Using Expert Assessments to Forecast the 2010 House Election

Sanford C. Gordon*

Associate Professor

Wilf Family Department of Politics

New York University

sanford.gordon@nyu.edu

September 28, 2010

Abstract

I propose a methodology for forecasting the 2010 House election using current and historical race-specific expert rankings. Using ratings from the Cook and Rothenberg Political Reports and adjusting to overcome a potential pro-Democratic bias, the method projects that the Democrats are likely to retain control of the House, albeit with a significantly narrower majority than currently enjoyed.

1 Introduction

Due largely to the absence of reliable, district-specific polling data, political scientists seeking to forecast the outcomes of House elections are often compelled to rely on information from national-level surveys. When combined judiciously with data on the observable features of congressional districts, forecasts that adopt this approach can be startlingly accurate. Using such an approach, for example, Bafumi et. al. (2008) forecasted a 32 seat gain for the Democrats in the 2008 election – a prediction off by only two seats.

A potential limitation of such an approach is that it is unable to take advantage of idiosyncratic district-level information possessed by professional observers of specific congressional races. For

example, using only observable measures of district competitiveness may lead an analyst to conclude that a race is in play, whereas a field-level observer might know that the incumbent was implicated in a major bribery scandal. In other case, the statistical analyst may lack data on the success of the out-party in recruiting quality challengers – something a field-level observer might judge to be important in his or her assessment of an individual race. In other words, the approaches typically taken by political scientists by necessity exclude the "on the ground" knowledge possessed by journalists and other veterans of political campaigns. At the same time, rational expectations suggest that such observers will take into account observable features of the environment held to be important by political scientists – although perhaps not in as systematic a fashion.

A number of organizations produce periodic assessments of individual House races based on such expert-level assessments. Most famous among these are the Congressional Quarterly Weekly Report's Race Ratings, the Rothenberg Political Report's House Ratings, and the Cook Political Report Competitive Seat Rankings. Because we can go back and examine the predictive success of *previous* expert predictions using past election results, it is possible, using the historical track record, to construct a forecast based on *current* expert projections. This approach can be seen as complementary to those such as the regression-based one employed by Bafumi et. al. To the extent that it yields projections similar to regression-based approaches, our confidence in the accuracy of both kinds of predictions increases. To the extent that its projections differ, then a comparison of the methods' performance is informative of the relative merits of the differing techniques.

In this note, I present the results of a simulation of House election outcomes based on past and current forecasts from the Cook and Rothenberg Political Reports.[1] Cook divides races into eight classifications: "solid," "likely," "leaning," and "toss-up" categories for both Democrats and Republicans. Rothenberg uses a similar scheme, but adds a category for "pure" toss-ups. The basic idea behind my approach is to treat the outcome of each current House race as the result of a weighted coin-flip, with the true weight hinted at by the predictive success of the ranking categories in previous election cycles. So, for example, suppose that in the past, Democrats have won seats classified as "leaning Democrat" 93% of the time. Then our best estimate of the chance that the

---

[1]Campbell (2010) and Silver (2010) also employ Cook Political Report projections to facilitate forecasting the 2010 House election, but their methodologies differs substantially from mine.

Democrat will win in a "leaning Democrat" district in the current race should be around 93%.[2] If we estimate weights for each of the Cook or Rothenberg classifications using historical rankings and election outcomes, we can let a computer flip the coins for each congressional district (depending on its classification) to come up with a hypothetical election outcome. Then, we can repeat the exercise, say, 10,000 times, to see what sorts of outcomes are more or less likely.

To benchmark the performance of this approach, I present results of simulations of the 2004, 2006, and 2008 elections (for which we know the outcomes). Two sets of results are presented: the projected number of Republican-held seats using the Cook and Rothenberg forecasts immediately preceding the election, and the projected number using forecasts available as of late-September for each of those years. The 2004 forecasts are calibrated using historical performance of rankings from 2002, the 2006 forecast from 2002 and 2004 rankings, etc. The benchmarking exercise implies that the approach would have worked quite well for the 2004 and 2008 cycles, but would have fared more poorly in 2006.

Subsequently, I present two forecasts of the 2010 election based on the Cook rankings from September 27, 2010 and the Rothenberg rankings from September 24. The first replicates the methodology described, while the second employs weights for "solid," "likely," "leans" and toss-up categories that do not distinguish by party. I adopt the second approach because the presence of numerous pro-Democratic surprises in 2006 and 2008 are likely to bias the 2010 forecast in a pro-Democratic direction. My results differ substantially from those presented by Bafumi et. al. While their forecast and mine predict sizeable Republican seat gain, my approach anticipates one insufficiently large to make a Republican takeover likely. In particular, the eight-category approach estimates an expected Republican gain of about 28 seats, and the four-category approach a gain of about 35 seats. By contrast, Bafumi et. al. anticipate a gain of around 50 seats.

Interestingly, these net seat projections differ from those offered by both Cook and Rothenberg – as of writing, Cook forecasts a Republican gain of "at least 40 seats," while Rothenberg projects a 37-42 seat-gain for Republicans. Cook attributes this difference to a strong likelihood that 2010 is a "wave" election year in which a substantial number of pro-Republican upsets are possible.

---

[2]The approach described in greater detail below also factors in the uncertainty regarding this and other probabilities.

Indeed, when I conduct my analysis assuming that Republicans will be as fortunate in 2010 as Democrats were in 2006 and 2008, a Republican takeover appears inevitable. However, I argue that this interpretation relies on the implicit assumption that the individual seat ratings do not (yet) accurately reflect the raters' prospective beliefs about the direction those races are heading (and thus the likelihood of Republican upsets). To the extent that the ratings fully reflect the raters' rational expectations, then my forecast is unbiased. To the extent that they do not, my forecast may significantly underpredict Republican gains.

2   A Bayesian Approach to Expert-Based Forecasting

Let c indicate race classification, with $c \in \{\text{Solid Dem., Likely Dem.}, \ldots, \text{Likely Rep., Solid Rep.}\}$. Let s_t^c be the number of successful predictions in classification c in races *prior to* election t, and f_t^c the corresponding number of failures. I will adopt the convention that every classification – even Solid Democrat – is a prediction that the *Republican* will win. This convention is adopted purely for notational convenience, and makes no difference for the results.[3]

Let $q_t^{i,c}$ represent the (Bernoulli) probability that the Republican in district i in classification c will win at election t. Then assuming a $U(0,1)$ prior distribution, $q_t^{i,c}$ follows a beta distribution with parameters $(1 + s_t^c, 1 + f_t^c)$.

Each iteration of the simulation proceeds in three steps, and represents a single simulated election summarized by the number of Republican seats won. Let r_i be an indicator equal to one if a simulated Republican wins in district i and zero otherwise, and R be the sum of Republican wins across districts.

(1) For all classifications c, for each seat i, draw $q_t^{i,c}$ from $Beta(1 + s_t^c, 1 + f_t^c)$.

(2) For each seat i, draw r_i from $Bernoulli(q_t^{i,c})$.

(3) Calculate $R = \sum_{i=1}^{435} r_i$.

This algorithm may be repeated an arbitrary number of times to produce a simulated probability distribution for R, from which quantities of interest (e.g. mean, median, standard error, etc.) may

---

[3]One way to think about this is that if we had a machine that predicts when a coin will turn up tails with incredible accuracy, we can use it to predict heads with the same degree of certainty.

be calculated.

## 3 Historical Results Using Immediate Pre-Election Projections

As a benchmarking exercise, I first employ the approach detailed in the preceding section using the Cook rankings from immediately preceding the 2002, 2004, 2006, and 2008 elections; and the Rothenberg rankings preceding the 2006 and 2008 elections. Because we know the outcomes of these races, we can assess how this exercise *would have performed* if conducted at the time. With the Cook rankings, I use the performance of the 2002 classifications to calculate s_t^c and f_t^c for the 2004 forecast, classifications from 2002 and 2004 for the 2006 forecast, and classifications from 2002, 2004, and 2006 for the 2008 forecast. With the Rothenberg rankings, I use the performance of the 2006 classifications to calibrate the 2008 forecast.

Results of the exercise appear in Figure 1. Each of the panels displays a histogram of 10,000 simulated elections using the algorithm described above. The red vertical line displays the actual election outcome. The figures suggest that the expert-assessment approach would have performed exceptionally well in 2004 and 2008, and less well (though the outcome still within the margin of error) in 2006. The 2004 simulation projects a Republican caucus of 233 seats, while the actual number of seats captured was 232. The 2008 Cook simulation projects 177 seats and the Rothenberg simulation anticipates 175 – the actual number was 178. The Democrats substantially outperformed expert assessments in 2006 to retake the House: whereas the simulation would have predicted, with near certainty, the Democratic takeover, it also would have anticipated a narrower margin between the parties. The simulation projected a Republican seatshare of 209 – seven seats off of the actual figure of 202. The eventual Republican seatshare does, however, just fall within the 95% confidence interval of the prediction.

4 Historical Results Using Contemporary Data

While the expert-based simulation approach does well using data from immediately prior to the election, an observer handicapping House elections may not wish to wait until the final rankings come in. Fortunately, both the Cook and Rothenberg Political Reports issue periodic rankings over the course of the election season.

Figure 1: Simulated 2004, 2006, and 2008 House Election Outcomes using Cook and Rothenberg Political Report Race Classifications Immediately Proceeding the Election
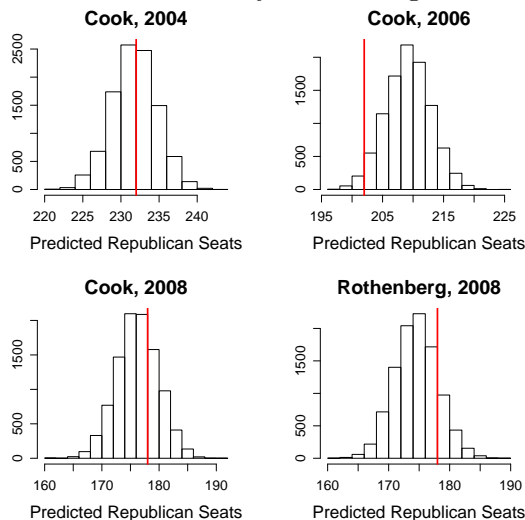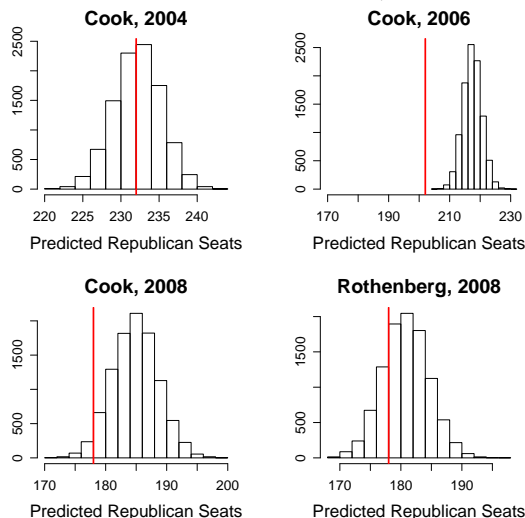

The Histograms display the predicted number of Republican seats from 10,000 simulated elections using the methodology described in the text. The red vertical lines display the actual number of Republican seats.

As a benchmarking exercise, I ascertain the accuracy of this approach in predicting past elections using the data that would have been available by late-September/early-October of the given election year. For the Cook simulation, I rely on rankings released August 7, 2002; September 7, 2004; September 20, 2006; and September 25, 2008. For the Rothenberg simulation, I use the rankings from October 6, 2006 and September 30, 2008. Needless to say, these rankings are typically less accurate in predicting race outcomes than the rankings posted later in the election cycle, as they cannot take advantage of the information generated in the two months leading up to the election – this will tend to be reflected in the poorer predictive accuracy of these forecast relative to those using the immediate data above.

Figure 2 displays simulations for the 2004-2008 House races. The figures indicate that the expert-based approach would have been highly reliable in predicting the race in September of 2004. In other words, there was little additional fluidity in the national outlook in September and October of that year that changed the outlook of the race. The predictions for 2008 based on early September projections overestimates the number of Republican seats: Cook predicts 185

Figure 2: Simulated 2004, 2006, and 2008 House Election Outcomes using Cook and Rothenberg Political Report Race Classifications as of Late September/Early October Proceeding the Election

The histograms display the predicted number of Republican seats from 10,000 simulated elections using the methodology described in the text. The red vertical lines display the actual number of Republican seats.

and Rothenberg 181 (vs. an actual outcome of 178). However, the outcome for that year falls well within each simulation's margin of error. 2006 is the clear outlier here – as of September, the expert ratings would have implied an almost perfectly split House, with 218 Republicans and 217 Democrats, with a 45% probability of a Democratic takeover. Significant developments occurring in the last two months of the race substantially improved the Democrats' outlook (as reflected by the final rankings that year) by the time of the election.

The failure of the simulation using the September 2006 rankings suggests that a party's electoral fortunes can change dramatically in the last few months of an election season. It is also noteworthy that 2006 was a midterm election year. While too much should not be made of this fact, given that the benchmarking is only done with three election cycles, it is consistent with the notion that midterm elections are often more difficult to predict than presidential-year congressional elections.

5   The 2010 Forecast

To assess the probability of a Republican House, we can use the most recent (as of writing) race classifications, which were released by Cook on September 27 and Rothenberg on September 24.

Cook currently lists 146 solid Democrat, 26 likely Democrat, 32 leaning Democrats, 43 Democratic toss-ups, 3 Republican toss-ups, 8 leaning Republican, 15 likely Republican, and 162 solid Republican districts. Rothenberg currently lists 172 solid Democrat, 17 favors Democrat, 19 leans Democrat, 8 toss-up Democrat, 15 pure toss-up, 14 Republican toss-up, 10 leaning Republican, 14 favors Republican, and 166 solid Republican districts.

The top two panels of Figure 4 displays the forecast for the current election cycle using the methodology described above, employing the Cook and Rothenberg forecasts, respectively. Using the Cook data, the approach anticipates a Republican caucus of around 209 members, plus or minus eight seats, or a loss of 31 seats for the Democrats. The Rothenberg data yields even better news for the Democrats: an expected Republican seat share of 200 seats – amounting to only a 22 seat loss. The methodology anticipates only a 3 percent probability of a Republican takeover using the Cook data, and a negligible probability using the Rothenberg. Compare this to the 79% probability estimated by Bafumi et. al. This prediction, of course, comes with the caveat – underscored by the 2006 simulation – that a lot can happen between now and election day.

A more significant source of bias in this approach stems from the fact that the parameters for solid, likely, leans, and toss-up Democratic and Republican districts are assumed to be drawn from different distributions – that is, the weights are "asymmetric." This approach may lead to a pro-Democratic bias because Democrats did so well in both 2006 and 2008. In other words, the approach's use of data from those elections may assign too high a likelihood to pro-Democratic upsets, and, correspondingly, too low a weight to pro-Republican upsets.

To address this concern, I reran the analysis assuming the weights are "symmetric:" Thus, if the Republican in a "leans Republican" district can be expected to win around 80% of the time, then the Republican would have around a 20% chance of victory in a "leans Democratic" district. The results of the simulations with symmetric weights appears in the bottom two panels of Figure ??. The symmetric weight simulation using the Cook data anticipates a Republican caucus of 215 members, plus or minus 18 seats. This corresponds to a roughly 37% probability of a Republican takeover. Using the Rothenberg data, the symmetric weight simulation predicts a Republican caucus of 213 members, plus or minus nine seats, or just a 14% probability of a

Figure 3: Simulated 2010 House Election Outcome using Current Cook and Rothenberg Political Report Projections
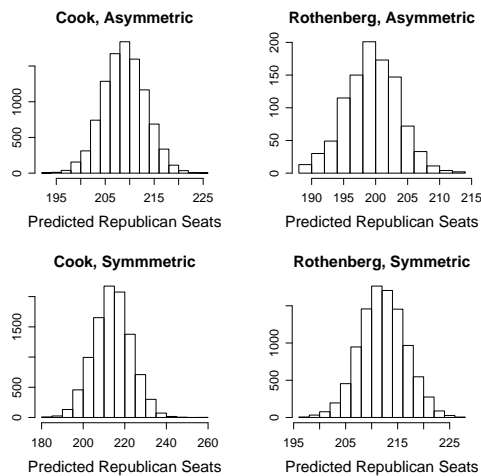
The top histograms display the predicted number of Republican seats from 10,000 simulated elections using the asymmetric weighting methodology described in the text. The bottom histograms repeats the exercise, but use symmetric weights to adjust for a possible pro-Democratic bias.

Republican takeover.

6 Is a "Wave" Responsible for the Methodology Underestimating Likely Republican Gains?

The above indicates that reports of the Democrats' imminent loss of the House may be exaggerated. It is possible, however, that my approach does not fully take into account the possibility that the 2010 election is fundamentally different from previous congressional elections in the last election, owing to a pro-Republican "wave." Cook, for example, argues that in wave years, toss-ups go disproportionately to the party benefiting from the wave, and substantial losses may even occur in the "lean" category for the party on the losing end of the wave.[4]

This raises an interesting question: do the current seat rankings reflect the raters' beliefs about the likelihood that this *is* a wave election year? The popular press tends to label game-changing elections such as 1994 and 2006 because the out-party did *surprisingly* well relative to expectations earlier in the season. Suppose, however, that the raters foresee a Republican wave that would lead to gains even in the "leans Democratic" category. Because they presumably wish to call elections correctly, they would then have an incentive to shift seats in that category to the toss-up column; likewise, they might likewise have an incentive, on the margin, to shift toss-ups into the leans Republican category. If this were the case, then we would be unlikely to see as many upsets as when the wave comes as a surprise to analysts.

On the other hand, it is also possible that raters are conservative at this stage of the game, and are reluctant to base their forecasts on their beliefs about the likely trajectory of the race. If that is the case, then the ratings do not reflect fully rational expectations; however, Cook's assessment about upsets is more likely to be correct.

To assess Cook's hypothesis, we can see what the model would forecast if we expected Republicans to fare as surprisingly well in 2010 as the Democrats did in 2006 and 2008. Accordingly, I reran my analysis on the Cook Political data using contemporary forecasts from those selection cycles using "reflected" weights: the simulation flips the parameters around in such a way as to assume that the 2010 Republicans will enjoy the 2006 and 2008 Democrats' good luck, and the

---

[4]E-mail communication, September 26, 2010.

Figure 4: Simulated 2010 House Election Outcome using Reflected Cook Projections
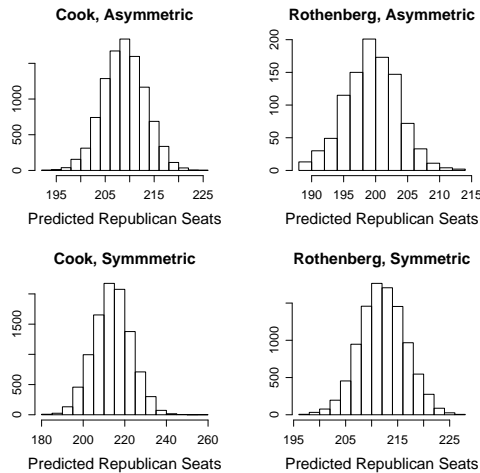
The histogram displays the predicted number of Republican seats from 10,000 simulated elections using the reflected weighting methodology described in the text.

2010 Democrats will suffer from the 2006 and 2008 Republicans' bad luck. In other words, while the "symmetric weighting" methodology described in the previous section attempts to erase the potential pro-Democratic bias, the "reflected weighting" seeks to create a pro-Republican bias. The results of the reflected weighting simulation appear in Figure ??.

As is evident from the figure, the reflected weighting approach yields a predicted 235 seat majority for the Republicans, plus or minus 10 seats. The corresponding expected seat gain of 57 seats is the largest predicted Republican gain that I am aware of, and corresponds to a 100% probability of a Republican takeover.

7 The Effect of Rankings Changes

As a final matter, it is useful to consider how changes in the rankings between now and election day could affect the outcome. While numerous transitions between categories are possible, the most dramatic change in the probability a seat changes hands occurs when a seat classified as "leans" is moved to the "toss-up" category. Thus, a useful question to ask is how many Democratic-leaning districts would have to shift to the toss-up category for the Republicans to even the odds of a takeover.

Figure 5 depicts the results of this exercise, employing the eight-category and folded four-category methods. Currently, there are 32 Democratic-leaning races. Using either the eight- or four-category approach, 10 or more would have to switch to the toss-up category in order for the Republicans to have a 50/50 chance or greater of reclaiming the House.

8 Conclusion

In this note, I have described a very simple approach to exploiting expert ratings of individual House races to generate a probabilistic forecast of partisan control of the chamber. This approach is simple and easy to implement. Using current and historical seat rankings by the Cook and Rothenberg Political Reports as of September 28, I project that the Democrats will retain control of the House of Representatives, albeit with a considerably smaller margin than that enjoyed in the current Congress.

An interesting disparity exists between the results described below and the interpretation of the rankings by the experts themselves. On its website, The Cook Political Report projects a Republican seat gain of 40 seats, which would be just sufficient for a Republican majority. If we assume, however, that the rankings reflect the polls, what accounts for the disparity between their forecast and the one contained here? One possibility is that Cook views the odds associated with its own ratings categories are different from those warranted by their historical performance. A more likely scenario is that their aggregate forecast includes expectations about a continued Republican surge that is not adequately reflected in the race-specific rankings. An open question, therefore, is the extent to which expert-based seat rankings capture such national tides. My simulation is

Figure 5: The Probability of a Republican Takeover as a Function of Hypothetical Switches of Races from the "Leans" to "Toss-up" Democratic

Cook, Asymmetric Weights

Cook, Cook Symmetric Weights

13
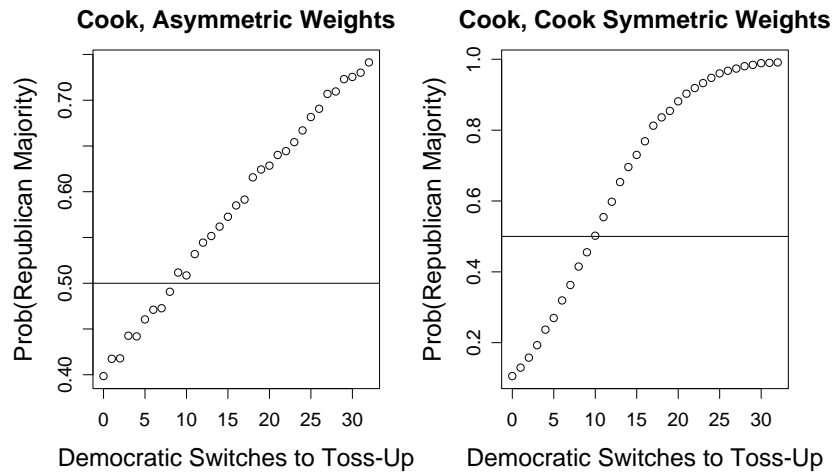
premised on the notion that raters exploit all available information when ranking districts, which would include factors affecting all districts simultaneously. To the extent that they do not, then the simulation's assumptions are false, and it is likely to fail given significant national momentum for one party or the other.

Two avenues for improving the predictive accuracy of this approach are immediately apparent. The first would be to use expert assessments from other sources, either in isolation or in conjunction with those from the Cook Political Report. The second would be to incorporate observable district-level metrics into the simulation in a hybrid regression-and-expert-based approach. At the very least, however, this note shows that there may be value to taking seriously what qualitative observers of U.S. politics pay attention to "on the ground."

References

Bafumi, Joseph, Robert S. Erikson, and Christopher Wlezien. 2008. "Forecasting House Seats from Generic Congressional Polls." In *Elections and Exit Polling*, eds. Wendy Alvey and Fritz Scheuren. New York: Wiley and Sons, 118-123.

Bafumi, Joseph, Robert S. Erikson, and Christopher Wlezien. 2010. "Forecasting the House of Representatives' Seat Division in the 2010 Midterm Election." *PS: Political Science and Politics*, forthcoming.

Campbell, James E. 2010. "The Seats in Trouble Forecast of the 2010 Elections to the U.S. House." *PS: Political Science and Politics*, forthcoming.

Cook Political Report. Available online at http://www.cookpolitical.com.

Rothenberg Political Report. Available online at http://rothenbergpoliticalreport.com.

Silver, Nate. 2010. "G.O.P. Has 2-in-3 Chance of Taking House, Model Forecasts." *FiveThirtyEight: Nate Silver's Political Calculus*, September 10, 2010, http://fivethirtyeight.blogs.nytimes.com.