



The Benefits of Random Sampling
Lessons from the 2015 UK General Election

Author

John Curtice

Introduction

At the 2015 general election, the opinion polls collectively underestimated the Conservative lead over Labour. In their final polls they suggested on average that the two parties were neck and neck on 34% each, when in the event the Conservatives proved to be seven points ahead. Moreover, even when the polling companies reinterviewed those whom they had polled before the election, they still obtained much the same neck and neck result.

This paper reports the results that were obtained by NatGen's 2015 British Social Attitudes survey (BSA) when it asked all of its 4,328 respondents whether and how they voted in the election. BSA is conducted very differently from the opinion polls. Whereas the polls interview those who can be contacted during a short period of time on the phone or over the internet, BSA is conducted face to face over a period of months amongst a randomly selected sample of respondents, with interviewers making multiple attempts to secure interviews with those selected for interview. The paper argues why this means the polls are at greater risk of obtaining samples of respondents who are unrepresentative of the nation at large.

In contrast to the polls, BSA 2015 replicates the outcome of the election quite closely. Just 70% said that they had voted, only a little above the official turnout figure of 66%. Meanwhile, amongst those who did say that they voted, support for the Conservatives is six points higher than that for Labour, very close to the near seven point lead actually recorded in the election.

Much of the speculation about the failure of the polls in the general election has focused on claims that voters were not being honest with the pollsters – and perhaps even to themselves. However, BSA's relative success in replicating the outcome of the 2015 election suggests that the problem of the polls lay instead in the unrepresentative character of the samples of people whom they interviewed, an error they were then unable to correct in their subsequent analysis and weighting of their data. The paper reports the results of some analyses of BSA's data that suggest where the source of that error may lie, looking in particular at the importance of identifying accurately those who are less likely to vote and the impact of an apparent tendency for Labour voters to be more accessible to pollsters.

The Problem of the Polls

There has been much speculation about why the polls collectively and significantly underestimated the lead of the Conservatives over Labour in the May 2015 general election (Cowley and Kavanagh, 2015, Chap. 9). Inter alia, it has been suggested that Conservative supporters were less willing to declare their intentions to a pollster, that those who said they were going to vote Labour instead stayed at home, and that polls were mistaken in asking people how they would vote at the beginning of the interview rather than towards the end, that is after they had been given an opportunity in other questions to express their lack of confidence in Ed Miliband and in Labour's economic competence (Morris, 2015).

Perhaps the samples of people who participated in the polls were not representative

What these explanations have in common is a presumption that the problem of the polls lay in the failure of respondents to reveal their true, underlying preference – not least perhaps because in some instances respondents themselves were not fully aware of the decision they would eventually make.

But there is, however, another possibility – that the problem lay in who was being asked the questions in the first place. Perhaps the samples of people who participated in the polls were not representative, containing too many people who did indeed vote Labour and too few who voted Conservative - and that the efforts pollsters made after collecting their data to correct any apparent inadequacies in their samples proved insufficient to remedy the problem.

One clue that this might be the case came from the results that a number of polling companies obtained when they went back to as many as possible of the people that they had interviewed in the weeks immediately before polling day and asked them how they had voted. As Table 1 shows, these exercises typically obtained much the same result as the polls had done in advance of polling day. Moreover, in addition to the exercises tabulated there, a similar exercise conducted by YouGov amongst those who had responded to one of their newspaper polls during the election campaign found that 84% of respondents voted the way that they said they would, and that the switching that did occur added no more than 1.1 points to the Conservative lead over Labour (Rivers and Wells, 2015).

Table 1 Reported Vote in the 2015 Election in Opinion Poll Recontact Exercises

	ICM %	Opinium %	Populus %	Survation %	BES/ YouGov %
Conservative	34	33	33	35	35
Labour	33	33	33	32	34
UKIP	12	12	15	15	10
Liberal Democrat	9	9	8	8	10
Green	5	4	5	n/a	4
Other	7	6	6	9*	7

* This figure includes the Greens.

Sources: National Centre for Research Methods (2015) except BES/YouGov: Internet panel study conducted by YouGov for the British Election Study (See Mellon and Prosser (2015)).

We can discount the notion that the error in the polls was primarily accounted for by a 'late swing' from Labour to the Conservatives

Of course, it is possible that those who voted Conservative were just as reluctant to declare their preference after polling day as they had been beforehand. However, if those who had said they would vote Labour did in fact disproportionately stay at home they would now have had to have lied about the fact that they had abstained. Meanwhile, if before the election respondents had overstated their willingness to vote Labour because they had been asked how they might vote before being asked what they thought of Ed Miliband or Labour's economic competence, that should now have revealed itself in a lower level of reported Labour voting in the ballot box. Certainly it seems that we can discount the notion that the error in the polls was primarily accounted for by a 'late swing' from Labour to the Conservatives, whether as a result of voters' dissatisfaction with Mr Miliband or for any other reason.

Unrepresentative Samples?

Why might the samples of people who respond to polls not be representative? Polls are not all conducted in the same way. A majority of those published during the 2015 election campaign were conducted over the internet, while the remainder were conducted by phone. In the case of the former, the respondents came from panels of persons who had indicated their willingness to take part in such surveys; in some cases the panel had been established by the polling company itself, while in others it was maintained and supplied by a commercial provider. The phone polls, in contrast, were conducted by ringing (landline and mobile) phone numbers at random, though rather than necessarily interviewing whoever answered the phone, interviewers would typically be given a target quota of people with different demographic characteristics that they would attempt to complete.

Although these two approaches are quite different from each other, both depart quite considerably from the approach that is ideally required by the statistical theory that underpins most polling and survey research (Frankel, 2013). This presumes that those who are interviewed have been selected at random from the group of people (such as all voters in Britain) whose views and attitudes we are trying to represent. More specifically it requires that selection is carried out using 'random probability sampling', that is by a random process that gives every potential respondent a known (though not necessarily equal) chance of being selected. If a sample has been drawn in this way then, as long as a sufficient number of people have been selected for interview, the views expressed by those who are interviewed (assuming a fair proportion of those who are selected actually participate) should provide a reasonably accurate portrait of the distribution of attitudes in the population at large. In particular, statistical theory allows us to calculate a margin of error around the survey estimate within which the level of support for a point of view amongst the public as a whole is likely to fall; for example, if a thousand people have been interviewed, this margin is typically three percentage points either side of the survey estimate.¹

1. More technically this is a 95% margin of error, which means that 95 times out of 100 the figure for the population as a whole will fall within ± 3 percentage points of the survey estimate.

Elements of random sampling are to be found in the way in which both internet and telephone polls are conducted. Those invited to participate in an internet poll may be drawn at random from the membership of the relevant panel of those who have indicated a willingness to participate in such polls. In a telephone poll the numbers that are rung are usually generated randomly. However, the members of an internet panel have not been chosen randomly – they have volunteered for or been successfully recruited into membership. Meanwhile, although over nine in ten people in Britain now own some kind of mobile phone, only just over four in five households have a working landline phone (Ofcom, 2015, Figure 1.6). That means some people are potentially accessible via both routes while others are not. In any event, as we have already noted, once a landline phone has been answered, who is chosen for interview is not necessarily determined at random; rather, the person selected is often whichever of those present who best helps the interviewer fulfill a target demographic quota.

Above all, however, both internet and phone polls are typically conducted over a very brief time period – often no more than two or three days. This means that there is relatively little time and opportunity to ensure that those who are selected do actually undertake the interview. The ‘response rate’ to polls, that is the proportion of those selected for interview who actually complete the survey, is typically not reported, but is thought to be low. Indeed, in the case of phone polls, rather than simply make repeated efforts to secure an answer from a preselected pool of telephone numbers, a company will often keep on ringing new numbers in the hope of securing an answer. The lower the response rate to a poll, the greater the risk that those who actually participate in a poll are not fully representative of the population, and in the case of opinion polls there is an evident risk that they are more likely to secure the participation of those who are easily accessible by phone or are regular users of the internet, as well as, perhaps, of those with a relatively strong interest in politics (Jowell et al., 1993, pp. 253-4).

Indeed, the relative difficulty that most polls have in securing the cooperation of some sections of the population can be seen in Table 2, which compares the demographic profile of ten polls conducted just before polling day in May 2015 with the profile that the company in question was attempting to achieve (as revealed by the demographic profile of the poll after it was weighted to remedy apparent deficiencies in its social and political profile). In all but one case (BMG) the poll interviewed fewer younger people than intended, while all of the polls to some degree at least interviewed more middle class voters than they wished. In subsequently weighting their polls so that they did match the intended demographic profile, the pollsters were trusting that the younger and working class people that they did manage to interview were typical of younger and working class people in general. However, there was no guarantee that this would be the case. It certainly meant that the younger and working class people that the pollsters did interview typically counted for more than one person when the final tally was being compiled, thereby ensuring that they had a particularly large impact on a poll’s estimate of overall voting intentions.

Table 2 Achieved and Target Demographic Profile of the Final Polls in the May 2015 UK General Election

	% Young People		% Middle Class	
	Achieved %	Target %	Achieved %	Target %
Ashcroft	24	29	56	55
BMG	32	29	58	55
ComRes	26	29	59	55
ICM	24	30	62	56
Ipsos MORI	22	29	59	54
Opinium	25	29	66	54
Panelbase	26	29	66	55
Populus	26	29	61	54
Survation*	28	29	44	40
YouGov**	30	37	61	56

*% Young People: Percentage aged 18-34, except YouGov (**), where it is percentage aged 18-39.*

% Middle Class: Percentage in social grades A, B or C1. Note that Survation () use their own social grade scheme for this purpose.*

Target: The percentage of respondents in the sample after weighting.

Source: Author's calculations based on computer tables posted by each company

In contrast, NatGen's annual British Social Attitudes (BSA) survey, for which the interviewing is undertaken by interviewers going out and talking to people face to face, is conducted according to the principles of random sampling (Ormston and Curtice, 2015). First, the addresses of potential respondents are drawn at random from the Postcode Address File, which contains virtually every residential address in the UK. Second, at each selected address the interviewer compiles a numbered list of all those aged 18 and over and the one that should be interviewed is then selected via a grid of random numbers. Nobody else can be interviewed. Meanwhile, every effort is made to ensure that a successful interview is conducted at as many of the originally selected addresses as possible. This is a time consuming process that lasts not just weeks but months, but it helps ensure that those interviewed are not just those who are easily contacted.

True, there are a few limitations to the process. Rather than selecting addresses at random from across the whole country, first of all a set of geographical areas (postcode sectors) is selected at random, and addresses are then selected from within those areas. This geographical concentration helps ensure that the selected addresses are all within reasonable travelling distance of each other, but slightly increases the risk that the sample is not as representative of the population as it might otherwise be. Similar geographical considerations also mean that no interviewing is conducted in the sparsely populated part of Scotland north of the Caledonian Canal. Meanwhile, we should bear in mind that despite the efforts that are made, the response rate to a survey such as BSA is far from perfect, and indeed has fallen since the survey was first conducted in 1983. In recent years, typically just over half of those whom it was hoped would participate have actually done so.

Still, we might wonder whether a survey like BSA that has been conducted according to the principles of random sampling, proves to be any more successful than the polls proved to be at replicating the actual lead that the Conservatives enjoyed over Labour in May 2015. If so this would add considerable weight to the suggestion that the error in the polls lay in the character of the samples that they obtained, a character that they proved unable to correct via the weighting and other reporting strategies that they deployed. As it happens, as part of the background information collected about each respondent, the 2015 BSA survey asked all of its respondents whether they voted in the May 2015 general election, and if so for which party. Moreover, thanks to the funding the survey secured (from diverse sources), the sample was rather larger than usual, comprising no less than 4,328 interviews. With interviewing conducted between 4th July and 2nd November, this represented a response rate of 51%. In the remainder of this briefing we examine how successful the 2015 BSA was in replicating the last general election result.

Demographics

To begin with we should look briefly at how successful the 2015 BSA was at achieving its target demographic profile. Was it any more successful than the polls at interviewing younger and more working class people? Or, despite the persistent efforts that are made to interview those who are difficult to contact, did it suffer much the same relative difficulty at securing the participation of certain sections of the population?

Table 3 Age Profile of the 2015 British Social Attitudes sample

	Achieved %	Target %
18-24	7	12
25-34	15	17
35-44	17	16
45-54	18	18
55-64	16	14
65+	28	22

Table 3 shows the age profile of those who were interviewed as part of the 2015 survey and the target age profile to which the data have subsequently been weighted. It is apparent that BSA is no more successful than the polls in securing interviews with those aged less than 35. Just 21% of all those interviewed fall into this category, compared with a target figure of 29%.² The deficit in the proportion of younger people who were interviewed does not, however, simply reflect a greater difficulty in getting younger people to participate in the survey. It also in part is a reflection of the way in which the random sampling was conducted. From our description above it is evident that only one person is ever interviewed in any household,

2. While the proportion of the achieved sample aged between 18 and 24 is 7% and that between 25 and 34 is 15%, rounding error means that the proportion of all those aged less than 35 is 21% rather than 22%.

irrespective of the number of people who live in that household. Thus those living in large households have a (known) lower probability of being interviewed than do those living in smaller households. Younger people are more likely than older people to be living in larger households, and thus are less likely to be selected for interview. If the 2015 BSA sample is simply weighted to reflect differences in the probability of each respondent being selected for interview, that alone is enough to increase the proportion aged less than 35 from 21% to 24%.

However, given that the BSA data are weighted to reflect the known age and gender (though not social class) profile of the population, but given also our interest in whether or not the sample of respondents is politically representative, in the analysis that follows we evidently should look at not only the weighted figures for participation and vote choice in the 2015 election but also at the unweighted ones. It may be the case that the survey's accuracy or otherwise is partly a consequence of the weighting that has been applied to the survey rather than simply because of the character of its sample.

Turnout

The turnout for Britain as a whole, calculated as the number of valid votes cast divided by the total number of persons registered to vote, was 66.4%. In practice, this will be a slight underestimate of the proportion of those eligible to vote who cast a ballot. A few votes will have been declared invalid (in 2015 these represented 0.7% of the electorate) while the electoral register will contain the names of some people who are not in fact eligible or able to vote (because, for example, they have recently died) (Electoral Commission, 2014; Rallings and Thrasher, 2015). At the same time, some people will (legitimately) be registered at more than one address but can only vote once. Thus we would anticipate that even the most accurate of surveys should report at least a slightly higher level of turnout than the widely quoted 'official' figure.

Table 4 Proportion Reporting Voted in the 2015 General Election in 2015 British Social Attitudes and British Election Study surveys

	Weighted %	Unweighted %
British Social Attitudes	70.3	71.0
British Election Study	73.6	73.8

Against that backdrop the proportion of BSA respondents that claim to have voted (see Table 4), 71% before the data are weighted, 70% afterwards, has to be regarded as demonstrating that the survey was highly successful at securing the participation of those who did not vote in the election. Indeed, it was even more successful than another random probability sample of how people have voted that has been conducted, a face to face post-election survey of 2,987 respondents conducted by gfkNOP for the academic British Election Study (BES) (British Election Study Team, 2015a). In this survey 74% said they had voted. One likely reason why the reported level

of turnout is lower on BSA than in the BES is that whereas the BES is entirely about politics and the election, the 2015 BSA covered a wide range of subjects including attitudes towards health, food, social class and work as well as more political subjects such as the European Union and the welfare state. As a result interviewers on BSA may well have found it easier than those on BES to persuade those with little interest in politics to undertake the survey. Certainly a similar pattern is to be observed at previous elections. For example, when the 2010 BSA asked (a randomly selected) one-third of its respondents whether or not they had voted, just 69.3% said they had done so, whereas the equivalent figure for that year's BES survey was 76.8%. Ironically it may be easier to secure a politically representative sample of respondents on a survey that is not wholly about politics.

In any event, both of the random probability exercises, BSA and the BES, were much more successful than non-probability ones at identifying those who did not vote. For example, the internet panel run by YouGov for the BES (to which reference was made in Table 1) reported that (even after weighting) 91.2% said that they had voted – a figure very much in line with the 89.7% turnout that YouGov report would have been anticipated if those they interviewed for their newspaper polls actually participated in the election in line with their previously stated intention to do so (British Election Study Team, 2015b; Rivers and Wells, 2015). Equally, on the same basis ICM's final election poll pointed to a 87% turnout while that conducted by ComRes pointed to a 90% one. Yet it is clear that there is no inherent difficulty in identifying and interviewing those who say that they did not vote in the 2015 election. This makes it unlikely that the reason why the polls could not come closer to the election result even when they recontacted their respondents after polling day was because some of those who in fact had abstained still insisted to the pollsters that they had voted Labour. Rather, the polls simply had difficulty in identifying those who never intended to vote in the first place.

Vote Choice

Still, in truth, polls are not primarily interested in forecasting the level of turnout. Rather their clients are looking for an accurate assessment of the balance of preferences amongst those who will participate in the election. It is thus their failure to estimate the relative size of Conservative and Labour support for which they have primarily been criticised.

But is a random probability approach any more successful at replicating the Conservative lead over Labour than the polls proved to be? Table 5 shows the reported level of support for each of the parties in the 2015 election amongst those BSA respondents who said that they did cast a vote. We can see immediately that their responses closely replicate the actual Conservative lead over Labour. In the BSA survey the reported level of support for the Conservatives is 6.1 percentage points higher than that for Labour, only a little less than the actual lead at the election of 6.6 points. Moreover, the survey's success in identifying a relatively large Conservative lead rests not on the weighting scheme that has been deployed – in fact, at 8.4 points, the Conservative lead in the unweighted data is

even bigger. In short, BSA had no obvious difficulty in finding the Conservative voters who apparently eluded the polls.

Table 5 Reported Vote in the 2015 British Social Attitudes survey

	Weighted %	Unweighted %	Election Result (GB) %
Conservative	39.7	40.8	37.8
Labour	33.6	32.4	31.2
UKIP	9.0	9.6	12.9
Liberal Democrat	7.3	7.4	8.1
Green	4.0	3.9	3.8
Other	6.4	5.7	6.2
<i>Con Lead over Lab</i>	<i>6.1</i>	<i>8.4</i>	<i>6.6</i>

The relative success of BSA in coming close to replicating the actual Conservative lead over Labour in the ballot box is also to be found in the BES post-election random sampling survey. Indeed, as Table 6 shows, that survey even slightly overestimated the Conservative lead over Labour. True, as is also the case for BSA, the proportion of UKIP supporters in the BES is a little lower than that in the ballot boxes, but between them the two surveys provide convincing evidence that the problem that plagued the polls – too few Conservative supporters relative to the number of Labour voters – was a consequence of the character of the samples that they obtained. If, instead, the problem arose because Conservative supporters are relatively reluctant to declare their preference, or Labour voters are overreporting their levels of attendance at the polling station, then both BSA and the BES should have had much the same (indeed given they were interviewing people face to face, maybe even greater) difficulty in replicating the Conservatives' lead.

The source of the polls' problem did not lie in how the pollsters' questions were answered but rather in who was answering them in the first place.

Table 6 Reported 2015 Vote in the BES Random Probability Sample

	Weighted %	Unweighted %	Election Result (GB) %
Conservative	40.6	39.9	37.8
Labour	32.7	32.3	31.2
UKIP	10.7	11.2	12.9
Liberal Democrat	7.1	7.5	8.1
Green	3.2	3.2	3.8
Other	5.8	5.9	6.2
<i>Con Lead over Lab</i>	<i>7.9</i>	<i>7.6</i>	<i>6.6</i>

Source: British Election Study Team (2015b)

That they have not done so suggests that the source of the polls' problem did not lie in how the pollsters' questions were answered but rather in who was answering them in the first place.

Why Might BSA Have Estimated the Conservative Lead Correctly?

Our earlier discussion suggested that there were two key differences between the way in which BSA is conducted and the approach taken by the polls. First, BSA is more likely to identify those with less interest in politics. Second, the survey is more likely to make contact with those who are not immediately accessible by phone or over the internet. How far might these differences help explain BSA's relative success in replicating the Conservatives' election lead?

We have already shown how the first of these features results in a reported level of turnout that is much closer to the actual level of participation in the election. But is there any reason to believe that the polls' overestimate of the likely level of participation may also have helped occasion their error in the estimate of the Conservative lead? One possibility, for example, is that the polls not only overestimated the overall level of participation in the election but also failed to estimate correctly the differences between social groups in their level of participation – and that as a result they overestimated the level of participation amongst those social groups in which Labour support was relatively high.

Table 7 Reported Turnout by Age in Random Probability Surveys

	BSA %	BES %
18-24	56	57
25-34	55	58
35-44	64	68
45-54	75	78
55-64	80	83
65+	84	88
<i>Age Gap*</i>	28	31

* difference between reported turnout amongst those aged 65 plus and those aged 18-24.

One persistent feature of turnout at elections is that younger voters are less likely to vote than their older counterparts (Phelps, 2004). Table 7 shows that this pattern was certainly identified at the 2015 election by both BSA and the BES. In both surveys those aged 18-24 were about 30 points less likely to say that they had voted than were those aged 65 or older. This is by no means exceptional. In the 2010 BSA, when whether or not people had voted in the election was asked of one-third of those who participated in that year's survey, the gap between the level of turnout amongst the youngest and the oldest age group was no less than 43 points.

What, however, does appear to be the case is that the relationship between age and whether people voted Conservative or Labour strengthened in 2015. Labour appears to have gained ground amongst younger voters while it was less popular than five years previously amongst older voters (see also Ipsos MORI, 2015). This pattern is certainly apparent in Table 8, which is based on the 2010

and 2015 BES surveys. Equally, if we compare the age profile of party support in BSA 2015 with that in the 2010 survey, we find that Labour support increased from 32% to 39% amongst those aged 18-34, but eased back from 33% to 30% amongst those aged over 55. Conversely, Conservative support held steady at 32% amongst the younger age group but increased from 42% to 46% amongst the older group.

Table 8 Conservative and Labour Support by Age

	% vote 2015		Change since 2010	
	Con	Lab	Con	Lab
18-24	29	45	+1	+17
25-34	28	47	-3	+6
35-44	33	42	-3	+10
45-54	44	28	+8	-5
55-64	44	27	+3	-2
65+	50	23	+4	-7

Source: British Election Study 2010 and 2015

Thus, any poll that overestimated the propensity of younger voters to participate was at particular risk in 2015 of overestimating Labour support. We have already seen that most polls certainly had difficulty finding younger voters to interview – though we have also seen that the polls are not unique in that respect. But perhaps those younger voters who participated in the polls were atypical of younger voters in general in that they were relatively interested in politics and thus more likely to turn out and vote, while at the same time they reflected the relatively strong Labour sympathies amongst those in their age group who did vote. If so, then such a pattern could have contributed to the error in the polls, especially given that typically those younger voters who were interviewed by the polls were subsequently given a relatively high weight in calculating the final estimate of voting intentions.

But did the polls underestimate the age gap in turnout? Typically the polls attempted to ascertain people's propensity to vote by asking them to say how likely they were to vote on a scale from either 1 to 10 or 0 to 10, with 0/1 meaning they were certain not to vote, while 10 indicated they were absolutely certain to vote. In calculating their final estimates polls either only took into account those who said they were certain or almost certain to vote, or else weighted respondents according to their reported likelihood of voting. In Table 9, therefore, we show two ways of estimating the age gap anticipated by the final polls conducted just before polling day. First, we show the difference between the proportion of 18-24 year olds who said they were certain to vote and the proportion of 65 year olds doing so. Second, we show the difference between these two age groups in the level of turnout implied by the score that those belonging to these two age groups gave themselves on average on the 0/1 to 10 scale.

Table 9 Age Gap in Anticipated Levels of Turnout in Final Election Polls

	Difference between those aged 18-24 and those aged 65 plus	
	% certain to vote	% turnout based on mean scores
Ashcroft	-27	-12
ComRes	-39	-22
ICM	-34	-21
Ipsos MORI*	-31	-25
Populus	-24	-9
YouGov**	-18	-10

* Ipsos MORI first asked their respondents whether they were registered to vote before asking how likely they were to vote. The gaps reported here are based on reported likelihood of voting multiplied by the proportion who said they were registered.

** Differences are between those aged 18-24 and those aged 60 plus.

Source: Author's calculations based on polling companies' published computer tables.

The polls certainly anticipated that younger people were less likely to vote than their older counterparts. Indeed, if we look at the differences in the proportion who said they were certain to vote then in the case of the first four companies listed in the table, all of whom conducted their polls by phone, the anticipated age gap in turnout is not dissimilar to the 30 point or so difference identified by both BSA and BES. However, the anticipated gap is somewhat smaller in the final two entries, both of which were conducted over the internet. Meanwhile, the difference in turnout anticipated by the average of the scores that respondents gave themselves on the 0/1 to 10 scale is in each case rather smaller, and in some instances is well short of the BSA and BES estimates of what actually happened. Focusing entirely on those who said they were certain to vote (as Ipsos MORI did) thus appears to have been a more successful strategy for identifying the age gap in turnout than weighting respondents by their reported likelihood of voting. However, most polls did not use this approach and thus for the most part the polls do appear to have been at some risk of underestimating the age gap in turnout. And given also that most polls were reporting higher levels of Labour support amongst younger voters than amongst older voters that potentially put them at risk of overestimating Labour support.

Meanwhile, BSA certainly confirms the impression that collectively Labour voters were more likely to stay at home. As well as asking people whether or not they voted, and for whom, BSA also asked all of its respondents a sequence of questions designed to ascertain with which party, if any, they identified or felt close to. As a result we have an indication of the partisan sympathies of those who did not vote. Unsurprisingly, nearly half (47%) of those who abstained said they did not identify with any party or did not know whether they did or not. But amongst the remainder Labour identifiers (24%) were more numerous than Conservative ones (14%). Indeed, whereas 86% of those who said that they identified with the Conservatives turned out to vote, just 76% of those who stated that they were Labour supporters did so. The BES also suggests that there was a similar ten

BSA certainly confirms the impression that collectively Labour voters were more likely to stay at home

point difference in the relative propensity of Conservative and Labour identifiers to participate in the election.

But what of the second key feature of BSA, viz. that it goes to considerable trouble to interview those who are not immediately contactable. Interviewers were instructed to make at least six attempts to make contact with someone at each address with which they were issued, and in some instances further efforts were made thereafter, sometimes by a different interviewer. In a few instances (comprising 2% of the sample), as many as nine calls were made before an interview was successfully obtained. Only in one in eight cases was an interview obtained at the first time of calling.

So far as turnout is concerned, those who were interviewed on first contact were not markedly more likely to say they had voted – just 70% said that they did so. Indeed, the highest level of participation was reported amongst those who were interviewed after between three and five calls; as many as 73% of this group said that they voted. However, turnout was markedly lower amongst those who were only interviewed after six or more calls; just 64% of this group said that they voted. In part at least this reflects the age profile of those who were most difficult to contact; just 17% of this group were aged 65 and over, compared with no less than 35% of those interviewed after just one or two calls. But this simply underlines the extent to which getting hold of those who are difficult to interview – who constituted nearly one in five (18%) of the BSA sample – is important if a representative sample is to be obtained, and especially so when it comes to securing the participation of those who are less politically interested and engaged.

However, is there any reason to believe that those who were more difficult to contact were distinctive in the way in which they voted? It seems that there is. Table 10 shows the distribution of reported vote in the general election broken down by the number of calls that had to be made before an interview was conducted. We can see immediately that those who were the most accessible for interview, that is they were interviewed the first time an interviewer called, were markedly more Labour and less Conservative in their sympathies than were even those who were interviewed on the second call, let alone those who were only interviewed after between three and six calls. Only amongst those who were the very hardest of all to reach (a group that as defined in Table 10 consists of less than one in ten of all those who voted and which consists disproportionately of younger voters) was the pattern reversed.

Table 10 Reported Vote By Number of Calls Made to Achieve Interview

	No. of Calls			
	1 %	2 %	3 to 6 %	7 to 9 %
Conservative	35	39	42	34
Labour	41	33	31	42
UKIP	9	10	9	6
Liberal Democrat	6	6	8	10
Green	4	5	4	2
Other	4	6	7	6
<i>Sample Size (unweighted)</i>	371	678	1619	272

Source: British Social Attitudes 2015

The relatively high level of support for Labour amongst those who were interviewed first time around cannot obviously be accounted for by the social character of these respondents. As we have already noted, they contain an above average proportion of older people, a group that in general was relatively unsympathetic to Labour in the election. True, those interviewed on the first call were rather more likely to be engaged in a routine or semi-routine (working class) occupation (35% of first time respondents fall into this category as compared with 30% in the sample as a whole), but even if we weight these respondents such that their class profile matches that of the sample as a whole, we still obtain a five point Labour lead.

Those who were most easily interviewed by BSA interviewers appear to have been more likely to support Labour and less likely to support the Conservatives

In short, those who were most easily interviewed by BSA interviewers appear to have been more likely to support Labour and less likely to support the Conservatives to a degree that cannot be accounted for by the social profile of these respondents. If indeed this group in any way mimics the kind of person who was most likely to respond to the polls, then we can begin to understand why the polls might have overestimated Labour's strength. Those who are interviewed most easily tend to be distinctive in their political views. Indeed, if we widen our scope and look at those who were interviewed on either the first or the second call, perhaps as close as we can get to replicating the kind of person most likely to be interviewed by the polls, we obtain an election outcome very similar to that reported by the polls – a Conservative lead over Labour of just two points.

There is then some support for both of our possible explanations as to why the random sampling approach used by BSA has been more successful than the polls in replicating the Conservatives' lead over Labour in the 2015 election. First, because BSA was more successful at identifying non-voters it also uncovered a stronger relationship between age and turnout, and in 2015 at least any underestimation of that relationship ran the risk of overestimating Labour support. Second, by making strenuous efforts to contact as many people as possible BSA was better able to ensure that those it interviewed were indeed representative of the partisan mood of the country. Labour voters, it seems, are too easy to find, and as a result a poll that focuses on the easily accessible runs the risk of overestimating Labour's support.

Conclusion and Implications

Opinion polls are intended to provide their journalistic clients with a relatively inexpensive way of securing a reading of the very latest political weather. They need to be inexpensive because of limited budgets, while a timely reading is often wanted because of a wish to establish whether a recent political event or development has changed public opinion – even though few such events or developments prove to have any electoral consequence. As a result polls are conducted in a way that does not fully meet the requirements of random sampling. A key lesson of the 2015 election is that, as a result, they run the risk of failing to take the political temperature correctly. Their approach at that election resulted in deficiencies in their samples that subsequent weighting and filtering of the data failed to correct. Whereas the polls still largely put Labour and the Conservatives neck and neck even when they asked people after the election how they had voted, two major post-election surveys that used random sampling, BSA and the BES, have both been able to replicate the Conservatives' 6.6 point lead reasonably accurately.

British Social Attitudes has, in truth, very different aims and objectives from most polls. Conducted annually for more than 30 years, it attempts to identify long-term changes in the climate of public opinion rather than short-term changes in the political weather. It also endeavours to provide high quality data that make it possible to identify who is more likely to hold a particular viewpoint and why. To achieve those objectives the survey needs to ensure that it covers all sections of British society - not just those with a strong interest in a particular subject, be it politics or anything else, and not just those who can be got hold of easily. And while lower response rates have undoubtedly made random sampling more challenging, BSA's relative success in replicating not only the Conservative lead over Labour in the 2015 election but also the level of turnout at that election has demonstrated that, relatively expensive and time-consuming though it may be, random sampling remains a far more reliable method for securing representative samples - and thus building a solid foundation upon which to make claims about what Britain thinks. The time and effort involved in random sampling pays dividends, and thus it is clearly the approach that should still be used by anyone with a serious interest in understanding British public opinion.

BSA and the BES have both been able to replicate the Conservatives' 6.6 point lead reasonably accurately.

References

British Election Study Team (2015a), 'British Election Study 2015 Face-to-face survey v1.0: Release note'. Available at <http://www.britishelectionstudy.com/bes-resources/f2f-v1-0-release-note/#.VpKt1VI6zTM>

British Election Study Team (2015b), 'The Benefits of Random Probability Sampling: The 2015 BES Face-to-Face'. Available at <http://www.britishelectionstudy.com/bes-resources/f2f-release/#.VpKyill6zTM>

Cowley, P. and Kavanagh, D. (2015), *The British General Election of 2015*, Basingstoke: Palgrave Macmillan.

Electoral Commission (2014), *The Quality of the 2014 Electoral Registers in Great Britain*, London: Electoral Commission.

Frankel, M. (2013), 'Sampling Theory', in Rossi, P., Wright, J. and Anderson, A., *Handbook of Survey Research*, New York: Academic Press.

Ipsos MORI (2015). 'How Britain Voted in 2015'. Available at <https://www.ipsos-mori.com/researchpublications/researcharchive/3575/How-Britain-voted-in-2015.aspx?view=wide>

Jowell, R., Hedges, B., Lynn, P., Farrant, G. and Heath, A. (1993), 'The 1992 British Election – The Failure of the Polls', *Public Opinion Quarterly*, 57: 238-63.

Mellon, J. and Prosser, C. (2015), 'Why did the polls go wrong?'. Available at <http://www.britishelectionstudy.com/bes-resources/why-did-the-polls-go-wrong-by-jon-mellon-and-chris-prosser/#.VpKdDFI6zTO>

Morris, J. (2015), 'Late Swing? Labour's private polls showed the Tories ahead before Christmas', *New Statesman*, 11 May.

National Centre for Research Methods (2015), Polling Inquiry Documents. Available at <http://www.ncrm.ac.uk/polling/document.php>.

OfCom (2015), *The Communications Market Report 2015*, London: OfCom

Ormston, R. and Curtice, J. (eds) (2015), *British Social Attitudes: the 32nd report*. Available at <http://www.bsa.natcen.ac.uk/latest-report/british-social-attitudes-32/key-findings/introduction.aspx>

Phelps, E. (2004), 'Young citizens and changing electoral turnout, 1964-2001', *Political Quarterly*, 75: 238-48.

Rallings, C. and Thrasher, M., *The 2015 General Election: aspects of participation and administration*. Available at http://www.electoralcommission.org.uk/our-work/publications/policy-and-research-reports?&form=subsearch&collection=electoral-commission&profile=default&formtype=&formtypedefault=policy-research&&scope=&start_rank=11

Rivers, D. and Wells, A., *Polling Error in the 2015 UK General Election: An Analysis of YouGov's Pre and Post-Election Polls*. Available at www.yougov.co.uk