British Politics and Policy

5/7/2015

At 10pm on election day, the ballots will close, and the counting will begin. It may be many hours before it becomes clear which party has the most seats in the new parliament, and many days until we know which parties will form the next government. However, shortly after the 10 o'clock deadline, broadcasters will release a forecast based on the exit poll, which will give the first insight into how voters have actually voted in the general election. In this post, Jouni Kuha explains the detail behind the exit poll forecast, and indicates some of the difficulties in putting together a forecast under considerable time pressures.

After the many months of pre-election polls and forecasts, the first post-election survey will be an exit poll done just outside the polling stations. There will be only one face-to-face exit poll in the UK on May 7th. It is a collective effort, commissioned jointly by BBC, ITV and Sky News, and the interviewing carried out by Gfk-NOP and Ipsos-Mori.

You will not hear any results from the exit poll during the day, as it is illegal to release them while the voting is ongoing. But as soon as the polling stations close at 10pm on Thursday, the broadcasters will announce a forecast based on the exit poll. This will be for numbers of seats (not national vote shares) for the main parties (Conservatives, Labour, Liberal Democrats, SNP, Plaid Cymru, UKIP, Greens, and all others combined). It will be almost the only hard evidence of the results for several hours, and will be extensively discussed in the election broadcasts.

In 2010 the exit poll forecast was received with scepticism, especially because it predicted a result for the Liberal Democrats which was quite different from pre-election expectations. In the end the forecast proved to be very accurate. The predicted seat total for the largest party – 307 for the Conservatives – was exactly correct, as it – 356 for Labour – had also been in 2005. Such precision requires a large amount of luck, and the forecast will not be as good again this time. Predicting the largest party's total to within 20 seats is a reasonable ambition for a well-conducted exit poll. We are hopeful that this can be achieved, as the methodology is well-founded and has been proved in the past.

So how is it done? The methodology of the exit poll has two crucial elements, the design of the poll itself and the statistical modelling of the data from it. Both are working toward one goal: predicting *variation in changes* in vote shares from 2010 to 2015.

The exit poll will interview nearly 20,000 voters. They are not evenly distributed across the country, but concentrated on a smaller number of carefully selected voting stations, around 140 of them in 2015. The design focuses on change by re-visiting where possible the same stations as in 2010, and on variation by selecting the stations from constituencies with different characteristics in dimensions – such as location, previous election results and demographic features – which are expected to affect the extent of the parties' gains and losses of vote share from 2010 to 2015.

The statistical analysis of the exit poll data consists of identifying regularities in the variation in the changes. For example, in 2010 we concluded that Labour was doing relatively better (i.e. losing less of their share from 2005) in constituencies where incumbent Labour MP was standing for re-election or which had a population with a large proportion of ethnic minorities, and relatively less well in England than in Scotland and Wales (we should really have combined Wales with England, but this was not easy to conclude from the data).

These models are then turned into the forecast in three steps: (1) the model for change is applied to all constituencies (except for ones in Northern Ireland and a few others, which are treated separately) given their

characteristics and their vote shares in 2010, to produce predicted vote shares in 2015; (2) the predicted shares are turned into predicted probabilities of winning the constituency, for every party in every constituency; (3) the probabilities are added up across constituencies to produce the forecast of the total numbers of seats for each party. You can find more information about the details of these steps in David Firth's excellent exit poll explainer, and the academic articles cited there.

On the day of the election, many of the challenges of the analysis have to do with time pressure. The predictors used in the forecast model are selected only then, when the 2015 data are actually seen. While the day is long, most of the voters go to the polls late in the day and so most of the data arrive toward the evening. At the end of the day, there is only a small interval between the times when the last pre-forecast data arrive and when the forecast must be sent to the broadcasters to be ready for publication immediately after 10pm. This makes for a thrilling but nerve-wracking experience for the modelling team.

Once the command *send.exitpoll.seats()* has been issued, the exit poll team will then join other forecasters and the rest of the nation to follow the results of the 2015 General Election with eager anticipation.

Note: This article gives the views of the author, and not the position of the General Election blog, nor of the London School of Economics. Please read our comments policy before posting.

About the Author

Jouni Kuha is Associate Professor of Statistics and Research Methodology and Head of the Department of Methodology at the LSE.

(Featured Image Credit: David Holt CC BY-SA 2.0)