling and the New Media Culture: A Case of More Being Less." *Public Opinion Quarterly* 69(5):698–715.

Rosenstone, Steven J. 1983. *Forecasting presidential elections.* New Haven, CT: Yale University Press.

Shaw, Daron R. 1999. "A Study of Presidential Campaign Event Effects from 1952 to 1992." *Journal of Politics* 61(2):387–422.

Stevenson, Randolph T. and Lynn Vavreck. 2000. "Does Campaign Length Matter? Testing for Cross-National Effects." *British Journal of Political Science* 30(2):217–235.

Stovall, James Glen and Jacqueline H. Solomon. 1984. "The Polls as a News Event in the 1980 Presidential Campaign." *Public Opinion Quarterly* 48(3):615–623.

Strauss, Aaron. 2007. *Florida or Ohio? Forecasting Presidential State Outcomes Using Reverse Random Walks.* Working Paper, Princeton University.

Sturtz, Sibylle, Uwe Ligges and Andrew Gelman. 2005. "R2WinBUGS: A Package for Running WinBUGS from R." *Journal of Statistical Software* 12(3):1–16.

Traugott, Michael W. 2001. "Trends: Assessing Poll Performance in the 2000 Campaign." *Public Opinion Quarterly* 65(3):389–419.

Traugott, Michael W. 2005. "The Accuracy of the National Preelection Polls in the 2004 Presidential Election." *Public Opinion Quarterly* 69(5):642–654.

Traugott, Michael W. and Christopher Wlezien. 2009. "The Dynamics of Poll Performance During the 2008 Presidential Nomination Contest." *Public Opinion Quarterly* 73(5):866–894.

Traugott, Michael W. and Paul J. Lavrakas. 2008. *The voter's guide to election polls.* Lanham, MD: Rowman & Littlefield Publishers, Inc.

Vavreck, Lynn. 2009. *The Message Matters: The Economy and Presidential Campaigns.* Princeton, NJ: Princeton University Press.

Wlezien, Christopher and Robert S. Erikson. 1996. "Temporal Horizons and Presidential Election Forecasts." *American Politics Research* 24(4):492–505.

Wlezien, Christopher and Robert S. Erikson. 2001. "Campaign Effects in Theory and Practice." *American Politics Research* 29(5):419–436.

Wlezien, Christopher and Robert S. Erikson. 2002. "The Timeline of Presidential Election Campaigns." *The Journal of Politics* 64(4):969–993.

Wlezien, Christopher and Robert S. Erikson. 2007. "The Horse Race: What Polls Reveal as the Election Campaign Unfolds." *International Journal of Public Opinion Research* 19(1):74–88.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion.* Cambridge: Cambridge University Press.