

Different methods, similar outcome: comparing the Poll of Polls with MRP



The star of the show in 2017's general election polling landscape was YouGov's MRP model, which produced remarkably accurate estimates of the results in seats across the country. The equivalent model for this year's election, explains [Joe Greenwood](#), produces quite similar estimates to approaches based on the average figures from standard national polls.

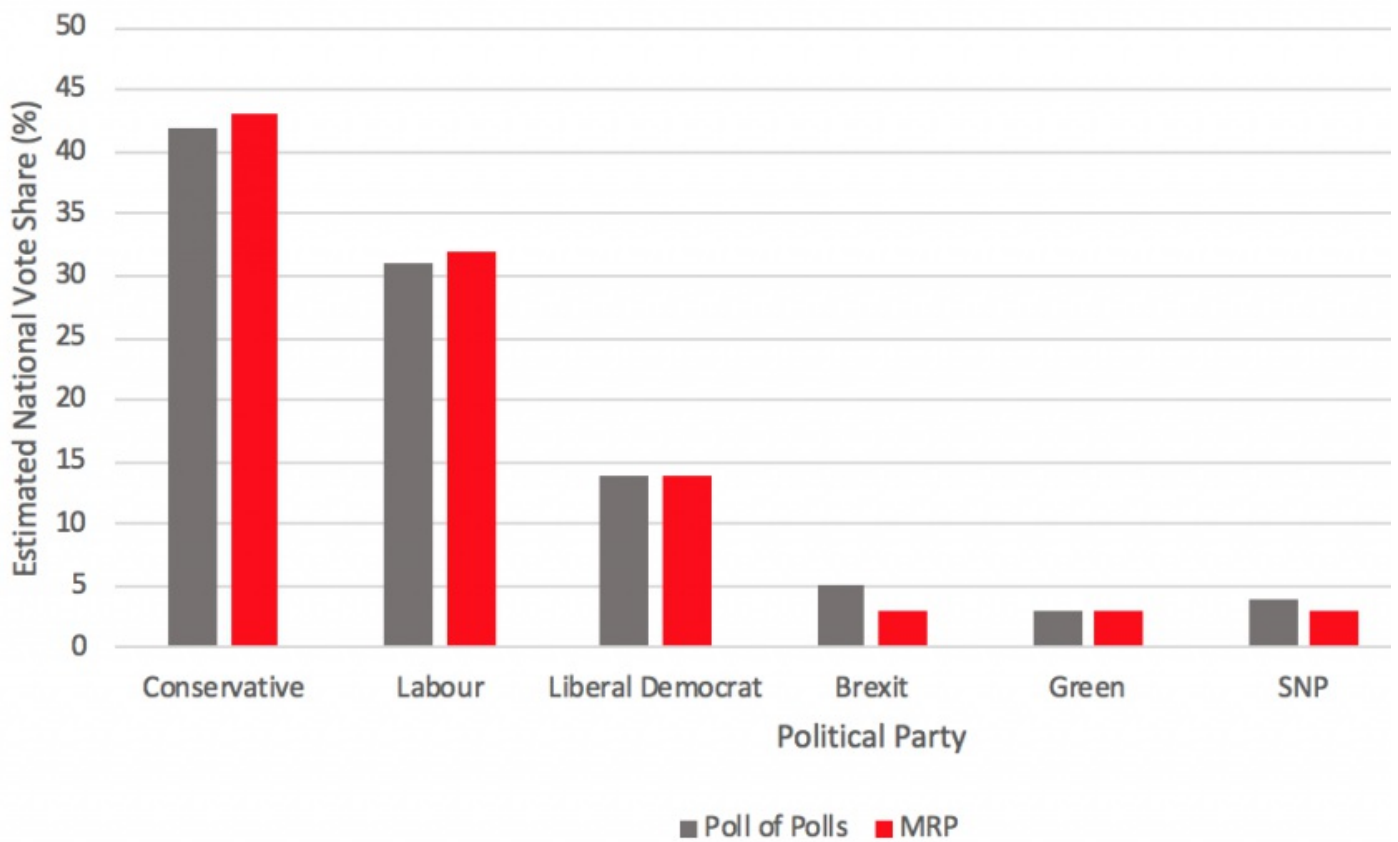
At an excellent event on '[Reading the 2019 election polls](#)' at LSE last week, the sense of anticipation regarding the shortly-to-be-released YouGov Multilevel Regression and Post-stratification (MRP) model was palpable. Indeed, the two academics who run the model using YouGov's data, [Ben Lauderdale](#) and [Jack Blumenau](#), made it a running theme of their presentation that they could not even hint at the model's estimates of party vote shares or numbers of seats. Why the anticipation? Well, back in 2017 when many polling companies had a bad time estimating the vote shares of parties (especially Labour) in that year's election, the previous incarnation of the [MRP model was remarkably accurate](#) at estimating, in particular, the seats that each party would end up with. Indeed, it predicted 93% of the constituency results correctly and anticipated shock outcomes such as those in [Canterbury](#) and [Kensington](#).

Later that night, [the model's national and seat-by-seat estimates were released](#) and indicated that, based on current voting intentions, the Conservatives would win a comfortable victory. Indeed, even at the lowest end of the estimates, Boris Johnson's party was estimated to reach 328 seats, giving him a small majority (or a slightly larger one, once we take account of Sinn Fein MPs who do not take up their seats, and the seats of the Speaker and Deputy Speaker of the House of Commons).

How did the model produce this estimate? Well, [MRP](#) begins by estimating the relationship between voting intention and key characteristics such as age, gender, and education. It then applies these relationships in different constituencies based on characteristics of their populations (obtained from official statistics such as the [Census](#) and subsequent [updates](#)), whilst also adjusting for the characteristics of the seats themselves (such as whether they voted Leave or Remain, and even how many fish and chip shops they have). It is this process that gives the model its particular benefit: estimates of the outcome in each parliamentary constituency.


At a national level, though, how different is the estimate of vote share produced by the MRP model from the average of the estimates produced by the standard polls that we have seen up until now? In short: not very different. I start by comparing the national level vote share produced by the [BBC's poll of polls](#) with the equivalent figure produced by the MRP. As can be seen in the first figure, the estimates are remarkably similar. Indeed, none of the party vote share estimates given by the MRP differ from the equivalent figures in the poll of polls by more than 2%. The MRP gives the Conservative and Labour parties each a 1% higher vote share than the poll of polls, whilst it estimates a 2% lower share for the Brexit Party and a 1% lower share for the SNP. The Liberal Democrat, Green, and Plaid Cymru vote shares are the same in both approaches.

Vote Share from the Poll of Polls and the MRP Model



But, this is not the most important part: MRP's main benefit lies in its seat estimates. Turning to the second figure, we can see a little more difference in terms of the number of seats the different approaches estimate each party will obtain, but still considerable similarity.

Seats based on Swing and the MRP Model



To calculate the seat estimates stemming from the poll of polls I took its national vote share figures and fed them, first, into [Election Polling's UK Swingometer](#) (labelled 'Uniform Swing' in the figure). Then, to make a basic adjustment for the differing voting patterns in Scotland from the rest of the UK, I fed those same national vote share figures into [Electoral Calculus' equivalent swingometer](#), which also allows separate vote share figures to be entered for Scotland (I took these from YouGov's latest national voting intention poll, which provides estimates for Scotland, and the columns a labelled as 'Separate Scottish Swing' in the figure).

The largest difference is that the MRP model's estimate of Conservative seats (359) is 13 higher than the estimate (346) stemming from the Scotland-adjusted swing. Otherwise, none of the seat estimates that the MRP model produces are more than eight seats different from either of the approaches using swing based on the average national vote share from the poll of polls. In line with its higher estimate of Conservative seats, the MRP model estimates fewer seats for Labour (seven below the Scotland-adjusted swing) and the Liberal Democrats (six below the Scotland-adjusted swing). In short, the MRP model, when compared with swing-based estimates drawing on average national polling vote share, suggests that the Conservative Party will do slightly better whilst Labour and the Liberal Democrats will do slightly worse.

Of course, neither of these estimates can account for something very important: what might change in the final days of the campaign. There could be major events and associated shifts in voting intentions in the population or certain sub-groups. This might lead the MRP model to predict a notably different outcome from an application of swing (uniform or otherwise) based on national voting intention figures. However, that seems unlikely, and if things stay roughly as they are now then, whatever estimate of seats we use, the Conservatives seem likely to have a comfortable majority.

About the Author

Date originally posted: 2019-12-02

Permalink: <https://blogs.lse.ac.uk/politicsandpolicy/poll-of-polls-mrp/>

Blog homepage: <https://blogs.lse.ac.uk/politicsandpolicy/>



Joe Greenwood [@niceonecombo](#) is an LSE Fellow in the LSE Department of Government, where he teaches on GV101 (Introduction to Political Science). He previously worked at YouGov and, before that, completed his PhD at the University of Essex. His research focuses on political participation, privilege, and perceptions in the British context.

All articles posted on this blog give the views of the author(s), and not the position of LSE British Politics and Policy, nor of the London School of Economics and Political Science.