

The Future of Election Forecasting: More Data, Better Technology

Drew A. Linzer, *Emory University*

Toward the end of the 2012 US presidential campaign, public poll aggregators converged on a forecast that President Obama would be reelected with victories in the same states he won in 2008, except Indiana and North Carolina, with Florida too close to call.¹ Meanwhile, most pundits and commentators took a contrary view: the race as a whole was no better than a toss-up. Some attacked the quantitative approach to election forecasting altogether (Nyhan 2012). But the poll-based predictions turned out to be correct—even about Florida’s closest-in-the-nation election outcome. The ensuing media coverage took on an almost triumphal tone: “The Poll Quants Won the Election,” declared a headline in *The Chronicle of Higher Education* (Bartlett 2012). *The New Yorker* called it a “Victory for the Pollsters and the Forecasters” (Cassidy 2012). In *The Economist* (2012), it was a “March of the Nerds.”

How satisfied should political scientists—or, for that matter, anyone interested in rigorous approaches to politics and elections—actually be about this accomplishment? The answer depends on what we consider to be the purpose of election forecasting. Careful modeling of trial-heat survey data certainly beat the intuitions of political pundits (Greenfield 2012). It also affirmed the credibility of public opinion surveys as a method for learning about voter preferences (Blumenthal 2012). On the other hand, forecasts issued just on the eve of the election arrived too late to help reporters, news consumers, or researchers understand how and why voters’ preferences had responded to the campaign. Nor were they much use to candidates and party strategists who needed to decide how to allocate resources for organizing, advertising, and voter mobilization much earlier in the election season.

Accurate election forecasts have considerable value to political observers as well as scholars and practitioners. To maximize these benefits, however, forecasting models need to generate predictions early in the race—preferably, three to four months before Election Day. They should also be accompanied by informative statements of uncertainty. Initial forecasts can be updated with newer information as the election nears; for example, from the results of public opinion polls. In addition, forecasting models should strive to generate predictions at the same level (state or district) as the election of interest. National forecasts that aggregate lower-level results are less useful than district-by-district predictions. And although most current forecasting models focus on the US presidential race, researchers should be exploring how to extend existing methods to other types of races and to contests outside the United States.

By any of these standards, forecasting elections is still far from a “solved” problem. In the lead-up to the 2012 election, Campbell (2012) catalogued twelve regression-based forecasting models that extrapolated from past presidential election outcomes to predict the upcoming vote. Only seven correctly foresaw an Obama victory. Among the incorrect forecasts, two gave Obama as little as a 10% or 11% chance of winning.² Nine models only offered predictions of national, rather than state, vote outcomes. None of the twelve models provided any mechanism for correcting inaccurate estimates closer to Election Day. The performance of presidential election forecasting models was somewhat better in 2008 (when Obama won by a larger margin), but completely missed the narrow victory of George W. Bush in 2000 (Campbell 2001; 2008). Forecasting models for US congressional and gubernatorial elections have faced similar challenges (e.g., Klarner 2009; Lewis-Beck and Rice 1984; Peltzman 1987).

The issue is not that the theories of voter behavior or campaign dynamics underlying any of these models are badly flawed or even terribly incomplete. Political scientists have long been aware that, broadly speaking, incumbents fare better when the economy is improving and when they (or their party) are viewed more positively in the electorate (e.g. Lewis-Beck and Stegmaier 2000; Nannestad and Paldam 1994). Factors related to candidates’ incumbency status, the prevalence of identity-based voting, and variation in electoral rules and governing institutions may attenuate or amplify the effects of these core structural variables (Duch and Stevenson, 2008).

The difficulty, instead, is in operationalizing and measuring each of these factors and figuring out how much each one contributes to the forecast. For example, should an “improving economy” be interpreted as increasing gross domestic product, rising household income, falling unemployment rates, or something else? If more than one variable is used to predict the election outcome, then which ones, and in what combination? Political science theories are rarely specific enough to say, but the small size of most election datasets precludes an empirical solution to the problem. Without more data, selecting variables based on the strength of their observed association with historical election outcomes runs the risk of over-fitting the forecasting model and degrading its ability to predict out-of-sample. Ultimately, a substantial portion of any election outcome is going to be random and unpredictable, no matter how good election science becomes.³ Elections are complex, stochastic events. This “noisiness” not only limits the accuracy of forecasts from even well-specified models but also makes it nearly impossible to adjudicate between alternative model specifications.

Different models can, and will, generate highly divergent election forecasts even if they are equally justifiable from a theoretical standpoint.⁴

Fundamentals-based election forecasting is running into the limits of what additional *theory* is going to contribute. The greatest impediment to the development of better election forecasting models is not a lack of theory; it is a lack of data. Short of waiting another 50 years until the signal-to-noise ratio in election data tilts somewhat more favorably in our direction (and hoping that in the meantime, election conditions remain consistent), forecasters must find new and better sources of data to inform their predictions. Advancements in election forecasting will come from researchers who identify—or gather—these

Frank Newport remarked that survey aggregators “don’t exist without people who are out there actually doing polls,” and that aggregation threatens to dissuade survey organizations from gathering these data in the first place (*Marketplace* 2012).

Beyond the polls, how can forecasters supplement historical election data? Some of the more exploratory recent efforts include data from prediction markets, which combine the judgments of large numbers of individuals about the likelihood of a future event; social media data; Internet usage or search data; or patterns of online political activism. We are only beginning to study how these data are generated, how they relate to election outcomes, and how they can be integrated into election forecasts. The potential value in these new types of information is not

Fundamentals-based election forecasting is running into the limits of what additional theory is going to contribute. The greatest impediment to the development of better election forecasting models is not a lack of theory; it is a lack of data.

data, devise and test theories about how the data relate to the election outcomes we care about, and build the modeling technology to produce forecasts in a timely manner.

MORE DATA: POLLS... OR WHAT?

One of the most promising new sources of data over the last few years has been the large number of trial-heat opinion polls that are conducted—and released publicly—by survey research firms, media organizations, advocacy groups, and others. During the 2012 US presidential campaign, the results of more than 1,200 state-level polls were published, representing more than one million Americans. In 2008, more than 1,700 state polls were made available. Hundreds more polls were conducted at the national level.

For forecasters and analysts, this increase has represented a tremendous breakthrough. By applying the basic principle that sampling error in the individual polls can be cancelled out by averaging the results of concurrent surveys, poll aggregators have been able to estimate smoothed trends in state- and national-level voter preferences during the campaign,⁵ and, in some cases, project these trend lines forward to Election Day (e.g., Linzer 2013; Silver 2012). The projections combine polling data with information from historical models, applying the principle that by Election Day voter preferences will “revert” toward the outcome implied by the election fundamentals (Kaplan, Park, and Gelman 2012). Another variant of this procedure shrinks the forecasts toward a tied result, which helps prevent overconfidence.

The long-term availability of public trial-heat polling data can not be taken for granted, however. Public polling is completely decentralized and growing more expensive. Fewer state polls were conducted in 2012 than in 2008. At least 20 polls per state are needed in the final three months of the campaign to generate consistently accurate forecasts, but most states in 2012 had nowhere near that many.⁶ In addition, within the polling industry there are complaints that aggregators are exploiting pollsters’ data. After the 2012 election Gallup editor-in-chief

primarily in testing theories of voter behavior (although accurate forecasts can reflect positively on the underlying theory), but rather in measuring voter preferences and forecasting outcomes in ways that could not be done before.

IMPROVING TECHNOLOGY TO DO BETTER SCIENCE

There are many questions in the study of voting, campaigns, public opinion, and elections that political scientists could investigate more thoroughly with additional behavioral data and improved campaign monitoring and forecasting technology: Why do people vote the way they do? Why do people turn out to vote at all? When do people decide who they will vote for, and why? What types of voters change their minds during a campaign? What effect do campaigns have on the election outcome? The technology might be a statistical forecasting model, a measurement technique, or any systematic approach to extracting quantitative information from noisy election data. The aim is to put theories from political science to practical use.

A burst of innovation along these lines has recently occurred in the professional political arena (Issenberg 2010, 2013). An empirical, experimental approach to understanding (and, sometimes, guiding) political behavior is at the core of the research being done at organizations such as the progressive Analyst Institute, or inside the analytics department of the 2012 Obama presidential campaign (Issenberg 2012; Scherer, 2012). The Republican National Committee recently hired its first-ever chief technology officer to revamp the party’s digital campaign infrastructure (Ward 2013). A host of other data-driven consulting firms have sprung up since the 2012 elections on both the Democratic and Republican sides (Wilner 2013a).

While the motivations and objectives of academic political scientists clearly differ from those of partisan political strategists, the data analytic methods and approaches that each use overlap. Political science can be a model for how this research and technological development proceeds. The academic field,

unlike the world of consultants, campaigns, and media pundits, has norms favoring transparency and replicability. Statistical models entail assumptions; a lack of transparency prevents us from evaluating those assumptions in any meaningful or constructive way, determining how sensitive the conclusions may be to particular methodological choices, or learning more about the relationships an analyst claims to be seeing in the data. There is less reason to trust research that is overly secretive or “proprietary.” Especially in politics, an open research model not only advances the science; it helps avoid charges of manipulation or bias (most likely from whichever side a forecaster is predicting to lose). Scholars who contribute to public discourse around campaigns and elections also offer a credible counterpoint to the often exaggerated or misinformed claims of political pundits.

Looking ahead, proving the value of political science to electoral politics might promote opportunities for similar contributions in other areas. Debates over international security, inequality, criminal justice, education, health, and the environment—to name only a few—can all benefit by drawing on empirical scientific evidence. The debate over whether sophisticated quantitative political research can be “relevant” to contemporary politics should be settled. ■

NOTES

1. These efforts included the work of Nate Silver at *The New York Times*; Sam Wang at the Princeton Election Consortium; Josh Putnam at FrontloadingHQ; the websites realclearpolitics.com, elections.huffingtonpost.com, and polltracker.talkingpointsmemo.com; and my own at votamatic.org, where I published state and national level presidential election forecasts, and tracked voter opinion, based on research in Linzer (2013).
2. Lauderdale and Linzer (2013) suggest that very low probabilities such as these badly understate both specification and estimation uncertainty in fundamentals-based forecasting models.
3. Most input variables are also measured with error: the results of trial-heat polls, or economic data that do not become “final” until revisions months after the campaign ends.
4. Montgomery, Hollenbach, and Ward (2012) discuss statistical approaches for improving forecasts by combining the predictions of multiple model specifications.
5. A common misconception is that poll aggregation is intended to allow “sound surveys to compensate for sketchier ones,” as suggested by Wilner (2013b). No single poll is either “right” or “wrong;” all polls contain sampling error—even the most methodologically rigorous. One benefit to aggregation is that it can help identify pollsters whose results contain systematic errors. If aggregation can cancel out individual firms’ house effects as well, that is a bonus.
6. For example, see <http://votamatic.org/final-result-obama-332-romney-2016>.

REFERENCES

- Bartlett, Tom. 2012. “The Poll Quants Won the Election.” *The Chronicle of Higher Education*. <http://chronicle.com/blogs/percolator/the-poll-quants-won-the-election/31722>.
- Blumenthal, Mark. 2012. “2012 Poll Accuracy: After Obama, Models and Survey Science Won the Day.” *The Huffington Post*. http://www.huffingtonpost.com/2012/11/07/2012-poll-accuracy-obama-models-survey_n_2087117.html.
- Campbell, James E. 2001. “The Referendum that Didn’t Happen: The Forecasts of the 2000 Presidential Election.” *PS: Political Science & Politics* 34 (1): 33–38.
- . 2008. “Editor’s Introduction: Forecasting the 2008 National Elections.” *PS: Political Science & Politics* 41 (4): 679–82.
- . 2012. “Forecasting the 2012 American National Elections.” *PS: Political Science and Politics* 45 (4): 610–13.
- Cassidy, John. 2012. “Cassidy’s Count: A Victory for the Pollsters and the Forecasters.” *The New Yorker*. <http://www.newyorker.com/online/blogs/johncassidy/2012/11/cassidys-count-how-obama-won-how-romney-lost.html>.
- Duch, Raymond M., and Randolph T. Stevenson. 2008. *The Economic Vote: How Political and Economic Institutions Condition Election Results*. Cambridge: Cambridge University Press.
- Greenfield, Rebecca. 2012. “The Best and Worst Pundit Predictors of 2012.” *The Atlantic Wire*. <http://www.theatlanticwire.com/politics/2012/11/best-and-worst-pundit-predictors-2012/58846>.
- Issenberg, Sasha. 2010. “Nudge the Vote.” *The New York Times Magazine*. <http://www.nytimes.com/2010/10/31/magazine/31politics-t.html>.
- . 2012. “How President Obamas Campaign Used Big Data to Rally Individual Voters.” *MIT Technology Review* <http://www.technologyreview.com/featuredstory/509026/how-obamas-team-used-big-data-to-rally-voters>.
- . 2013. *The Victory Lab: The Secret Science of Winning Campaigns*. Broadway Books.
- Kaplan, Noah, David K. Park and Andrew Gelman. 2012. “Understanding Persuasion and Activation in Presidential Campaigns: The Random Walk and Mean Reversion Models.” *Presidential Studies Quarterly* 42 (4): 843–66.
- Klarner, Carl E. 2009. “Forecasting Congressional Elections.” *Extension of Remarks; Newsletter of the APSA Legislative Studies Section*.
- Lauderdale, Benjamin E., and Drew A. Linzer. 2013. “Under-performing, Over-performing, or Just Performing? The Limitations of Fundamentals-based Presidential Election Forecasting.” Presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL.
- Lewis-Beck, Michael S., and Mary Stegmaier. 2000. “Economic Determinants of Electoral Outcomes.” *Annual Review of Political Science* 3: 183–219.
- Lewis-Beck, Michael S., and Tom W. Rice. 1984. “Forecasting U.S. House Elections.” *Legislative Studies Quarterly* 9 (3): 475–86.
- Linzer, Drew A. 2013. “Dynamic Bayesian Forecasting of Presidential Elections in the States.” *Journal of the American Statistical Association* 108 (501): 124–34.
- Marketplace. 2012. “Post-election, a Polling Conundrum: Interview with Frank Newport.” <http://www.marketplace.org/topics/elections/attitude-check/post-election-polling-conundrum>.
- Montgomery, Jacob M., Florian M. Hollenbach, and Michael D. Ward. 2012. “Ensemble Predictions of the 2012 US Presidential Election.” *PS: Political Science and Politics* 45 (4): 651–54.
- Nannestad, Peter, and Martin Paldam. 1994. “The VP-function: A survey of the Literature on Vote and Popularity Functions after 25 Years.” *Public Choice* 79 (3–4): 213–45.
- Nyhan, Brendan. 2012. “Pundits versus Probabilities; The Misguided Backlash against Nate Silver.” *Columbia Journalism Review*. http://www.cjr.org/united_states_project/pundits_versus_probabilities.php.
- Peltzman, Sam. 1987. “Economic Conditions and Gubernatorial Elections.” *The American Economic Review* 77 (2): 293–297.
- Scherer, Michael. 2012. “Inside the Secret World of the Data Crunchers Who Helped Obama Win.” *Time: SwampLand*. <http://swampland.time.com/2012/11/07/inside-the-secret-world-of-quants-and-data-crunchers-who-helped-obama-win>.
- Silver, Nate. 2012. “Methodology.” *FiveThirtyEight*. <http://fivethirtyeight.blogs.nytimes.com/methodology>.
- The Economist. 2012. “Politics and Statistics; March of the Nerds.” <http://www.economist.com/blogs/democracyinamerica/2012/11/politics-and-statistics>.
- Ward, Jon. 2013. “Republican National Committee Hires Senior Facebook Engineer As Chief Technology Officer.” *The Huffington Post*. http://www.huffingtonpost.com/2013/06/04/republican-national-commi_n_3386575.html.
- Wilner, Elizabeth. 2013a. “The Cook Political Report ‘Big Data At-A-Glance.’” *The Cook Political Report*. <http://cookpolitical.com/story/5804>.
- . 2013b. “The Survey Monkey on Our Back: Where Our Addiction to Polls May Take Us.” *The Cook Political Report*. <http://cookpolitical.com/story/5777>.