

What Do We Know at 7 PM on Election Night?

ANDREW GELMAN

Columbia University
New York, NY 10027
gelman@stat.columbia.edu

NATE SILVER

FiveThirtyEight.com
nsilver@baseballprospectus.com

On the evening of November 8, 1988, I (AG) was working with my colleague Gary King in his Harvard office. It was election night, and Massachusetts Governor Michael Dukakis was a candidate for President against Vice President George H. W. Bush. Gary somehow had gotten his hands on a pair of tickets to Dukakis's victory party in Boston, and we were trying to decide whether to go. Dukakis was expected to lose, but—who could say, right? We had the TV on, and the first state to report, at 7 PM, was Kentucky, which Bush won by over 10 points. Gary informed me that the election was over: Kentucky, at the time, was near the political center of America, and there was no way that Dukakis would do much better nationally than he did in Kentucky. So we saved ourselves a subway ride and kept on working.

What about the most recent election night—November 4, 2008, when the candidates were Senators Barack Obama and John McCain? Was it possible for a viewer to play along at home with the election and decide at 7 PM, or 8 PM, or 9 PM what the outcome would be? On the night before the election we (AG and NS) analyzed a probabilistic election forecast to make some guesses at what might be known at different times during the evening. This article is a report of that analysis. Most of the results presented were obtained on November 3, the night before the election, and were intended as a guide for interpreting what we would hear the next night, as the votes were counted.

We performed one set of calculations using statewide vote margins; that is, reports like we remembered from Kentucky in 1988; and another set of calculations using only the tally of states won or lost, without the margins of victory.

In the next sections we describe the model we used, and then present our results from November 3. We then discuss what actually happened and ways in which election-night reporting could be improved in the future.

Our model on the night before the election

If one wants to make probabilistic forecasts, one needs a model. We used the model developed by one of us (NS) to make election forecasts at the website fivethirtyeight.com. It is described at that site [6] but it is too rich a model for us to describe in detail here. We will describe it in general terms, and discuss the tools one might use to develop a similar model of one's own.

A joint probability distribution The model takes the form of a joint probability distribution on 51 variables, one corresponding to a state's popular vote result. (We are treating the District of Columbia as a state, since it has electoral votes.) For each $j = 1, \dots, 51$, let Z_j represent the margin of victory, expressed as a percentage of the two-party popular vote, in the j -th state. We take the margin as positive if Obama

wins and negative if McCain wins (an arbitrary choice) and, for convenience, we express the value of Z_j in percentage points. For example, if McCain wins Kentucky (the eighteenth of the 51 states on our list) by a margin of 58.2% to 41.8% (counting only these two candidates' votes) then we say that $Z_{18} = -16.4$. The outcome of the election is summarized by the vector of random variables $\mathbf{Z} = \{Z_1, \dots, Z_{51}\}$. If we know the values of these variables, then we know the number of electoral votes for each candidate, as well as (with reasonable assumptions about voter turnout in each state) the national popular vote margin.

In the next section we will describe how one might build a probabilistic model of this form, combining the best available mean forecasts of the variables with a realistic error model. In this section we describe how we used the model we had.

We used the model as the basis for a simulation. In the language of statistics, this is called a *Monte Carlo* method. We drew random vectors $\mathbf{z}_i = \{z_{i,1}, \dots, z_{i,51}\}$ for $i = 1, \dots, 10000$, independently from the joint distribution defined by the model. By "independently" we mean that the 10,000 vectors were independent of each other—not that the variables were independent, since their correlations were dictated by the model. These 10,000 independent simulations amounted to a discrete approximation to the essentially continuous distribution of election outcomes.

All of the results in this paper were derived from this 10000×51 matrix of numbers. Each row represents a possible result of the election. In effect, we act as if these are the only possible results, and that each row of the matrix is equally likely to occur. The full matrix represents the state of our knowledge just before the polls close. (Elsewhere we have analyzed a similar set of simulations to estimate the probability that a single vote in any given state would be decisive in determining the outcome of the election [4]).

For our calculations below, we will need to not merely average over our 10,000 simulations but also to use them for conditioning—that is, making probabilistic inferences about future events given information that has already occurred, or has been assumed to occur. We implement the conditional probability formula $P(A | B) = P(A \cap B) / P(B)$ in a numerical computation fashion by keeping only the simulations consistent with the conditioned event.

For instance, suppose we were to learn at 7 PM that McCain has won Kentucky by 14 percentage points; in that case, what can we say about the outcome in the remaining states? What we can do is restrict our inferences to those simulations in which $z_{18} = -14\%$. Because of the discreteness of the simulations, we have to allow some slack in the computations; instead of conditioning on the precise event $z_{18} = -14.0\%$, we would keep all simulations in which z_{18} is between -15.0% and -13.0% . The smoothness of the underlying distribution ensures that our results computed in this way are virtually identical to those that would be obtained by exact calculations on the continuous space.

How can one build a model to forecast the election?

There are several different ways that one could construct a probabilistic election forecast (see Campbell [1] and Erikson and Wlezien [2] for recent reviews). Three natural approaches are: (1) aggregation of pre-election polls, (2) forecasts based on what are sometimes called "the fundamentals" (forecasting the national vote based on the state of the economy and making state-level adjustments based on long-term trends in voting patterns), and (3) judgments of political experts.

The model we used is based on a combination of these sources of information, but primarily the recent pre-election polls. When considering forecasts made months

before Election Day, there is a distinct difference between the snapshots obtained by poll aggregation and the predictions obtained by historical forecasting models. In discussing a particular method for incorporating public opinion surveys into forecasting models, Lock and Gelman [5] find that early polls provide some information about the relative positions of the states in the final voting but, when it comes to the national vote, they reveal essentially nothing beyond forecasts based on the fundamentals. The day before the election, the balance between polls and external forecasts is slightly different, because the polls tend to converge to the true election outcome (Gelman and King [3]; see also FIGURE 3 below). This is why we can pretty much make our election-day forecast based on aggregated polls, with slight adjustments based on a regression model and simulations capturing state-level and national uncertainty. The third part of the forecasting process—judgments of political experts—comes in the form of adjustments to the raw survey numbers.

Suppose you wanted to create your own forecasts of a national election, which can be studied in aggregate (using national polls) or separately (with state polls). How would you do it? To start, you might want to put together these sources of information:

- National polls. Average them how you like, fit a trend line if you want, but in any case use these to predict the aggregate vote.
- A national-level forecast. It makes sense to pull your poll aggregate toward whatever external forecasts you have (for example, based on the historical relation between the economy and election outcomes).
- Historical patterns of state voting relative to the national average. If nothing else, you can take your national forecast and perturb it, adding and subtracting percentages for each state to get an estimated state-by-state vote assuming a uniform national swing from past elections.
- State polls. You can use these to adjust state-level estimates. There are different ways to actually combine all these numbers, but the key is to separate the national forecast from the relative positions of the states.

Now you have a point estimate—which becomes an estimate of the mean of the vector $\{Z_1, \dots, Z_{51}\}$. The next step is to add variation to capture uncertainty in the forecast. It is natural for this variation to take the form of a covariance matrix for these 51 variables. You want for the correlation matrix to reflect uncertainty at the national, regional, and state levels.

Our particular simulations are based on a poll aggregation method developed for the website fivethirtyeight.com and involving several layers of adjustments (Silver [6]). What is relevant for the present article is that our 10,000 simulations represent forecast uncertainty about the election, as expressed in a joint distribution of the candidates' vote shares in the fifty states.

The vote shares in these distributions are correlated—for example, if Barack Obama had outperformed the average forecast in Ohio, we would expect (with some level of uncertainty) that he would've outperformed the forecast in Oklahoma, Oregon, or any other state. This is a sensible property for a set of forecasts, and it implies that our inferences for the outcome of the election in any state is affected by our knowledge of what happened in any other states where the polls closed earlier.

As we demonstrate below, we can use the 10000×51 matrix of state-level forecasts to make conditional inferences, to step through the hours of election night making our best prediction for the future based on the data available at any given time. Whatever the form of the probability distribution you select, the computations can be performed most transparently and generally by operating directly on the simulations.

The prior distribution Let's start by asking what we could say, based on this model, before any election returns were in. In Bayesian statistics this is called the *prior distribution*, and we can describe it by averaging all 10,000 simulations. By this method Obama was expected to win the national popular vote by 4.8 percentage points, with an expected electoral vote total of 340 (compared to 270 needed for victory), and a 96% chance of winning the electoral vote (with a 0.2% chance of a tie in the electoral college). The top row of FIGURE 1 shows the forecast distributions of Obama's popular and electoral vote shares.

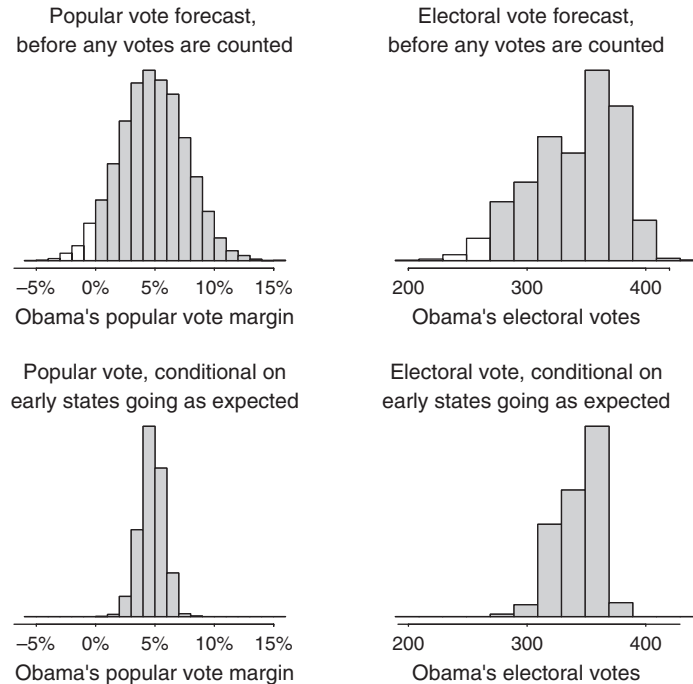


Figure 1 Uncertainty distribution for the presidential election outcome, expressed as Obama's popular vote margin and his electoral vote total. Top row is based on a poll-based forecast the day before the election; bottom row is based on these forecasts, conditional on the states whose polls close at 7 PM going as expected. The averages of the distributions in the top and bottom row are the same, but the distributions on the bottom show less variation: the vote margins of the early states tell us a lot.

Our model on election night

In 2008 the polls closed at 7 PM Eastern Time in six states, which we list in decreasing order of Obama's predicted vote margin: Vermont (Obama predicted to win by 21%), Virginia (+5%), Indiana (-2%), Georgia (-5%), South Carolina (-11%), and Kentucky (-15%). (The map has certainly changed since the days when Kentucky was a swing state!)

Real-time predictions given vote margins in the early states We start by assuming that the viewer would know the vote margin (either as exactly tabulated or as estimated from exit polls) in each of these states, which can be summarized by a simple unweighted average of Virginia, Indiana, Georgia, South Carolina, and Kentucky. (We excluded Vermont because it is the smallest of these six states and farthest from the national median.) Based on the last pre-election poll aggregates, the estimated five-state average vote margin was -5.7%; that is, McCain was expected to beat Obama by

an average of 5.7 percentage points on the way to winning Kentucky, South Carolina, Georgia, Indiana, and losing Virginia.

If the expected happens, what have we learned? That is, what if the vote margin in Virginia, Indiana, Georgia, South Carolina, and Kentucky were to equal the expected -5.7% ? We pipe this assumption through our model by calculating, for each of our 10,000 simulations, the average vote margin in these five states, and then restricting our analysis to the subset of simulations for which this vote margin is within 1 percentage point of its expected value (that is, between -6.7% and -4.7%). Out of our 10,000 simulations, 2800 fall in this range; that is, we predict there is a 28% chance that McCain's average vote margin in these five states will be between 4.7% and 6.7% . What is of more interest is what happens if this occurs. Considering this subset of simulations, Obama's expected national popular vote margin is $+4.7\%$, his expected electoral vote total is 343, almost the same as the prior distribution. But now the conditional probability of an Obama victory is 100%: he wins the electoral college in all 2800 simulations in this condition. The bottom row of FIGURE 1 shows the forecast distributions of Obama's popular and electoral vote shares, conditional on him doing exactly as expected in the first round of states.

So if the 7 PM states were to go as expected, we would know a lot.

How about other possibilities? We repeat the above calculation under scenarios in which Obama's average vote margin in Virginia, Indiana, Georgia, South Carolina, and Kentucky takes on each possible value between -12% and $+1\%$; based on our simulations, there is a 97% chance that the average 7 PM vote margin (excluding Vermont) will fall in this range.

In the unlikely event that McCain were to get an average vote margin of 12 percentage points in the five 7 PM states, we would have forecast a 1.3% margin for McCain in the national popular vote, an expected 265 electoral votes for McCain, and a 38% chance of him winning the electoral college (with a 60% chance of Obama winning and a 2% chance of a tie).

At the other extreme, if Obama happened to get an average vote margin of 1 percentage point in these early states, we would have predicted his national popular vote margin to be 10% with 394 electoral votes and a 100% chance of winning.

What about the possibilities in between? FIGURE 2 shows Obama's expected popular vote margin, the expected division of electoral votes, and the probability of each




Figure 2 Election predictions at 7 PM Eastern time, after the polls were to close in Virginia, Indiana, Georgia, South Carolina, and Kentucky. Obama's average vote margin in these five states was predicted to be -5.7 percentage points, but it could plausibly have fallen between -12 and $+1$ percentage point. For each of these possible outcomes, we computed Obama's expected share of the national popular vote, his expected electoral vote total, and the probability he would win in the electoral college. Unless McCain's average vote margin in these five states was at least 9 percentage points, we could confidently call the election for Obama at this point.

candidate winning under the different scenarios. According to these simulations, if McCain hadn't secured an average vote margin of at least 9 percentage points in Virginia, Indiana, Georgia, South Carolina, and Kentucky, he could pretty much have thrown in the towel. And if Obama had lost these states by an average of more than 10 points, we might have been up all night worrying (rather than celebrating or crying).

Real-time predictions given only the state winners What if the TV stations do not give vote margins but just report winners? With six states reporting at 7 PM, there are 64 possible outcomes. Some of these are impossible or uninteresting, however: if McCain had won Vermont, or if Obama had won South Carolina or Kentucky, the election would have been over. What remains are Virginia, Indiana, and Georgia.

TABLE 1 gives the eight possibilities, our forecast probability of each happening, and Obama's expected popular vote margin, electoral vote margin, and win probability under each scenario. The only interesting possibility is if McCain had swept all three states: then it would have been a contest, and he would have had an even chance of pulling it out.

TABLE 1: Scenarios of interest at 7 PM Eastern time, with the first states reporting. In order to have a chance, McCain needed to win Virginia in this first round (in which case he would almost certainly have won Indiana and Georgia as well).

VA	IN	GA	Prob. of scenario	Pop. vote margin	Electoral vote		Probability of...		
					Oba	McC	O.win	M.win	Tie
McC	McC	McC	7%	+0.2%	268	270	.47	.49	.04
McC	McC	Oba	0						
McC	Oba	McC	0						
McC	Oba	Oba	0						
Oba	McC	McC	66%	+4.0%	330	208	.99	.01	.00
Oba	McC	Oba	<1%						
Oba	Oba	McC	22%	+7.4%	375	162	1.00	.00	.00
Oba	Oba	Oba	5%	+9.9%	399	139	1.00	.00	.00

We would then have had to wait until 7:30 PM Eastern time, when we would hear from Ohio, North Carolina, and West Virginia. Our pre-election forecasts gave Obama less than a 1% lead in North Carolina and a 2% lead in Ohio, with McCain having a 10% lead in West Virginia. But in the (unlikely) event of McCain sweeping Virginia, Indiana, and Georgia, the story would change. At that point, Obama would be expected to lose Ohio, North Carolina, and West Virginia, by margins of 3, 5, and 15 percentage points, respectively. If McCain's chances were still alive at 7 PM, there was an 87% chance he'd win all three of these must-win states that close the polls at 7:30.

Onward to 8:00, when most of the remaining eastern states closed the polls. The 8 PM states range from Maryland and Massachusetts (where Obama was forecast to win by 23 and 19 points, respectively) to Oklahoma and Alabama (predicted McCain victory margins of 27 and 24 points). Conditional on McCain winning the key early states of Virginia, Indiana, Ohio, and North Carolina, these predictions shift by about 5 percentage points in his favor.

The states to watch at 8 PM—if there was anything worth watching at all—would be New Hampshire, Pennsylvania, and Florida, for which McCain's predicted vote margins, conditional on his previous success, would be -3%, -2%, and +4%. At this point, winning Pennsylvania would pretty much guarantee victory for McCain; his other possibilities are winning New Hampshire and Florida (which would give him an expected 277 electoral votes and a 79% chance of winning, with an amazing 12% chance of a tied electoral college), and winning Florida alone (the most likely

possibility, with an 8% chance of happening) which would take him to an expected 264 electoral votes and a 33% chance of winning.

At this point, McCain would have an 80% chance of winning both Missouri and Florida, which would move him up to an expected 281 electoral votes and a 66% chance of winning (compared to 30% for Obama and a 4% chance of a tie, at this point). The next most likely possibility is McCain winning Missouri but losing Florida, in which case his expected electoral vote count drops to 243 and his probability of winning declines to 3%.

If McCain were to win Virginia, Indiana, Ohio, North Carolina, and Florida, while losing Pennsylvania, we would have to keep the TV on. The news at 8:30 wouldn't help much: at this point, McCain would be expected to win handily in Arkansas, the only state to close the polls at that time. At 9 PM we would hear from a bunch of states further west, including Colorado (with Obama expected to win by less than 1% at this point), New Mexico (Obama expected to win by 4%), Minnesota (Obama by 4%), Wisconsin (Obama by 5%), and Michigan (Obama by 6%). If Obama won all five of these, he would have a 97% chance of winning. If McCain won any of them (most likely Colorado), he would have been almost home free, with a 90% chance of an electoral vote win, and if McCain had won two or more, he would have basically won the election.

Again, though, based on the polls the day before the election, we only estimated a 4% chance overall of this happening. The most likely outcome a priori was that the election would be over by 7 PM.

Discussion

With the ubiquity of polling and the rise of internet communication, state-level election forecasts and public opinion estimates are increasingly available to the general public. Attention has generally focused on point predictions and the odds of each candidate winning; here we demonstrate how the correlations in a forecast implicitly allow inferences for the election outcome, conditional on partial information. This is straightforward probability theory, or Bayesian inference; here we are applying standard practice and summarizing inferences using simulations. With 10,000 simulations, we can compute all the relevant conditional probabilities directly. If we wanted to work out some obscure possibilities (for example, Obama winning Virginia and Georgia with McCain winning Indiana), we would need to produce more simulations or else use some analytical approximations, but here we pursued only the more plausible scenarios.

Beyond whatever interest there may be in election-night dynamics, this work is relevant to understanding election forecasts and, more generally, inference about vector outcomes in which there is correlation (so that, for example, the outcome in Virginia is informative, not just about that state, but about national and regional swings). As we have demonstrated, simulation-based calculation allows us to condition on virtually any plausible outcome.

On a practical level, we are only interested here in moderate or large probabilities—we are not trying to pick out 1-in-1000 longshots—and so we only need to condition on events that have a nontrivial forecast probability of occurring. For example, we could safely ignore Vermont in the first set of results, and look only at the aggregate margin in the other five states—and then we only needed to look at one-percent intervals even though (according to the story) we would be observing exact values. These choices are not lazy but rather reflect a realistic understanding of the problem we are studying. If there is any worry on this point, one could always repeat the analysis with 10 or 100 times as many simulations and increase the precision.

What actually happened?

Obama won Vermont and Virginia, and McCain won Georgia, South Carolina, and Kentucky as predicted. As a small surprise, Obama won Indiana. The actual margins for Obama in these six states were Vermont +37.8%, Virginia +6.4%, Indiana +1.0%, Georgia -5.3%, South Carolina -9.1%, Kentucky -16.5%. The average for the last five states was -4.7%. Had we known that average, or just the fact that Obama won Virginia, we could have been nearly certain that Obama would win the election.

Of course, all of our inferences are only as good as the forecasting model. Unsurprisingly, given that the forecasts were based on the latest polls, they did pretty well; see FIGURE 3. And, in a sense, the election really was over by 7 PM—although it was not so easy to learn this from watching the networks' broadcasts.

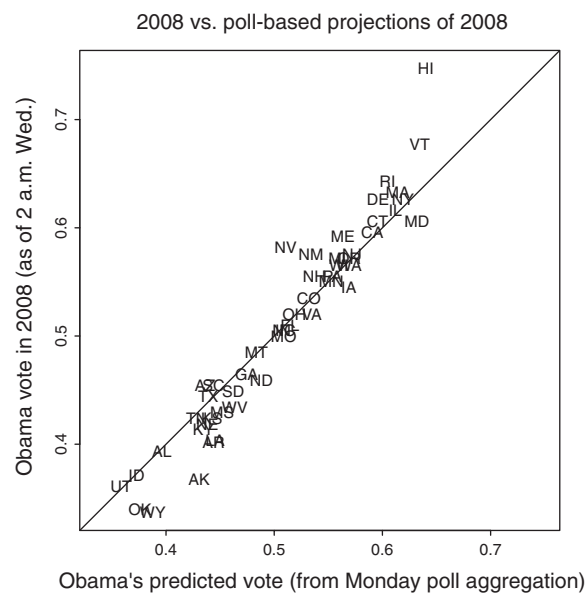


Figure 3 Obama's share of the two-party vote in each state, plotted vs. poll-based predictions made the day before. The poll-based model gave accurate forecasts of the national vote and also most individual states, although with a systematic pattern of underpredicting Obama's vote share where he was expected to do well, and overpredicting in most of the states where he was expected to lose.

Our conditional probability calculations worked reasonably well too—in theory. In practice, though, our advice on what to look for on election night was useless, because the TV networks did not report total vote margins and they did not immediately declare winners. Instead, they reported partial results as they came in from each state (14% of the precincts in Virginia, 12% of the precincts in Indiana, and so forth). As the tallies mounted it became clearer who would win each state, but the broadcasters missed out on the opportunity to provide additional information in the form of conditional probabilities.

In particular, what if the newscasters, after reporting the election results based on 14% of the precincts reporting in Virginia, were to also tell us the results in the previous election in those same precincts? The viewers would then know the swing in vote proportions, which would be much more informative than the raw numbers. (When watching the election returns on TV, we heard occasional comparisons with the 2004 Kerry-Bush numbers, but only sporadically.) Our suggestion would give the informa-

tion needed to make inferences about the national election without the speculation involved in calling states early based on exit poll results.

For future elections it would be fun to set up an online widget so that users could enter election returns as they are happening, and the relevant probabilities would pop out. It would also be desirable to connect this to election returns by county and even precinct. TV networks aren't supposed to make early calls of the election but maybe there's be some way of doing this informally.

Acknowledgment We thank the Columbia University Applied Statistics Center, National Science Foundation, National Institutes of Health, National Security Agency, Department of Energy, Institute for Educational Sciences, and Yahoo Research for partial support for this research.

REFERENCES

1. J. E. Campbell, Forecasting the 2008 national elections, *PS: Political Science and Politics* **41** (2008) 679–681.
2. R. S. Erikson and C. Wlezien, Leading economic indicators, the polls, and the presidential vote, *PS: Political Science and Politics* **41** (2008) 703–707.
3. A. Gelman and G. King, Why are American Presidential election campaign polls so variable when votes are so predictable? *British Journal of Political Science* **23** (1993) 409–451. doi:10.1017/S0007123400006682
4. A. Gelman, N. Silver, and A. Edlin, What is the probability your vote will make a difference? *Economic Inquiry* (to appear).
5. K. Lock and A. Gelman, Bayesian combination of state polls and election forecasts, *Political Analysis* (to appear).
6. N. Silver, Frequently asked questions, <http://www.fivethirtyeight.com/2008/03/frequently-asked-questions-last-revised.html>.


Summary We use a probability forecasting model to estimate the chance of different branches on the tree of state-by-state outcomes on election night. Forecasting models can use data from pre-election surveys as well as extrapolation based on previous election results. We implement conditional probability calculations numerically using a matrix representing 10,000 simulations of the outcomes in the 50 states.

ANDREW GELMAN is a professor of statistics and political science at Columbia University. His books include “Bayesian Data Analysis,” “Teaching Statistics: A Bag of Tricks,” and “Red State, Blue State, Rich State, Poor State: Why Americans Vote the Way They Do.”

NATE SILVER is a statistician at Baseball Prospectus and the founder of FiveThirtyEight.com.

Yes!

Follow the complete 4-graphs shown in bold.



—Suzanne Dorée