

Piecing it all together and forecasting who governs: the 2015 British General Election

Stephen D. Fisher
Department of Sociology, University of Oxford
Trinity College, Oxford OX1 3BH.
stephen.fisher@trinity.ox.ac.uk

Abstract

This paper discusses a new probabilistic forecasting method that was designed for the 2015 British general election. It proceeds in a series of steps from opinion poll averaging, forecasting national-level vote shares and uncertainty estimates, and subsequent simulation of hypothetical election results, through modelling of constituency polls and survey data to identify and adjust for patterns in the constituency-level variation in party performance, and finally to probabilistic forecasting of seat outcomes and of different combinations of parties commanding relevant governing majorities in parliament.

Key words:

opinion polls; election forecasting; UK elections; British politics; government formation

Highlights:

- A new probabilistic method of forecasting British general elections from opinion polls is presented.
- Method discusses how both national, constituency and individual-level data are used.
- The model is exploited to produce forecast probabilities for relevant combinations of parties holding governing majorities.

Introduction

This paper describes the method for the final of a 20 month series of forecasts that were published at ElectionsEtc.com.

The methodology for this forecast is based on a series of steps that pieces together information from various sources and combines different estimates. The paper is structured as a step-by-step sequence, with the forecasts for various aspects presented in the appropriate locations.

1. Calculate averages of recent Britain-wide and Scotland-only vote intention polls

For GB polls I use an average of various different methods of averaging. The idea is to look for consistency and robustness across different methods. This includes checking how things change after excluding outliers, excluding particular pollsters one-by-one, weighting for past performance or not, and varying how far back and how many polls per pollster were used. The aim was to get a polling

average that treats the pollsters as relatively but not completely equal and averages over enough polls that sampling variation can be assumed to cancel out. For much of the time this also has the effect of smoothing over small short-lived blips.

The polling average at midnight before the election was Conservative 34%, Labour 33%, Liberal Democrat 9% and the UK Independence Party (UKIP) 13%.

Scottish polls are fewer and further between so I use the whatscotlandthinks.org method of taking the average of the last four polls. This polling average at midnight before the election was Conservative 15%, Labour 25%, Liberal Democrat 5%, Scottish National Party 50%, and UKIP 2%.

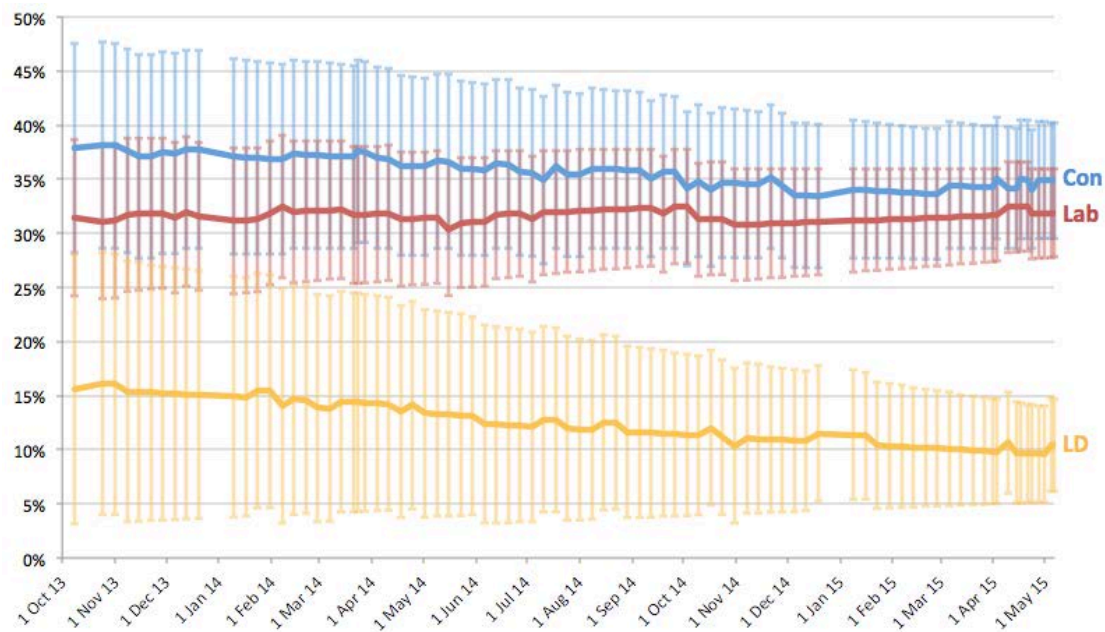
The choice of these more informal methods stands in marked contrast to the use of house-effects or state-space models (Jackman 2005, Fisher et al. 2011). These models are clearly appropriate and have many advantages for the analysis of the dynamics of public opinion. But they are risky in the contest of forecasting. They essentially assume that house effects are relatively or totally stable without announced changes in methodology. But there have been instances of dramatic changes in house effects shortly before elections, such as during the Scottish independence referendum campaign. Many pollsters use slightly different adjustments for their final pre-election polls, and concerns have been raised that the choice of these special adjustments might be influenced by polls published by other pollsters (herding). Under these circumstances it seems safest to assume that house effects are not stable and apply an average of the most recent or most recent two polls, perhaps with reputation weighting. All these lead to similar estimates.

2. Use regression analysis of historical votes and polls to forecast how GB vote intention will change from a given number of days before the election, and to estimate prediction intervals for those changes.

For 20 months before the election I applied the method described in Fisher (2015a) for projecting current public opinion to a forecast election share of the vote at the GB level with prediction standard errors. For the Conservatives and Labour the forecast is based on an average of two elaborations of vote-on-polls type regression models (Erikson and Wlezien 2002). Since a votes-on-polls model with the addition of a government status dummy and a weighted relative change since the last election model both had similarly good out-of-sample prediction properties, an average of the two is used.

The figure below shows how the forecast progressed over time with 95% prediction interval bars. The Labour forecast stayed steady as they declined in the polls at the rate history would suggest, but the Conservatives and Liberal Democrats did not recover as expected so their forecasts fell.

Figure 1 Forecast GB vote shares by date of forecast, with 95% prediction intervals



Source: Author's own calculations published at Fisher and Jones (2015)

The final forecast GB shares of the vote with 95% prediction intervals based on pooled standard errors are Con 35% (31% – 39%), Lab: 32% (28% – 36%), LD: 10% (7% – 14%), UKIP: 12% (8% – 16%) Others: 11% (9% – 12%). Forecast shares for the two main parties in Scotland with similarly constructed 95% prediction intervals are SNP 48% (44% – 52%) and Labour 27% (23% – 31%).

For Labour and the Conservatives at the national level, the main reason the forecast shares differ from the current polling average is because the polls have tended to overestimate Labour and underestimate the Conservatives, both by about a point and a half (Fisher 2015). The model also predicts swing back towards 2010 levels for the Liberal Democrats, UKIP, SNP and Scottish Labour.

3. Use the forecast vote shares and uncertainty estimates, and between party correlations in the opinion polls, to simulate hypothetical election results.

For this I use a multivariate normal distribution with variances for each party estimated by pooling the forecast standard errors from the previous step. Since parties do not go up or down independently I use the average correlations between changes in party shares in successive polls as estimates of covariance for the simulations. So in a hypothetical election where UKIP does particularly well it is more likely that the Conservatives especially will do badly. There was also a big negative correlation between Conservative and Labour performance, which widens the range of possible outcomes in the simulations.

4. Use Ashcroft constituency polls and individual-level data kindly provided by YouGov to identify constituencies where parties are doing

particularly well or particularly badly, and apply adjustments to the hypothetical results accordingly.

The analysis of the Ashcroft and YouGov data is done at the constituency level using simple regressions of party changes in the share of the vote since 2010 with binary predictor variables. This mimics the approach taken for the exit poll prediction (Curtice et al. 2011).

The most important constituency level factors within England and Wales are to do with incumbency.

Those Conservative MPs who took their seat from an MP from another party in 2010 appear to be doing a couple of points better than other Tory candidates. This seems to be an instance of the classic sophomore surge, which is common in the US and also seemed to help many first-term Labour MPs hold on in 2001 despite a swing to the Conservatives.

Also, incumbency effects for Liberal Democrat MPs seem to be strengthening, by about 7 points above and beyond the personal vote bonus they got in 2010. But this is against a backdrop whereby the party is falling more where it started stronger, not least because there are many seats where they are starting with fewer votes than uniform swing suggests they should lose. But these effects depend on the principle challenger. Liberal Democrat MPs appear to be doing worse where Labour came second in 2010 but better where the Tories were second. Correspondingly the Labour share is up more where they are challenging a Lib Dem MP and the Tories seems to be falling a bit more where they are starting second to the Liberal Democrats.

Labour are the chief beneficiaries of the decline in Lib Dem support but this means they seem to be advancing more in places where the Lib Dems did well last time than they are in their target seats. Moreover in Scotland Labour are apparently falling further the higher they started. Again, this is partly because there are places where they did not start high enough to fall in line with the national average. This means the SNP are forecast to do better and Labour worse than would be expected from uniform swing with the Scottish forecast share of the vote.

All these effects are applied systematically to the hypothetical elections by making the vote share projection equations dependent on the specified constituency characteristics.

5. Convert forecast seat shares into probabilities for different parties winning in each constituency in each hypothetical election.

For each simulated election I apply the formula for estimating win probabilities from seat-by-seat shares developed for the exit poll and set out in Curtice and Firth (2008). With so many hypothetical elections, this step does not make much difference. But it does bring down the SNP forecast from 55 to 53 out of 59 Scottish seats. There are more opportunities for the SNP to under than over

perform a such a high forecast, gains from above average swings are likely to be more than offset by the consequences of below average swings.

6. Set subjective probabilities for a very small number of particularly unusual seats where the constituency variation models are problematic.

There are not many of these, but they are especially important for UKIP. Some of the Ashcroft polls in UKIP target seats are relatively old and the scale and pattern of change for UKIP is so dramatic that it is hard to predict what will happen in a few relatively unusual seats where UKIP are running intensive campaigns. So the estimate of the total number of UKIP seats is effectively a guess based on various polling and non-polling sources of information. The sum of these probabilities suggests 3 UKIP seats. There is also a subjective probability set for George Galloway to be more likely than not to win Bradford West.

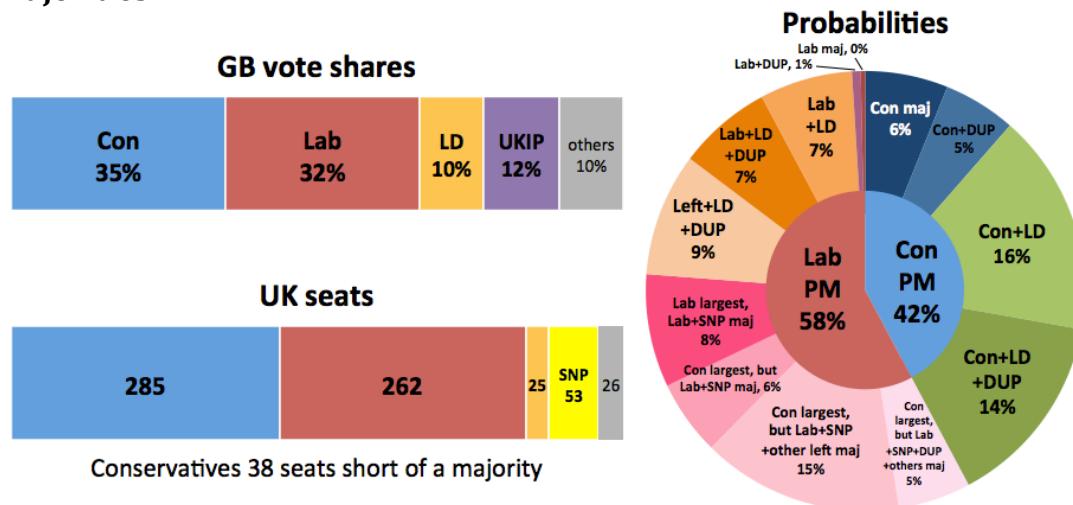
On this basis the central GB seat forecast, taking the mean across simulations, with 95% prediction intervals from the percentiles of the simulated seat totals, is Con 285 (245 – 326), Lab 262 (223 – 300), LD 25 (17 – 33), SNP: 53 (45 – 57), PC: 3 (2 – 3), UKIP: 3 (3 – 4), Grn: 1, Other: 1. These figures do not sum to 632 due to rounding of sums of probabilities.

7. Analyse the distribution of seats forecasts across the hypothetical elections to calculate probabilities for key events (e.g. hung parliament, Con largest party etc.)

For any single simulated hypothetical election the forecast number of seats for a party is just the sum of the probabilities that the party will win each seat. The distribution of these forecast seat totals across the 10,000 simulated hypothetical elections tells us the relative chances of different events. In particular this approach indicates which combinations of parties will have a majority (more than 323 seats given that Sinn Fein do not take theirs).

The pie chart in the figure below shows how complex this gets even when we use strong simplifying assumptions as to which grouping is most likely to control the government when several different combinations are possible. Whether it would be a coalition, confidence-and-supply or some other kind of agreement is not addressed here; that is less important for policy than which parties are involved.

Figure 2 Central forecast and pie chart of probabilities for governing majorities



Source: Author's own calculations published at Fisher and Jones (2015)

Starting from the top of the pie chart and working clockwise: there is a 6% chance of a Conservative majority. If they are just short of a majority they could call on the DUP, or if that is not enough then the Liberal Democrats or perhaps both if needs be.

Conversely, starting from the top of the pie chart and working anti-clockwise shows the relative chances of different Labour led governments. As with the Tories, Labour would also call on the DUP and Lib Dems first if those parties were sufficient to yield a majority. This assumes that the SNP will be harder for Labour to do work with, especially given their differences over Trident. But if it were necessary to appeal to the SNP and SNP support was sufficient to give Labour a majority then given the SNP pronouncements on the issue it is assumed that support would be given by the SNP to a Labour led government.

Since experience from other countries suggests that a majority-commanding group of parties is more likely to control the government the fewer parties there are in the group as well as greater ideologically similarity between them, the pie chart suggests that Labour might rely on the SNP alone before turning to other parties. But all the pink sections are Labour led governments depending on SNP support in one form or another. Although some of the them would also allow for Liberal Democrat cooperation on the left, Nick Clegg ruled out arrangements with a Labour led government that depended on the SNP for "life support", which is the case in all the pink sections in the pie chart.

Overall there is a 94% chance of a hung parliament, a 77% chance of the Conservatives having the most votes and a 68% chance of them having the most seats, but because Labour has more potential governing parties there is simultaneously a 58% chance of a Labour led government.

Post-election postscript

My final forecast was far from the actual result primarily because the vote-intention polls failed to reflect the 6.5 percentage point Conservative lead over Labour in the eventual GB vote share. The choice to use a bespoke average of not just the final polls helped slightly since there was no Tory lead in the final polls. On this issue there is an error in the above text: while it was usually the case that my polling average was practically insensitive to the choice of range time over which to pool polls, final polls only would have suggested a smaller Conservative lead.

Although there were no dramatic changes in house effects by party during the campaign, there was a noticeable contraction in the variation in Con-Lab leads between the penultimate and final polls across the 11 pollsters that published a final poll. Whereas the standard deviation across final polls was 0.94, the corresponding figure was 2.06 for penultimate polls. This was primarily the result of movement towards zero by those with noticeable Conservative leads in their penultimate polls, which worsened the polling average and somewhat legitimates the concerns raised above about the stability of house effects. But this issue is small compared with the overall polling error on the lead.

The vote share forecasting model was right to suggest that the polls were likely to underestimate the Conservatives and overestimate Labour, but since the actual error was much greater than the historical average the Conservative lead over Labour was underestimated by 3 percentage points. This inaccuracy was the biggest single source of error in the seats forecast. Had I estimated the Tories 1.5 points higher and Labour 1.5 points lower, the forecast would have been Conservative 303, Labour 244.

The constituency level model within Scotland was crude but improved the forecast, compensating for the underestimate of the Labour to SNP swing. While the win probabilities formula worsened the forecast, the principle was still right: deterministic forecasting with the right shares and constituency pattern overestimates the SNP. Overall the forecast of 53 SNP seats was close to the 56 they achieved.

Constituency level factors in England and Wales identified from modelling the Ashcroft and YouGov data all proved to be statistically significant with coefficients reasonably close to those for equivalent models for actual results, with one big exception: the Conservatives did not do much better where they started in third place. This inaccuracy together with 1 point underestimates of both the Conservative first-term incumbency effect and the impact of LD incumbency on the Tories jointly had the effect of bringing the forecast Conservative seat total down by 9 from what it would have been with accurate estimates of these effects.

While the constituency level model improved the forecast only slightly overall, and in truth worsened the LD forecast, the YouGov and Ashcroft data did reflect

the eventual constituency pattern fairly well, especially relative to the much more serious overall polling error.

The subjectively assigned probabilities did more harm than good, but only increased the error by 3 seats.

While conducting this post-mortem I discovered that what were presented above and on ElectionsEtc.com as 95%- were in fact 90%-prediction intervals. This is true of the text above regarding both seats and shares intervals, but Figure 1 is accurate. This is an embarrassing mistake, especially since there were some earlier forecasts where the prediction interval did not include 323 for Labour and so implied less than 2.5% chance of a majority while the corresponding pie charts showed Labour with a 4% or greater chance (e.g. Fisher and Jones 2015b). Also the final seats prediction interval for the Conservatives must be a 90% one to be reconciled with the 6% chance of 323 seats or more shown in the pie chart.

The 90% prediction intervals for vote shares included the actual result for all those parties forecast, including shares of the Scottish vote for Labour and the SNP. The prediction intervals for seats covered the final outcome for Labour, PC and the SNP but not the Conservatives, Liberal Democrats or UKIP. The latter was just based on subjectively assigned probabilities, so a prediction interval should never have been published. The failure to include the actual Conservative seat total was partly the result of the underestimating both their lead in votes over Labour and over the Liberal Democrats, but also the product of their more efficient votes-to-seats relationship (largely as a result of unusually strong first-term incumbent "sophomore surges").

Pooling the forecast standard errors for the simulations led to slightly narrower prediction intervals for the Conservatives than should have been provided, but correcting this only raises the probability of the Tories getting the 331 or more seats to just 4%, up from the 3.3% implied by the final forecast simulations.

Sadly of all 85 of my forecasts over 20 months the closest to predicting the eventual Conservative majority was the first. In October 2013 I forecast Conservatives to get 337 seats (Fisher 2013). However using the February 2014 revised methodology the Conservative forecast from the same point was just 315 (Fisher 2015). It is some comfort that my forecast always had the Tories ahead on votes. On seats the Tories were predicted either clearly ahead or neck-and-neck, but consistently in the lead since the start of March 2015. Although, at the end, Betfair had a slightly greater implied probability of a Tory majority (>325 seats) than my final forecast (6% to 5.4%), up to that point my model had the same or better (and before 2015 much better) probability of a Tory majority than the betting markets (Tapper 2015).

Acknowledgements:

I am very grateful for comments from and discussions with various people on various aspects of the forecasting methodology including (in alphabetical order): Richard Coggins, Andrew Cooper, Philip Cowley, David Cowling, John Curtice,

Robert Ford, Matthew Goodwin, Chris Hanretty, Will Jennings, Ron Jonston, Jouni Kuha, Ben Lauderdale, Matt Lebo, Eilidh Macfarlane, Iain McLean, Jon Mellon, Andreas Murr, Clive Payne, Mark Pickup, Chris Prosser, Petra Schleiter, Mike Smithson, Jon Tonge, Nick Vivyan, Anthony Wells, Chris Wlezien and Chris Zorn. I have also much appreciated numerous and wide ranging comments on Twitter and elsewhere online. I am grateful to Joe Twyman for access to individual-level YouGov data, Ben Lauderdale, Tom Lubbock and Eilidh Macfarlane for spreadsheets with Ashcroft constituency poll data, Will Jennings for the historical polling data. Special thanks to Jonathan Jones who maintained the ElectionsEtc.com website, helped with various aspects of the forecasting project, commented on earlier versions of this text and compiled the graphics for Figures 1 and 2.

Disclosure:

Stephen Fisher is not a member of any political party and does not receive money from any person or organisation that might create a conflict of interest.

References

Curtice, John, and David Firth. 2008. "Exit polling in a cold climate: the BBC-ITV experience in Britain in 2005." *Journal of the Royal Statistical Society. Series A: Statistics in Society* 171(3): 509–539.

Curtice, John, Stephen D Fisher, and Jouni Kuha. 2011. "Confounding the Commentators: How the 2010 Exit Poll Got it (More or Less) Right." *Journal of Elections, Public Opinion & Parties* 21(2): 211–235.

Wlezien, Christopher, and Robert S Erikson. 2002. "The timeline of presidential election campaigns." *The Journal of Politics* 64(04): 969–993.

Fisher, Stephen D. (2013). *A Long-Range Forecast for a 2015 British General Election Based on Current Polls and Historical Polls and Votes*.
<http://electionsetc.blogspot.co.uk/2013/10/a-long-range-forecast-for-2015-british.html> . Accessed 25 October 2013.

Fisher, Stephen D. (2015). "Predictable and Unpredictable Changes in Party Support: A Method for Long-Range Daily Election Forecasting from Opinion Polls." *Journal of Elections, Public Opinion & Parties* 25(2): 137–158.

Fisher, Stephen D. and Jonathan Jones (2015) *Election-Day Forecast*.
<http://electionsetc.com/2015/05/07/election-day-forecast/>. Accessed 7 May 2015.

Fisher, Stephen and Jonathan Jones (2015b) *Forecast update: 20 March 2015*.
<http://electionsetc.com/2015/03/20/forecast-update-20-march-2015/>. Accessed 1 June 2015.

Fisher, Stephen D, Robert Ford, Will Jennings, Mark Pickup and Christopher Wlezien. 2011. "From polls to votes to seats: Forecasting the 2010 British general election." *Electoral Studies* 30(2): 250–257.

Jackman, Simon. 2005. "Pooling the polls over an election campaign." *Australian Journal of Political Science* 40(4): 499–517.

Tapper, Albert. 2015. "Predicting GE 2015: which group of forecasters should we rely upon: 'the expert crowd', 'the betting crowd' or the 'mass public crowd'?" : 1–17. http://alberttapper.blogspot.co.uk/2015/05/predicting-ge-2015-which-group-should_7.html (Accessed July 3, 2015).