31 March 2016

Ipsos MORI's analysis of our 2015 general election polling

Ipsos MORI reiterate today our gratitude to Professor Sturgis, and the rest of his British Polling Council (BPC) inquiry team, for their work on this inquiry.

Ahead of today's report we, along with many other polling companies, have been carrying out our own analysis of the reasons behind the error in our final general election poll and the specific issues that apply to our polls, some of which may be different to those of other companies.

Whilst much of this analysis is already in the public domain, we set out in this document a summary of our findings so far, which we hope will also be useful in understanding the steps we are taking to improve our polls for the future. Of course, we will continue to review our methodology and take into account the recommendations of the inquiry, published today, as far as we can.

1) Issues which affected our polling

Early on in our investigations, we identified the main error in our final poll to be an over-estimation of Labour voters, and an under-representation of non-voters. However, it is important to note that we did not under-represent Conservative voters as a whole within our total sample, as the table below shows.

	Election result		Ipsos MORI final poll
	millions	%	%
Total	50.5		
Conservative	11.3	22.4%	24.7%
Labour	9.3	18.5%	24.1%
Other parties	9.3	18.5%	20.7%
Will not/did not vote	20.5	40.6%	30.5%

Since the election our findings on our respondents' past voting behaviour also suggest our samples continue to include too many voters and not enough non-voters. However, much like the British Social Attitudes (BSA) and British Electoral Study (BES) reported figures, they do perform much better when reflecting which parties people say they voted for in the general election (as seen in the table overleaf, which shows our figures from June and July 2015[1], which used the same weighting and sampling procedures as before the election). Although, vote recall questions in post-election surveys after the outcome cannot be fully relied upon to accurately reflect the voting behaviour of the sample.

[1] We have seen the same pattern continue in our Political Monitor surveys since July too, with a clear Conservative lead over Labour in reported 2015 vote even in the unweighted samples – but to avoid confusion have only shown June-July here as those months used exactly the same weighting procedures as before the election

	IM Political Monitors				Error		
	IM June- July	BES May-Sept	BSA July-Nov	Actual	IM June-July	BES	BSA
Con	37.9	40.6	39.7	37.7	0.2	2.9	2
Lab	32.5	32.7	33.6	31.2	1.3	1.5	2.4
LibDem	9.4	7.1	7.3	8.1	1.3	-1	-0.8
SNP	4.4	4.5		4.9	-0.5	-0.4	
UKIP	8.7	10.7	9	12.9	-4.2	-2.2	-3.9
Green	5.8	3.2	4	3.8	2	-0.6	0.2
Plaid	0.7	0.5		0.6	0.1	-0.1	0.1
Other	0.6	0.7	6.4	0.8	-0.2	-0.1	0.1
				Average error	1.2	1.1	1.6

					Overclaim		
					IM June-July	BES	BSA
Implied turnout	77.2	73.3	70.3	66.4	10.8	6.9	3.9

Our turnout predictions also worked less well in 2015. In previous elections, a higher proportion of Conservative than Labour supporters voted. We have corrected for this in the past by asking respondents how likely they are to vote and adjusting accordingly. In 2005 and 2010, the proportions of each party's supporters saying that they were "absolutely certain to vote" proved a good indicator of relative turnout levels and gave us voting projections that closely matched the election result. Between 2010 and 2015 there was a significant increase in the proportion of Labour supporters saying that they were absolutely certain to vote. However, this expected increased turnout of Labour supporters (and turnout overall) did not materialise on Election Day. Had we assumed that the turnout differential between the Conservatives and Labour would be the same as it had been in 2010, our final poll would have forecast the shares as Conservative 38%, Labour 32% - pretty close to perfect.

2) Changes we've already implemented to improve our polling

There are two main ways the error in our polls may have arisen; through under-representing politically-disengaged people (particularly among groups most likely to support the Labour Party), and in our predictions of turnout among those we did reach. While the two are related, and have similar impacts, they need to be tackled in different ways. Under-representation of politically-disengaged members of certain groups is a sampling problem, and needs to be addressed by improving our sampling or weighting methods. Turnout exaggeration by respondents is a problem in projecting behaviour from the collected data, and needs to be addressed by improving the questions that we use to predict respondent turnout or the way in which we interpret their responses. We have adopted interim measures that incorporate both of these factors to improve our ongoing

monthly polling, but we will be considering the recommendations of today's report together with our own further experimentation refine these going forward.

2.1) Improving representation of politically disengaged people in our samples

The under-representation of politically disengaged people can be addressed either by changing sampling procedures to increase the numbers of this group included in the raw samples, or by weighting the data after fieldwork. Improving the raw samples is preferable because heavily-weighted samples have a higher "margin of error". Although doing so could increase the time and resources required to carry out a poll and could be impractical, especially during a campaign when quick-turnaround surveys are required. However, given there will never be a 100% response rate to political polls, it is likely some weighting will always be necessary.

Finding measures of political engagement which could be used to adjust data appropriately is a challenge. As an alternative, we have sought to address the issue of political disengagement indirectly, by using other measures which are related to levels of engagement.

For example, we have found on investigation that the claimed levels of readership of various newspapers in our Political Monitor surveys are in some cases significantly different from the levels indicated by the National Readership Survey (NRS), in ways which suggest over-representation of those people who take an interest in politics. For example, the numbers saying they read some of the main quality newspapers are much higher than the NRS suggests should be the case in a representative sample.

As we announced in September, we have started applying additional weights to our polls to draw the results in line with the NRS. Having run several waves since then, we can see this has had an impact on the proportion of claimed voters, reducing them by an average of three percentage points each month. This has had an effect on voting share, primarily for Labour, reducing their headline vote share by an average of 2.67 points, and increasing the Conservative headline vote share by 1.33 percentage points. (We cannot measure exactly the impact this change would have had on our final pre-election poll, as it did not include the more detailed questions which we are now using to measure readership.)

2.2) Improving predictions of respondent turnout

We have now replaced the single-question turnout indicator which we used at the election with a two-question measurement. Respondents are now asked "How likely would you be to vote in an immediate general election, on a scale of 1 to 10, where 10 means you would be absolutely certain to vote, and 1 means that you would be absolutely certain not to vote?", and also "Q. Which of the following best describes how often you vote in general elections?" with a choice of answers[2].

Previously we defined expected voters as being those who answered "10 – absolutely certain to vote". We now define them as being those who answer "9" or 10" to the first question *and* that they "always" or "usually" vote or that "it depends" to the second question. This model is informed by Ipsos' polling experience in elections around the world, and the rationale behind this change is that

[2] "I never vote", "I rarely vote", "I sometimes vote", "I usually vote", "I always vote" and "It depends".

when likely turnout is being exaggerated we expect that exaggeration to be greater among those who are not in the habit of voting. We do not assume that this definition literally identifies those who will vote from those that will not, but that the relative proportions of voters among the supporters of the different parties will reflect the difference in turnout.

During our three campaign polls (two in April and the final eve-of-election poll in May), the average effect of the new turnout filter would have been to correct for some of the error, by increasing the Conservative share of the vote by 1 percentage point and reducing the Labour share of the vote by 2 percentage points, increasing the Conservative lead by an average of 2.3 percentage points overall.

Since the general election, the effect of the new filter on average in our polls was to increase the Conservative share of the vote by 0.5 percentage points and to reduce the Labour share by 0.3 points, thereby increasing the Conservative lead over Labour by just under one percentage point. One possibility therefore is that the filter has more impact during campaigns than during 'peacetime', perhaps because during campaigns people who are not regular voters are more likely to feel they should vote (indeed, the biggest impact the new filter had was on our final poll).

Nevertheless, although we think that the current measure is an improvement, we will continue to search for new evidence on which questions will best help us better identify which respondents will vote. We may also conclude that different measures are necessary to correct for turnout inequalities in devolved elections, European elections or referendums from those which are suitable for polling in general elections.