

YouGov Polling at the 2024 Election

December 2024

Prepared by: Anthony Wells

Head of European Social and Political Research



YouGov[®]

Contents

Performance of YouGov Polling and MRP at the 2024 Election	1
Potential Causes of Error.....	4
Who respondents voted for	4
Whether Respondents Voted	8
Whether the respondents interviewed accurately represented the public	11
MRP modelling decisions	17
Conclusions and recommendations	19

Performance of YouGov Polling and MRP at the 2024 Election

Our polling for the general election focused upon our MRP modelling. We produced three large MRP models at the beginning, middle and end of the election campaign. These used samples of between 39,979 and 58,875 respondents and produced party vote intention estimates for each seat in Great Britain, as well as an implied national vote share.

In between these large MRP models we produced twice-weekly voting intention numbers for Sky and the Times. These used a smaller version of our MRP model, using a sample of c. 2000 and producing only GB-level support estimates.

Prior to the general election we had produced voting intention figures using a more traditional method (standard sample using rim weighting and additional turnout weighting). We dropped this method in favour of the small MRP at the beginning of the campaign and did not publish voting intention figures using the old method after 5th June.

Our final call for the general election was a large MRP model published on 3rd July. Our final published figures using the smaller MRP model was 25th June. However, we continued to process these figures internally. Our final internal run of the smaller MRP was also on 3rd July. Finally, while we ceased to use our traditional method on 5th June, we have been able to calculate what the figures would have been under the old method, had we continued to use it.

Comparisons between these and the final result in terms of seats and vote share are below:

Seats		CON	LAB	LDEM	SNP	RefUK	GRN
Actual Result	4th July	121	411	72	9	5	4
Final call (Large MRP)	19th Jun- 2nd July	102	431	72	18	3	2
	<i>Error</i>	-19	+20	0	9	-2	+2

Vote Shares		CON	LAB	LDEM	RefUK	GRN	<i>Average error</i>
Actual Result	4th July	24	35	13	15	7	
Final call (Large MRP)	19th Jun- 2nd July	22	39	12	15	7	
	<i>Error</i>	-2	+4	-1	0	0	1.4
Small MRP (unpublished)	1st-2nd July	21	34	14	17	9	
	<i>Error</i>	-3	-1	+1	+2	+2	1.8
Traditional method (unpublished)	1st-2nd July	18	38	11	20	9	
	<i>Error</i>	-6	+3	-2	+5	+2	3.6

YouGov’s final seat projection was the most accurate of the MRP polls in terms of seats called, with 92% correct (and was marginally more accurate than the BBC/ITN/Sky exit poll). It was also the second best across the whole industry in terms of minimising error on the national vote share estimates – including all MRP and regular vote intention polls.

Our final small MRP model numbers would have been more accurate on Labour, but slightly less accurate overall. Our figures produced using our old methodology would have been substantially out, confirming that the methodology switch at the beginning of the campaign was the correct decision.

Overall, the wider polling industry was perceived as having suffered a significant polling miss. All companies overstated Labour support and understated Conservative support. While our final figures were better than most other companies, they were out in the same direction as everyone else, and overstated Labour by four points. In this instance this error did not change the outcome – it was still a Labour landslide – but in a closer race it could have produced a far more serious error.

This paper will seek to diagnose what went wrong in the traditional old method and what could have been done better using the new small and large MRP methods, and to make recommendations on how we should proceed in the future in terms of both our standard polling samples and our vote share predictions.

Analysing the causes of polling error is difficult. The secret ballot means there is no official data on how people of different demographics voted to compare polling results to. In practice, the only

approach is that taken by the 2015 polling inquiry and those that came before it - that is, to take possible hypotheses on potential cause of error one at a time and examine whether there is evidence that they contributed to the error. That is the approach that this paper takes.

In terms of data, YouGov collected a large amount of data using the campaign, both to drive the MRP model and to accommodate a large number of Omnibus surveys for clients. The dataset used for most of the comparisons below is a merge of our data from political omnibus surveys for the two weeks leading up to election day, providing a dataset of 31,072 respondents. After the election we recontacted as many as possible of these respondents to ask if they had voted and for whom. Where appropriate, this sub-sample of recontacted respondents is used. Some analysis also uses the dataset of respondents to the final MRP model, a dataset of 47,751 respondents.

The potential sources of error are roughly grouped into four sections:

- Who a respondent voted for
- Whether a respondent voted
- Whether the respondents interviewed accurately represented the public
- The MRP modelling decisions

Potential Causes of Error

Who respondents voted for

“Late Swing” – Respondents changing their mind (or making up their minds late)

Labour’s vote declined throughout the campaign. This downwards trend in Labour support has led to suggestions that part of the polling error was the result of a “late swing” – that is, respondents continued to desert Labour between the final eve-of-election polls and actually casting their votes.

Another commonly suggested potential cause of error is don’t knows, that is, people who tell polling companies they don’t yet know how they will vote, but do actually vote in the end. If these people break heavily in favour of one party or another it has the potential to cause error.

A final consideration hypothesis is the question of tactical considerations – whether people correctly report local tactical considerations when responding to a survey, or give the party they “really” support.

These are separate potential causes of error, but in terms of diagnosing a polling error they are all issues that end up being diagnosed in the same way – by finding a difference between how respondents answer the voting intention question before the election and how they report having actually voted after the election.

The falling trend in Labour support and a gradual increase in Conservative support prior to the final polls is evident. The table below breaks down the raw, unweighted and unmodelled data gathered for our final MRP poll by day of fieldwork, and shows a clear upwards trend for the Conservatives and downwards trend for Labour.

Sample date	Con share	Lab share	Lab lead	Sample size
19 th June	16.3	41.8	25.5	1289
20 th June	18.2	42.5	24.3	1253
21 st June	17.4	42.3	24.9	1377
22 nd June	17	44.6	27.7	1710
23 rd June	18.3	43.2	24.9	1798
24 th June	16.7	40.3	23.6	1984
25 th June	16.6	41.2	24.6	3229
26 th June	17.8	39.7	21.9	3356
27 th June	19.9	39.1	19.3	3827
28 th June	19.9	38.6	18.7	4738
29 th June	20	38.2	18.2	4500
30 th June	20.5	39	18.4	3341
1 st July	19.5	38.9	19.4	5491
2 nd July	20.5	37.4	16.9	3852

Turning to the data from our political omnibus samples during the campaign and recontact data from the same respondents, there is also a clear shift from their pre-election vote intention to their post-election recalled data.

	CON	LAB	LDEM	RefUK	GRN
Pre-election VI (last two weeks)	17.8	39	11.6	20.6	7
Post election recalled vote	18.5	36.3	13.2	18	6.9
Change	0.7	-2.7	1.6	-2.6	-0.1

	CON	LAB	LDEM	RefUK	GRN
Pre-election VI (final poll)	18.2	38.1	11.2	20	8.8
Post election recalled vote	18.5	35.7	12.2	18.7	8.6
Change	0.3	-2.4	1.1	-1.3	-0.2

This is consistent with part of the overstatement of Labour and Reform UK being the result of people changing their behaviour between the final polls and the actual election. This would also be potential cause for the overstatement of Labour in our final large MRP, which included data collected from 19th June to 2nd July.

Don't knows – respondents who did not give a pre-election voting intention

In our final sample on 3rd July, 7% of respondents said that they did not know how they would vote. In some cases this response will be a more socially acceptable way of saying they will probably not bother to vote, in other cases these will be voters who are genuinely undecided or do not wish to tell us their true intentions. This includes the phenomenon of people saying “don't know” because they are reluctant to admit supporting an unpopular party (commonly referred to as “shy Tories”). “Don't knows” are a potential source of error if a significant number of these respondents do vote, and vote differently from those voters who did provide a voting intention.

Looking at the observed behaviour we saw in 2024 and looking at the post-election self-reported behaviour of those respondents who said don't know in pre-election voting intention questions, only 16% of those who said don't know subsequently reported not having voted. The table below compares how the “don't knows” in our political omnibus surveys in the two weeks prior to the election reported having voted when recontacted after the election.

	CON	LAB	LDEM	RefUK	Other
National vote shares	24	35	13	15	16
How "don't knows" report having voted	25	24	18	13	21

Compared to the public at large, our don't knows were significantly less likely to end up voting Labour and more likely to vote Liberal Democrat or Other.

There are different approaches to dealing with don't knows. In our traditional methodology we used to exclude respondents saying don't know from our topline figures and did not replace them. The practical impact of this in 2024 would have been to overstate Labour by excluding a group who were less likely to vote Labour than average.

For the small and large MRP methods used at the 2024 election, respondents saying don't know were also removed from the model, but before the post-stratification, which had a different practical effect. In practice, the MRP modelled their behaviour, probabilistically modelling the extent to which they would vote the same way as other people of the same demographics, who were acting in the same political context. This would have had less of a negative impact.

Tactical voting – respondents voting for a party other than the one they supported

The final broad area of potential error in respondents giving their voting intentions is tactical voting. The difference in swing between areas where Labour and the Liberal Democrats were best placed to beat the Conservative Party suggests a high level of tactical voting at the 2024 election and this would be a potential cause of error if, for example, a respondent who supported the Labour Party but was voting tactically for the Liberal Democrats answered Labour to the voting intention question. For the purposes of predicting election results, our aim should be to collect data on who respondents will actually vote for, rather than the party they may genuinely support.

YouGov asked voting intention using two different question formats during the 2024 campaign. The standard VI question asked people how they would vote in the general election. The constituency VI question was asked as a follow up and asked respondents to specifically consider the parties and candidates standing in their own constituency. These were deliberately asked one after the other to allow respondents to express support for a national party should they wish, in the hope that this would make respondents more comfortable with giving a tactical vote without it necessarily being interpreted as support for that party nationally. This approach was developed in order to address an observed problem with polling in Liberal Democrat target seats, where a standard VI question often suggests a collapse in Liberal Democrat support and an increase in Labour support, which we theorised was actually just a failure of the standard question to pick up tactical considerations.

Our traditional methodology had used the standard VI question wording. Both our small and large MRP models used the constituency VI wording, and this was one of the reasons behind the differences in their figures.

Small MRP (average of polls over last week)

	CON	LAB	LDEM	RefUK	GRN
Standard VI Q	19.6	38	12.7	18.2	7
Constituency VI Q	19.6	36.9	13.8	17	6.9
Difference	0	-1.1	1.1	-1.2	-0.1

Comparing how the different questions would have worked, the differences at a national level are meaningful. The constituency prompt is more accurate for both Labour and Reform, and less accurate for the Liberal Democrats, but overall this results in a lower average error.

Whether Respondents Voted

Turnout has been a common factor in historical polling errors. Opinion polls consistently find more people saying they will or have voted than the actual turnout at general elections. Because turnout correlates with demographics that also correlate with voting intention (e.g. young people are both more likely to vote Labour and less likely to vote at all) this is a major potential cause of polling error.

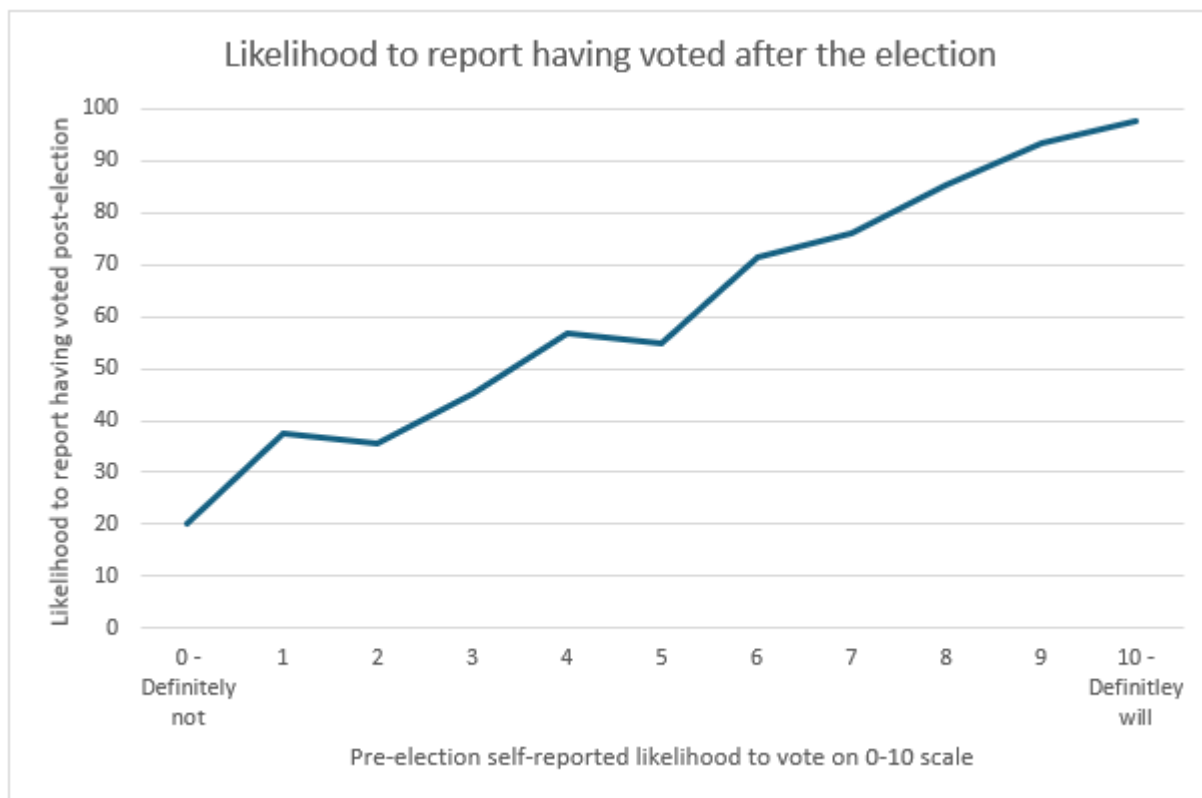
In the 2015 general election the error in the polls occurred partially through the mechanism of turnout - samples contained too many engaged people and, therefore, had too many people who voted, which consequently understated the gap in turnout between Labour and Conservative voters and understated the Tory lead. Those polls that got the 2017 election wrong also largely erred on the issue of turnout, having sought to correct the problems of 2015 by adopting demographic turnout models, but doing so in a way that skewed their samples towards the Conservative Party.

It is important to note there are two different ways of getting turnout wrong:

- 1) Wrongly predicting individual respondents' likelihood to vote – that is, having a model that assumes a respondent will vote when they will not, or vice-versa.
- 2) Having a sample with the wrong proportion of people who will vote – that is, having too many of the sort of people who vote, and not enough people who will stay at home.

It would be possible to ask questions or create a model that very accurately predicts whether individual respondents will or will not vote, but still be wrong if there are too many likely voters in the sample.

In the past our standard voting intention method used an approach where turnout was predicted based upon a combination of people's self-assessed likelihood to vote and whether they voted at the previous election. Comparing self-assessed likelihood to vote with post-election recalled vote suggests that this continues to be a broadly accurate way of predicting respondents' likelihood to vote relative to one another.

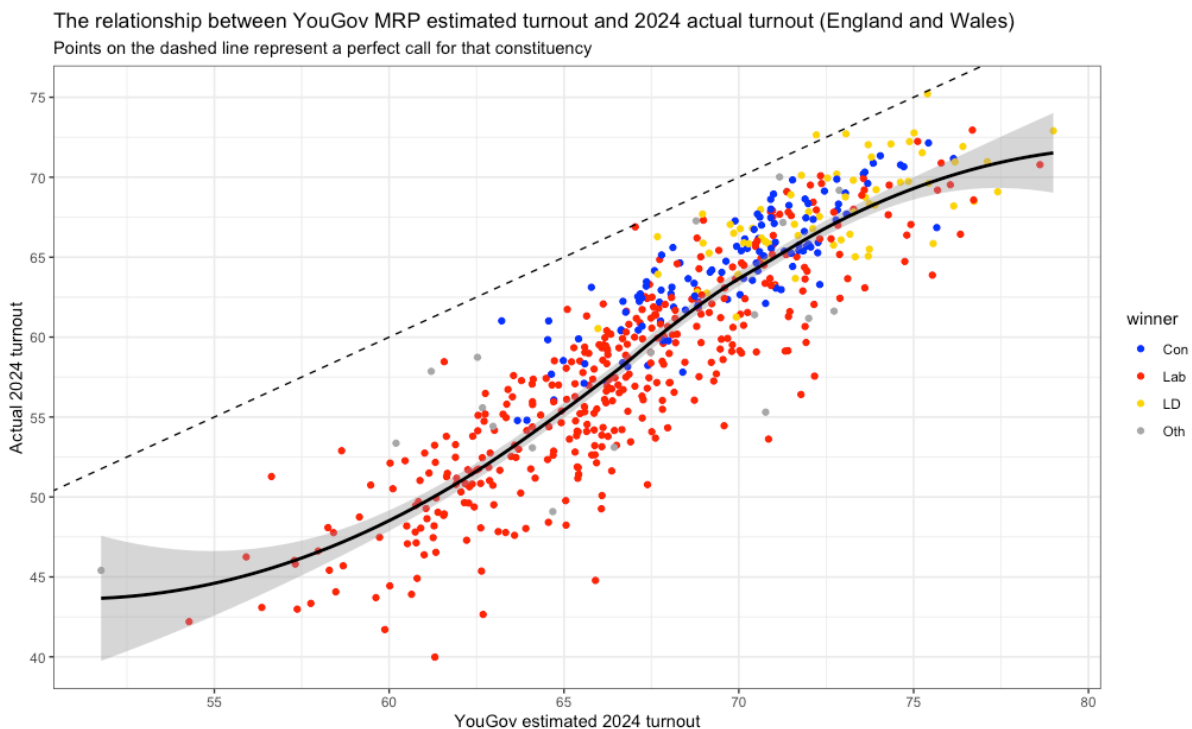


However, while this approach was broadly accurate, the proportion of respondents in our sample who vote remains far too high. Our final poll using the old method would have shown an implied turnout of 72%. In comparison, the actual turnout was 60% among those registered to vote. (Given the electoral registers are significantly incomplete, the turnout among the entire adult population would have been significantly lower than this). This suggests the root problem is having too many people who are likely to vote in the sample, rather than wrongly predicting respondents' likelihood to vote.

This is likely to still be having the impact of overstating Labour's vote. Typically those demographics who are more likely to vote Labour (most significantly younger respondents) are also less likely to vote. Overstating turnout risks too small a turnout gap between low and high turnout demographics, ultimately overstating Labour support.

In the absence of individual-level data from the face-to-face British Election Study, it is not possible to be certain that the overstatement in turnout particularly overstated Labour support, but the pattern of turnout at the election implies a continuation of historical patterns of lower turnout in more deprived, younger, urban constituencies that tend to vote Labour, and that having too many likely voters is likely to have contributed to an overstatement of Labour support.

Turnout also had a significant impact upon the accuracy of the MRP model. Our MRP model used a mixture of demographic modelling and self-assessed turnout to predict turnout at the group level. While it is not the intention nor design of MRP models to predict actual overall aggregate turnout, the fact that the final call MRP assumed a turnout of 68% is indicative perhaps of an underlying problem. This figure was closer than the implied turnout from the vote intention polling, but still an overestimate. If we had instead relied purely upon demographics to model turnout it would have produced lower turnout probabilities for certain groups, which would in turn have reduced both the level of Labour support and Labour seats in our final MRP model.



As we can see from the graph above, the MRP model systematically overestimated turnout most, in seats which Labour went on to win. In seats won by the Conservatives and Liberal Democrats, our turnout error is significantly lower. This strongly suggests that we overestimated Labour voter propensity in particular, which in turn would have been a strong contributor to the vote share errors at the national level.

Whether the respondents interviewed accurately represented the public

The Brexit Party and past-vote weighting

Since the 1990s past vote weighting has become a dominant feature in voting intention polling. Past vote weighting remains somewhat controversial, particularly in the USA, because of concerns over false recall. However, panel-based methodologies such as YouGov's largely address this issue as respondents' previous vote can be asked contemporaneously with the election, rather than relying on later recall. Almost all, if not all, companies using an online panel methodology in the UK currently use past vote weighting.

Past vote weighting is a useful tool because it correlates well with current behaviour. For example, the people who are most likely to vote Labour are people who voted Labour last time. If supporters of any political party are more likely to end up participating in polls, then past vote weighting is a highly effective tool at correcting for this.

In 2024 it is possible that past vote weighting was less effective because of the Brexit Party's decision to stand down in most seats in 2019. In practice, the Brexit Party effectively endorsed the Conservative Party. The practical implication of this is that in 2024, past vote weighting was not an effective tool in distinguishing between Conservative supporters who voted Conservative in 2019, and Brexit Party supporters who voted Conservative in 2019. Given one of the consistent errors in the 2024 polling was overstating Reform UK party support and understating Conservative support, this hypothesis would explain a substantial amount of error across all companies.

By definition this is impossible to test by looking at recalled 2019 vote, but we can compare recalled 2017 and 2015 vote among the samples we used in 2019 (which used 2017 past vote), and the samples we used in 2024 (which used 2019 past vote). Comparing these figures does suggest that our 2024 samples contained a larger proportion of 2015 and 2017 UKIP supporters among 2019 Conservative voters – a number that should have remained steady between the two samples. This evidence is consistent with past vote weighting in 2024 having been less effective at ensuring the correct balance of Conservative and Brexit Party/Reform UK supporters in the sample.

	% of 2019 Conservative voters who voted UKIP in 2017	% of 2019 Conservative voters who voted UKIP in 2015
YouGov samples at the 2019 election	1.7%	12.0%
YouGov samples at the 2024 election	3.9%	16.5%

Unlike other potential issues described in this section, this source of error would not suggest any immediate methodology change. Reform UK properly contesting the 2024 removes the anomaly and past vote weighting should, at the next election, return to being an effective tool. It should nevertheless be a warning of the shortcomings of relying upon past vote weighting in scenarios where a new party has emerged, or a party did not, in practice, contest the previous election.

Age

Grouping respondents aged over 65 together rather than looking separately at over 75s was highlighted as a potential source of difference between polling companies during the campaign. The highest age band in the weights for our traditional samples was 65s, while our MRP model included a separate age band for over 75s.

Prima facie there were differences between the voting behaviour of 65-74 year olds and over 75s, with over 75s being more consistently Conservative and less likely to vote Labour and Reform than 65-75 year olds.

	CON	LAB	LDEM	RefUK
Pre-election VI (raw data) 65-75	34	26	11	17
Pre-election VI (raw data) 75+	44	19	13	14

Our evidence was inconsistent about whether this would actually have had an impact. While over 75s were more likely to vote Conservative, they were also more likely to have voted Conservative in 2019, and the size of the fall in Conservative support was similar among both 65–74-year-old Conservatives and 75+ Conservatives. This meant that in practice, adding an over 75 band would have made less difference to the Conservative share of the vote than might have been expected, as it would have interacted with past vote weights (increasing the number of over 75s would have increased the number of 2019 Tory voters aged over 75, but reduced the number of 2019 Tory voter aged under 74, who behaved in a similar fashion). Weighting or using quotas for over 75s would, however, have reduced the overstatement of the Reform UK Party in our traditional method. The difference is because of ineffectiveness of past voting for the Brexit Party/Reform UK (as discussed in the previous section).

Our voting intention figures from our small MRP polls and large MRP polls already included a separate age band for voters aged over 75. Our post-election analysis suggested including the higher age bracket did help the constituency-level large MRP differentiate between Conservative – Conservative versus Conservative – Reform voters in particular, and this is something we will continue doing into the future.

Ethnicity

Our traditional voting intention method outside London has not included quotas or weights on ethnicity. The small and large MRP approaches we adopted during the campaign did include ethnicity as a variable, using Black, Asian, Mixed, White & other categories as targets.

At this election ethnic minority respondents did appear to vote substantially differently to white respondents, in particular among Asian respondents. Labour’s vote fell by the largest amount among voters of Pakistani and Bangladeshi ethnicity, while voters of Indian ethnicity were most pro-Conservative. The specific underperformance of Labour among voters of Pakistani and Bangladeshi ethnicity is widely ascribed to the impact of Gaza as an issue being greater among Muslim voters. The reason for the strong Conservative performance among Indian voters is less clear, though may possibly be due to positive perceptions of Rishi Sunak among the community. Whatever the drivers, it is clear these two groups of voters behaved very differently.

	CON	LAB	LDEM	RefUK	GRN	OTH
Pre election voting intention (traditional method)						
Black	8	69	10	4	8	2
Indian	26	47	13	4	7	3
Pakistani/Bangladeshi	1	44	8	5	19	21
Mixed	7	57	17	6	10	1
Other	14	49	13	8	13	0
Change since 2019						
Black	-5	-5	+2	+4	+6	-1
Indian	-6	-9	+2	+4	+6	+2
Pakistani/Bangladeshi	-12	-31	+4	+5	+17	+16
Mixed	-10	-1	-3	+5	+8	+1
Other	-12	+1	-2	+6	+9	-2

Adding basic ethnicity weights to our traditional approach would in practice have made an impact, albeit a modest one. Without turnout, adding ethnicity to the targets would have increased Con support by 0.4%, reduced Labour by 0.3% and reduced Reform by 0.5%.

	CON	LAB	LDEM	RefUK	GRN	OTH
Basic weights (without turnout)	16.8	39	11.6	20.2	7.2	1.7
With Ethnicity added (without turnout)	17.2	38.7	11.5	19.7	7.2	2.3
Difference	0.4	-0.3	-0.1	-0.5	0	0.6

The small and large MRP methods already included ethnicity targets, but these grouped respondents into categories of Black, Asian, Mixed and White. Grouping respondents in this way is unlikely to have properly represented the radically different voting patterns that were evident in 2024 between areas with a large number of voters of Indian ethnicity and areas with a large number of voters of Pakistani and Bangladeshi ethnicity.

In future we intend to break this down into more detail and treat respondents of Indian ethnicity separately to those of Pakistani and Bangladeshi ethnicity.

Income and Social Grade

Comparing income data with data available from the ONS there is some prima facie evidence that our samples may be under-representing the most affluent respondents and overrepresenting poorer respondents. Straight comparisons are hindered by the poor availability of official data and the fact that around a quarter of our respondents will not provide household income.

An overstatement of low-income people seems unlikely to be a cause of an understatement of Tory support in 2024. At the aggregate level poorer respondents were more likely to vote Conservative and higher income respondents are more likely to vote Labour. This is, however, largely due to income and age interacting with one another and pensioners having low incomes. Within age bands, there is a more complex relationship, with higher incomes relating to stronger Conservative support and lower Reform UK support.

Higher incomes also relate to higher Labour support though, possibly related to a lower proportion of people saying they would not vote. Having too many low-income respondents would help explain our Conservative understatement and Reform UK overstatement, but increasing income across our sample would likely have made the turnout pattern and level of Labour support worse.

<i>Age & Household Income</i>	CON	LAB	LDEM	RefUK	GRN	WNV	DK
<i>Under 40 HI under 25k</i>	4%	33%	9%	11%	12%	18%	10%
<i>Under 40 HI 25-50k</i>	5%	40%	8%	11%	9%	13%	10%
<i>Under 40 HI 50k-70k</i>	6%	43%	11%	10%	8%	11%	7%
<i>Under 40 HI over 70k</i>	8%	44%	13%	9%	7%	6%	10%
<i>40-65 HI under 25k</i>	9%	27%	6%	18%	4%	16%	14%
<i>40-65 HI 25-50k</i>	13%	29%	8%	18%	4%	11%	12%
<i>40-65 HI 50k-70k</i>	11%	36%	9%	19%	4%	6%	11%
<i>40-65 HI over 70k</i>	15%	38%	13%	13%	4%	4%	11%
<i>65+ HI under 25k</i>	25%	21%	8%	21%	3%	8%	10%
<i>65+ HI 25-50k</i>	29%	20%	10%	18%	3%	6%	10%
<i>65+ HI 50k-70k</i>	29%	26%	12%	15%	2%	4%	10%
<i>65+ HI over 70k</i>	31%	24%	12%	13%	2%	3%	10%

While there are some signs that there may be issues with income, but it would be challenging to add as a weight or quota. Official data on income is patchy and often available only in equivalised format. Our own data is limited by respondent reluctance to answer questions on income. This means weighting by income is unlikely to be a viable option.

An alternative but associated route is socio-economic measures such as social grade. While this does not correlate directly to income, it has some relationship.

We have until now used the MRS Social Grade as a socio-economic measure, but this also suffers from growing problems with a lack of official or reliable data to base targets upon. In 2011 the MRS recoded the 2011 census results by approximated social grade in order to provide targets for the market research industry. In 2021 this proved impossible to do for older respondents, so the 2021 census approximated social grade covers only those under the age of 65, making it unsuitable as a source for quotas or targets for the whole adult population.

We are recommending instead switching to the use of NS-SEC categorisation – this has more robust statistical underpinning, and is available from the census. It also has more finely graded categories than traditional Social Grade, allowing us to ensure that the highest socio-economic category is not under-represented. Testing so far does not suggest that using NS-SEC would have brought voting intention in our sample any closer to the actual result at this election. NS-SEC would still provide more robust targets for quotas and weights given the absence of approximated

social grade from the 2021 census, and we will look to switch our samples to use NS-SEC in the future.

Our large MRP model did not include social grade due to the lack of accurate target data at the constituency level. Post-election analysis of MRP suggests that including allowances for social class (as defined by the approximated social grade measure) within the MRP process would have made little difference to our final estimations and projections. In the future we will explore the potential for incorporating NS-SEC into our MRP model.

Education

Education has been suggested as a contribution towards previous polling errors, specifically the over-representation of graduates in survey samples. Education has increasingly correlated with voting behaviour and turnout, even when accounting for age. It is possible that an overrepresentation of graduates would result in an overstatement of Labour support.

Examining our data suggests there may have been some over-representation of education due to the question wording used to collect data.

Our education quotas and weights are currently based upon a single question asking respondents their highest qualification (similar to the question used on the paper version of the census). The Census in 2021 changed the way they asked the question for respondents filling in the form online, asking it as a series of five separate questions. Testing the two different methods, they produce significantly different results. Some respondents who are currently coded as “other” using our existing question were found to be graduates using the step-by-step question. This suggests that in the past we may have been inadvertently over-representing graduates.

Testing by reweighting the data using new education weights suggests it would have made only minor differences using the old methodology – it would have decreased the share of Labour vote and decreased turnout, but would have increased support for Reform. Nevertheless, we will update our approach to collecting education from to maximise the accuracy of our panellist data.

MRP modelling decisions

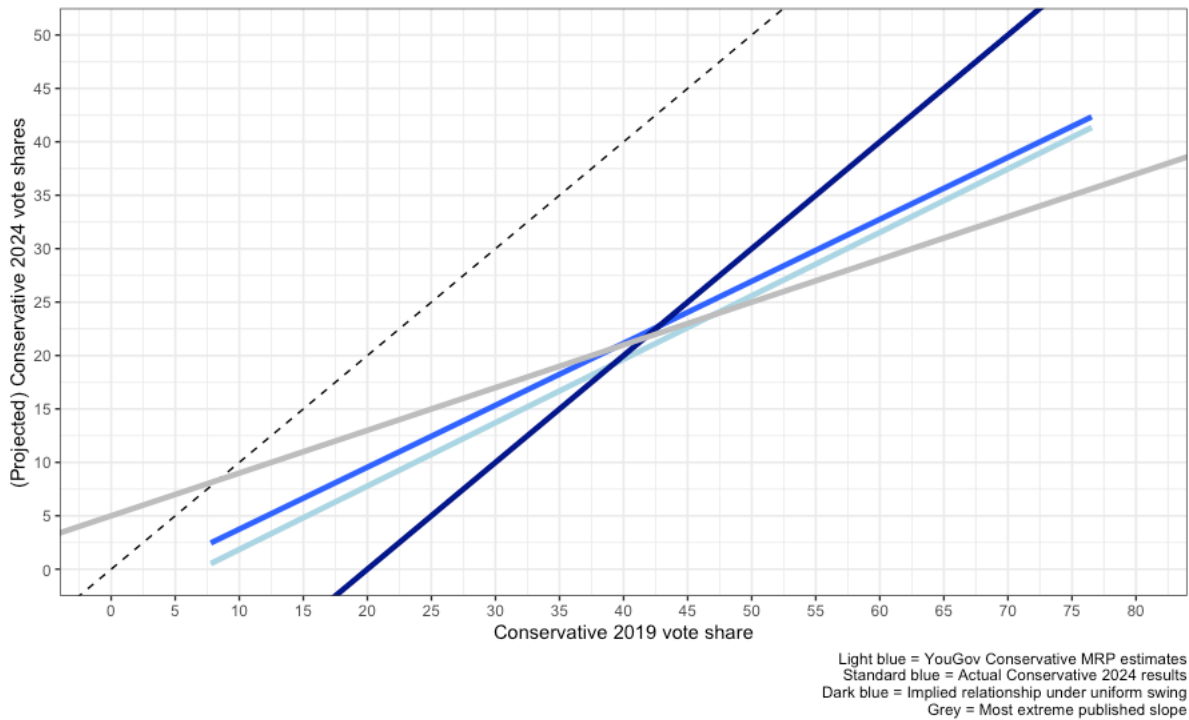
Failing to estimate 'the slope' of proportional decline

Prior to the election campaign there was discussion around MRP models showing a proportional drop in Conservative support at the seat level. The standard in British elections in the past has been for party support to fall uniformly, but MRP models were consistently suggesting that in 2024 the decline in Conservative vote would be overwhelmingly proportional.

During the campaign itself there were significant differences between MRP models regarding the precise strength of the slope between Conservative 2019 performance and their 2024 performance. Estimating this relationship wrongly could have been a serious source of error at the seat level for MRPs, with the Conservatives able to retain more seats in a scenario where their vote share decline was more uniform than in scenarios where it was more directly proportional.

Our post-election analysis suggests our constituency-level MRP model approximated the actual relationship extremely well, as the graph below shows. The mid blue line represents the actual relationship between Conservative 2019 and 2024 performance, showing the party clearly dropping most in its strongest seats – proportional swing. The light blue line shows the approximation YouGov gave this relationship in our final MRP, while the grey line shows the most extreme projection of this relationship from another company's MRP. The dark blue line shows what this relationship would have looked like in a purely uniform swing scenario.

The relationship between 2019 and 2024 Conservative seat shares
 Lines of best fit between Con 2019 shares, MRP shares, UNS-implied shares,
 and actual Conservative 2024 results



Given the proximity of the light blue line to the mid blue line, we can conclude that the MRP model correctly modelled the proportionality of the Conservative vote, and this was unlikely to have been a contributing factor to any MRP error.

Conclusions and recommendations

Polling error at the 2024 was, as with most past cases, a combination of factors. There is no single “silver bullet” explanation. Based upon the analysis of our traditional method we conclude that the main factors that contributed to polling error across the industry were likely to have been:

- A late swing against the Labour Party
- Poll samples containing too many people who were likely to vote, especially in low-turnout demographics, resulting in an overestimate of Labour support
- Past vote weighting being less effective due to the Brexit Party standing down in 2019, leading to an overestimate of Reform UK Party support and underestimate of Conservative support

YouGov’s MRP constituency model performed strongly, as did the small MRP model that we used for regular voting intention figures during the campaign. The model itself appears to have performed very strongly in translating levels of support into seats, but in terms of projected vote share it still overstated Labour support. We believe this was mainly, but not exclusively, down to a combination of three factors:

- A late swing against the Labour Party, which our sampling and modelling approach did not properly capture
- Inclusion of self-assessed likelihood to vote biasing voter-group turnout probabilities in an incorrect direction, as opposed to relying wholly upon modelled-demographic factors
- Our raw sample data did not quite contain the correct balance of Conservative loyalists versus Conservative defectors

A smaller factor may also have been insufficient granularity in our ethnicity modelling, where treating voters with Indian ethnicity and Pakistani/Bangladeshi ethnicity separately would likely have better modelled contrasting trends amongst Asian voters, better distinguishing the drop in support for Labour among Muslim voters and the comparatively strong performance of the Conservative Party among voters of Indian ethnicity.

Recommendations

YouGov made a significant switch to our polling methodology at the start of the election campaign, moving from our traditional approach to an approach using a smaller version of our MRP model. This approach was significantly more accurate than our traditional approach would have been.

The reasons for this improved performance were largely because it allowed us to better model turnout using demographics and to better reflect don’t knows behaviour based on how respondents

of the same demographics behave. It also used the “constituency prompted” version of our voting intention question, which performed slightly better than the traditional question.

- For voting intention polls, we will continue to use the MRP method that we adopted during the election campaign.

While our large MRP poll was the most accurate, it did still overstate the level of Labour support. We intend to make a number of minor adjustments to improve its performance in the future, these will also be reflected in the smaller MRP model we use for regular voting intention polls.

- Improving the ethnicity modelling in our MRP to model Indian and Pakistani/Bangladeshi respondents separately.
- Basing our turnout modelling wholly upon demographic modelling rather than self-reported likelihood to vote.
- We will look into approaches to deal with the problem of late swing in large MRP datasets that are collected over a period of time. Our analysis so far suggests that addressing this problem using model-based solutions such as giving more weight to more recent respondents would have been counter-productive, as “fast respondents” can themselves be different. Solutions are more likely to be based on the scale and timing of our fieldwork.
- Further, we will use improvements made in the sampling process for standard polling to help correct for imbalances in the types of Conservative voters in our MRP samples.

In the process of exploring potential sources of error we also found some areas that probably did not contribute towards the error at this particular election, but which nevertheless could be improved. These include how we approach social grade and education.

For our non-voting intention polls we will also make some technical, but substantial changes in to our samples to address these issues. These will be rolled out during 2025.

- The question wording we use to collect education data will be changed to reflect the approach used in the online version of the 2021 census, moving from a single multi-code question to a step-by-step question, something that in practice will reduce an overrepresentation of graduates.
- For political polling we will move from Social Grade weighting (the traditional “ABC1C2DE” breakdown), which suffers from a lack of robust target data, to using the ONS NS-SEC socio-economic categorisation.
- We will also incorporate the updated education data into our MRP model and will explore the possibility of incorporating NS-SEC data into the model.