




Bradshaw Advisory's Election Prediction Model

What is it? How does it work? What does it show?



Introduction

There is always massive interest in the likely results of elections, especially in times such as today - over the past year, the UK has seen a war in Europe, economic crisis, mass industrial action, multiple prime ministers and Matt Hancock's run to become a reality TV star. With all this turmoil, commentators and strategists are watching voting intention polls with bated breath.

When it comes to giving an accurate view of the likely outcome of an election, polls of the voting intention of the public at large are not the most sophisticated of tools. After all, not all votes are equal in the UK's electoral system.

Many votes in safe seats have little, if any relevance for electoral outcomes, whereas just a small number of voter decisions in the most hotly contested seats can have dramatic impacts. It's the behaviour of constituencies and not individual voters which will determine electoral outcomes.

Bradshaw Advisory's Election Prediction Model estimates the probability that parliamentary constituencies would change hands were an election to be called today and models the likely overall results of that election. The model is updated fortnightly, quickly responding to the fast changing political environment, allowing users to track changes in electoral outcomes over time.

The model takes into account voting intention data across a number of different pollsters, as well as historical election data going back to the 1980s. This rich information contributes to the model's reliability, with it correctly predicting 92% of constituency results over the past three general elections.

Bradshaw Advisory's wealth of experience in Westminster politics is part and parcel of the model, with it being built with the underlying shape of British politics in mind. Results change to reflect shifting constituency boundaries, by-elections and changing voter behaviours around elections.

The rest of this paper will first present some of the headline findings of our modelling work, as well as outlining the key uses of the Election Prediction Model. A detailed analysis of the results and what they mean for the UK's electoral landscape can be found on page 6. An overview of the model's methodology and an assessment of its reliability can be found on page 13, while a detailed technical methodology can be found in the appendix.



The headlines as of May 2023

The final quarter of 2022 saw one of the biggest swings in voter intention in the UK's history. Many pollsters predicted that, were an election to be called, a Labour landslide would be a foregone conclusion. Our modelling calls this into question. While at the height of the instability, peaking in late October our model did predict a reasonable working majority of 343 seats (± 13 seats), this is still a far cry from the forecasts of 400+ seat majorities pointed to elsewhere.

As the polls have begun to close slightly in the run up to May, we see that were an election to be held in the middle of the month Labour would just lose out on a majority. As of the 8th of May they would be set to gain 313 seats (± 13 seats), the Conservatives to take away 229 (± 9), the Liberal Democrats 39 (± 2) and the SNP to gain 46 (± 2).

This comes off the back of significant losses for the Conservatives at the local elections, where they saw 1063 councillors and control of 48 councils leave their grips. While these figures have been touted by Keir Starmer as signs of Labour being on track to win a majority at the next election, it is worth noting here that Labour's performance is no better than it was in last year's local elections even after a year of complete political turmoil. Many of the seats lost by the Conservatives in the locals went to Liberal Democrat challengers, something closely reflected in our modelling, with the Lib Dems set to take a number of seats in the commuter towns of the South East.


In fact, based on the projected national vote shares stemming from the local elections, calculations by Sir John Curtice would suggest that were an election held now, we would expect to see Labour taking 312 seats, the Conservatives taking 226 seats and the Liberal Democrats taking 40 - just within a hair's breadth of Bradshaw's predictions.

Such results today make the results of the actual election set to be held in late 2024 nigh on impossible to predict, hammering home that the UK's flirtation with political instability and uncertainty is far from over.

This prediction of a hung parliament, even today, may surprise some readers, with Labour currently coming 15 points ahead of the Conservatives in the polls. As we should all know, however, a strong showing amongst the electorate at large does not necessarily translate to electoral success - elections are decided by winning constituencies, not voters.

We have identified four main factors preventing the Labour party from translating their poll lead into likely electoral success:

- The SNP's continued dominance in Scotland and the high degree of competition between unionist parties.
- The likelihood of the Liberal Democrats winning over disaffected Conservative voters, especially in Southern England.
- Labour support deepening in areas of existing strength, rather than widening to new areas.



- Boundary changes, which may cap the number of seats Labour take home in Wales.

Most of these problems facing the Labour party are not new or surprising. But the fact that they have the potential to disrupt the party's ability to win a majority, even with a 15+ point lead in the polls should be of keen interest to strategists across Westminster and to the businesses and constituents which stand to be affected by the results of the election.

What can you use the model for?

Bradshaw's Election Prediction Model can be used as a means of closely tracking the development of electoral politics in Great Britain and can give fortnightly updates to users of how elections are likely to swing. While we do provide a 'central estimate,' indicating the most likely outcome based on our analysis, users are able to adapt the results to see the electoral outcomes swing more strongly to Labour or the Conservatives, based on their own hunches regarding the electoral landscape. Users can also look at constituencies of their choice to see how likely they are to change hands.

While a point of interest for anyone engaged with British politics, the model and its outputs can provide important insight as a tool for political strategists. It provides clear insight into which areas are the most closely contested at any point in time, providing a view of where incumbent parties will want to defend most strongly or where challengers will want to send resource to tip the scales in their favour.


In addition to the core outputs from the model, we have also provided analyses to help identify the top performing MPs from each party in each region, looking to where particular MPs have dramatically outperformed their parties in their given regions either during the last general election or in by-elections. This can provide both insight into which strategies and personalities can be leveraged towards electoral success in the future, as well as highlighting where potential anomalies may lie relative to the model's core results.

We also provide a unique overview of 'voter spread' in each region. With our new Voter Spread Index, strategists can see which regions see a significant clustering of their voters in particular seats, or where they are spread more broadly. A key risk to some parties is spending resource to deepen support in constituencies where they already have a high number of voters, but missing out on winning over the electorate in swing seats.

Why Bradshaw's approach?

This is by no means the first foray into election forecasting. At present, the most common form of model used to predict electoral outcomes in the UK are multiple regression post-stratification models (or MRP models for short). These models work by taking detailed polling data, with a large sample of individuals' voting intention recorded against a set of data regarding their demographic, social and economic characteristics.

These figures can give a sense of how likely voters of a given set of characteristics are likely to vote for a given party. This information is then matched to the demographic, social and economic characteristics



of constituencies across the UK, to give a sense of which parties will perform better or worse in each area.

These models have proven to be reliable in predicting electoral outcomes in some cases, with a YouGov MRP model predicting around 93% of the constituency-level results correctly in the 2017 election. They are, however, incredibly data intensive.

Attempting to estimate the vote shares of multiple parties across hundreds of constituencies on the one hand and requiring expensive, detailed and large-sample surveys on the other. The fact that these models are so demanding can make their reliability sensitive to the amount of resource put behind them. If a sample is slightly too small or a few marks off being representative of the true voting population, or, if the identification strategy misses out key voter characteristics, results can begin to lose accuracy very quickly.


Our approach, to simply model the probability of a given seat changing hands, based on national-level polling data and the historical electoral behaviour of each constituency dramatically reduces the data-intensity and sensitivity of the process to slight lapses in polling quality or other factors. The significantly reduced data requirement of the approach also means that it can be updated more quickly in response to changes on the ground, allowing for a close and up-to-date following of British electoral politics without the need for masses of investment.

In predicting the probability of seats changing hands, the model does not require an overview of exactly how a given party is performing amongst specific demographic groups at any point in time. It can take the last election's results as a given and simply ask 'have conditions for the incumbent party in this constituency changed enough for the seat to change hands?' In a political system where seats have a strong tendency to remain under the control of the same party for multiple election periods, investigating the seat-by-seat performance of individual parties is liable to be overkill. In fact, were you to assume that all the seats from the 2017 general election remained in the same hands for the 2019 general election, you would get 87% of your guesses correct.

What does the model tell us about the electoral landscape in early 2023?

As noted above, the UK was home to three prime ministers, an economic crisis and some of the most extreme industrial action in decades in 2022. This turmoil has been reflected across opinion polls, with pollsters tracking confidence in the UK's political institutions to some of their lowest levels since records began.

Perhaps most stark has been the lead gained by the Labour party when it comes to voting intention. They began the year with a modest lead over the Conservatives, with 39% of those interviewed saying they'd vote Labour, compared to 33% saying they'd vote Conservative. By October, that changed to 52% gunning for Labour and only 23% saying they'd still vote Conservative.



Many might think that this would be hinting at a landslide victory for the Labour party were a snap election called, but this is unlikely to materialise. The question why exactly the current 20+ point poll lead is unlikely to turn into a strong majority for Labour provides core insight into the real landscape of electoral politics in the UK and will be the main focus of this analysis.

As can be seen from figure 1, our modelling suggests that the Labour party would indeed have broken away from the Conservatives at around the time of Boris Johnson's departure, but that their lead would only serve to provide a fairly narrow majority. Labour's performance was at its strongest in late October, during the aftermath of the announcement of the government's Growth Plan, where they would be set to gain 343 (± 14) seats - 17 more than the 326 necessary for a majority. Over this same period, the Conservatives would be set to lose 197 seats compared to their 2019 majority. Following the peak of Conservative turmoil in late September, the projected Labour hovered on the borderline of a majority, before falling further in March.

Figure 1: modelled seat winnings by party, January 2022 - May 2023

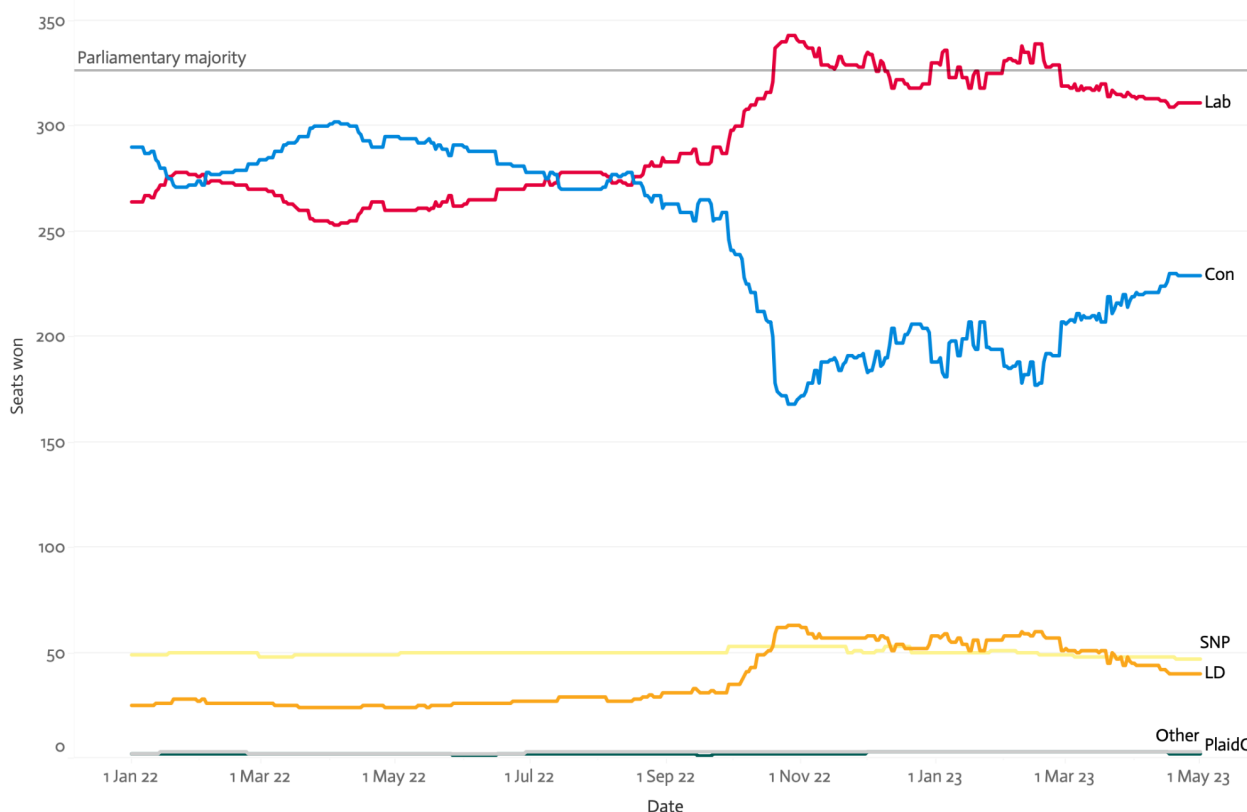


Figure 2 below shows the projected results of an election held on 1st of May 2023 by constituency. Figure 3 shows those results in Conservative-held areas. As can be seen, many of the gains made by Labour are in the Red-Wall areas, such as Penistone and Stocksbridge in

Yorkshire and Bishop Auckland in the North East, taken by the Conservatives in the 2019 general election. Seats across North Wales, also lost in 2019 would also be set to go back into Labour hands.

As anyone reading this will know, metropolitan areas are often strongholds of Labour support. Our modelling suggests that support across cities and larger towns would strengthen were a snap election called, with Labour set to take constituencies like Lincoln, Gloucester, York Outer and Peterborough from Conservative rule.

Figure 2: modelled results by constituency, 1st May 2023

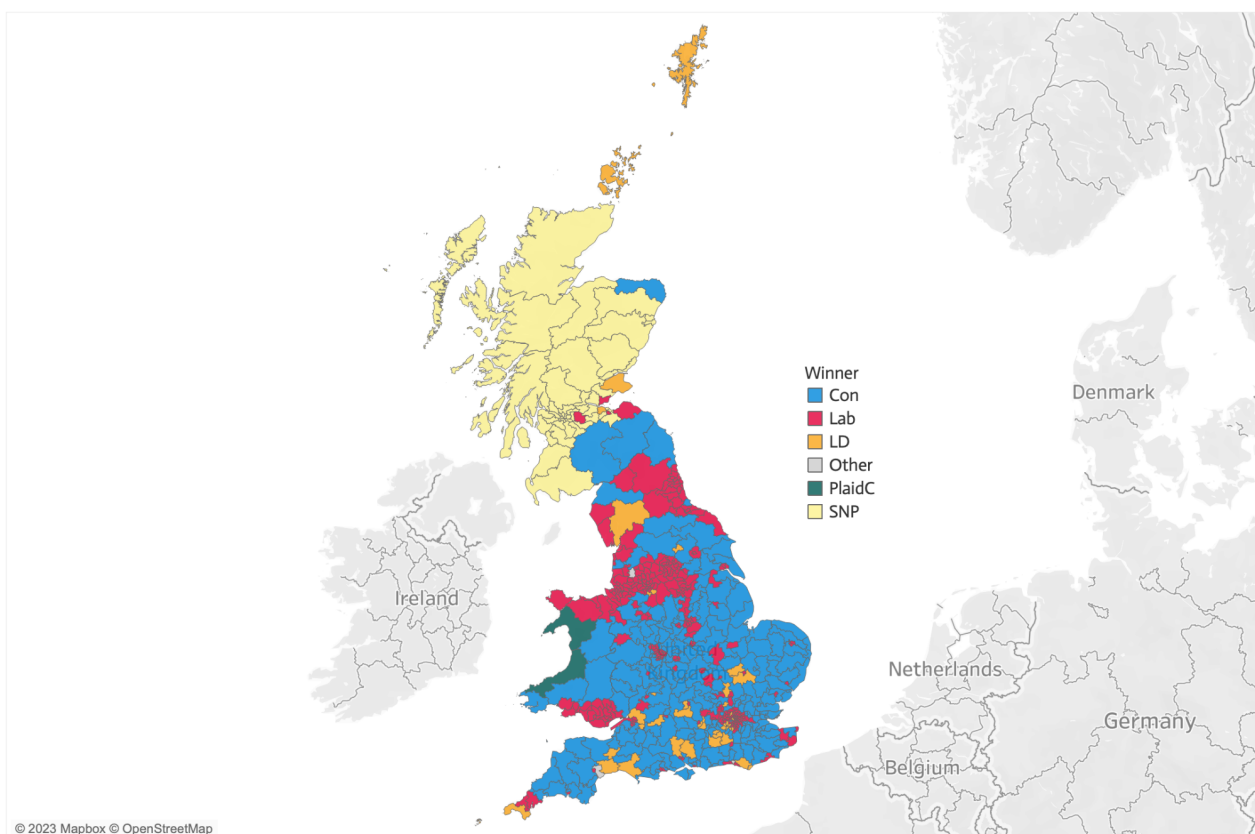
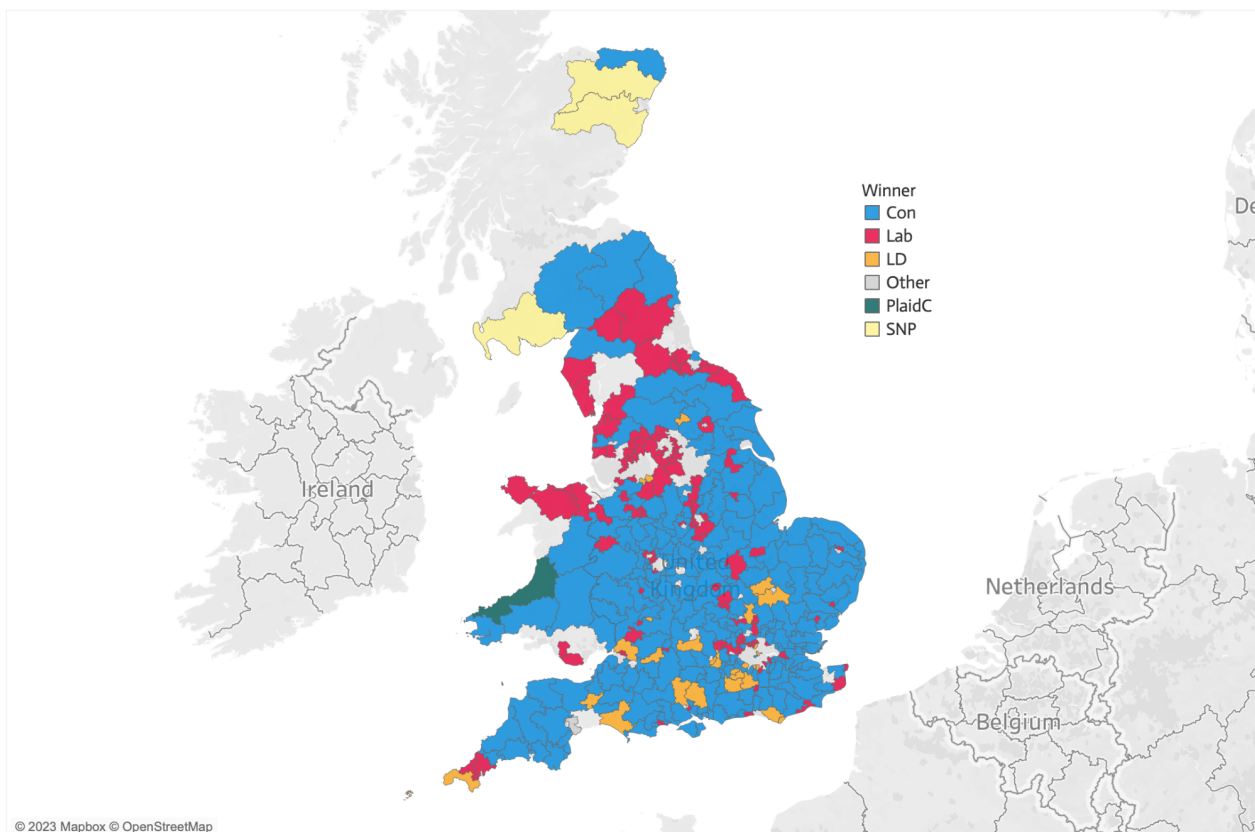



Figure 3: modelled results in Conservative-held areas by constituency, 1st May 2023



In essence, the Labour party is set to strengthen its hold in areas in which it traditionally does well. These areas of traditional strength are not, however, sufficient to deliver a secure majority that might be implied by the 20+ point poll lead the party has enjoyed over the final quarter of 2022 and early 2023. There are a few key reasons for this: the SNP's dominance in Scotland, the ability of the Liberal Democrats to win over Conservative voters, Labour's vote being concentrated across too few constituencies and the proposed changes to electoral boundaries.



Scotland

There has not been a time in living memory when a Labour majority has not been accompanied by significant support in Scotland - during Tony Blair's run to victory in 1997, the party gained 56 seats North of the border, while in the 2005 election, Scottish voters handed the party 41 seats. The Labour party currently holds only 1 Westminster constituency seat in Scotland, leaving them in fourth place in the country when it comes to seats, despite currently coming in second in the Scottish polls. Our modelling suggests that even in the ongoing turmoil amongst the SNP's ranks, Labour would only stand to increase their Scottish total to 4 seats.

The SNP holds 45 of the seats in Scotland, gaining 45% of the vote overall in the 2019 election. Support for the SNP is fairly evenly spread across Scottish constituencies. In 50 of the 59 constituencies in Scotland, the SNP captured between 40 and 54% of the vote, dropping below 40% in only 9 areas.


A key factor differentiating a number of the seats in Scotland is that many of them are 3-horse races, with 2 parties vying for the unionist vote, all directly competing against each other and the SNP. For example, in Cumbernauld Kilsyth & Kirkintilloch East, the SNP gained 52.9% of the vote in 2019, compared to 24.5% for Labour, 16.2% for the Conservatives and 6.6% for the Liberal Democrats. Any real strategy for unionist parties wishing to eat away at the SNP's hold must concentrate on where their individual performance is strongest, diverting resources away from areas where competition with other unionist parties is too fierce for any individual party to challenge the SNP's base.

Another factor to consider here is that Westminster elections are separate from those used to select representatives within Holyrood. This is a particularly important factor when considering the potential effects of the scandals currently surrounding the SNP's top brass. It is entirely possible for Scottish voters to punish the SNP in Holyrood elections, but continue their support during Westminster elections, where the battle lines between the pro and anti-independence vote is likely clearer. This is already being reflected in polling results, with the SNP's support dropping swiftly in Holyrood voting intention polls, but remaining relatively static in Westminster polls.

The Liberal Democrats

The second core factor to consider here is the support for the Liberal Democrats. They only won 11 (3% of the total) seats in the 2019 election, but gained 11.6% of the popular vote. Of the Conservative seats our modelling has highlighted as being likely (>50% chance of changing hands) to change hands, the Liberal Democrats were the main contenders to the Conservatives in 25% of cases.

Much of the narrative around the declining popularity of the Conservative vote has generally been, quite naively, assumed to signal a near 1-for-1 gain for the Labour party. This would tacitly assume that the Labour party is the second choice party for Conservative voters, over the Liberal Democrats, Reform UK or the prospect of simply not voting at all. As noted, the Liberal Democrats regularly appear in strong Conservative seats as being the second party of choice, with this pattern being especially strong around the home counties. Our model suggests that many of these areas would be won by the




Liberal Democrats in the event of a snap election, with them gaining 58 (± 2) seats compared to their position at the last general election, leaving them in a similar position to that of the 1997 election.

Voter concentration

As noted above, our analysis suggests that the Labour party will generally strengthen its hold in areas it has historically done well in. In many cases, this process will work to send seats which have been lost in recent years back into Labour hands. In many cases, however, the deepening of support in areas of historical strength will do little to change overall election results - new territory must be won.

As an indication of the potential threat posed to any party by this problem, we have constructed a simple Voter Spread Index for each party. The Index simply asks the question 'if I were to draw two voters at random from those who voted e.g. Conservative, what is the probability that they would be registered to different seats?' The higher this number is, the less acute the issue of support concentrated in areas of existing strength is. The results for Labour and the Conservatives over the past 4 elections can be seen in figure 4. As shown, Labour has seen an increasing degree of concentration over time.

Figure 4: Voter Spread Index





Boundary changes

A final point to note here is that of the changes to constituency boundaries set to be brought into force in the summer of 2023, particularly those which are due to come into force in Wales. Our model has been calibrated to fit these updated boundaries.

Wales, like Scotland has been in previous years, has been a significant help to the Labour party, with it consistently ranking as the country's most popular party. Wales is currently home to 40 constituencies, but that is set to fall to 32 following the proposed boundary changes.

For the most part, this will not change the landscape of politics within Wales too strongly, our modelling suggests that the proportion of seats taken by each party would be fairly similar if the model were calibrated to the new boundaries or the ones currently in force. The change may, however, be of significance to the political landscape of the country overall, given the paper thin margins currently predicted for a Labour majority, 8 seats either way can very easily mean the difference between a clean transition of power and a hung parliament.


Conclusions

Bradshaw Advisory's Election Prediction Model gives constituency-by-constituency predictions of the results of snap general elections, with those results updated fortnightly. The results of the model show that over the course of 2022, while we do see the Labour party clearly pulling ahead in terms of the seats it is likely to win, those winnings may be significantly more limited than many are predicting.

A good analogy for the shape of results here would be '1997, but with the SNP.' Where we see similar support for Labour (as per 1997) across their strongholds in England and Wales, the Liberal Democrats returning to the strength they saw prior to the coalition government, but with the landscape of Scottish politics still divided by the independence vote. In fact, were we to take the results of 1997 and hand all but one of Labour's Scottish seats to the SNP, the results would be remarkably similar to what Bradshaw's model predicts over the latter part of 2022 and early 2023.

Such a pattern and the inherent difficulty in maintaining a 20+ point poll lead for **2 years**, would imply that we may well be in for a hung parliament come January 2025. The only real chance of seeing a parliamentary majority being called would be under an incredibly successful campaign for the Labour party, which would be able to overcome:

- The SNP's dominance in Scotland and the high degree of competition between unionist parties.
- The likelihood of the Liberal Democrats winning over disaffected Conservative voters, especially in Southern England.
- Labour support deepening in areas of existing strength.
- Boundary changes, which may cap Labour winnings in Wales.



For the most part, these factors should not be a surprise to anyone. The new insight here is that even with a 15+ point poll lead for Labour, the influence of these factors is strong enough to generate a hung parliament, something that should be taken into close consideration by strategists in all the UK's parties.

How does the model work?

In broad strokes, Bradshaw Advisory's Election Prediction Model provides estimates of the likelihood of parliamentary constituencies changing hands were a general election called immediately. The model takes into account the national-level polling performance of the parties in each seat, as well as data regarding the past behaviour of that seat, for example, how many times in the last 4 elections has the given seat changed hands? These figures help build a picture of how likely the seat is to change hands.

Once each seat has had its change probability estimated, those figures are run through a number of scenarios which provide us with a picture of how a general election would play out at different points in time. We provide a 'central estimate' of what is liable to happen were an election called tomorrow, as well as a range of other results which allow users to 'swing' the election slightly, depending on their own views of how things could change in the near future.

A full technical methodology can be found in the Appendix for those who want a detailed look at how exactly the model's results are derived. For now it is worth seeing how well the model performs when it comes to predicting previous elections.

How accurate is it?

We have tested the model's ability to predict the seat changes of the past 3 general elections (2015, 2017 and 2019) to test the model's effectiveness. These elections were selected in particular, given that their underlying electoral landscapes are most similar to the ones we face today.

Model performance across these three elections is remarkably consistent, with 92% of seat changes predicted correctly across each election. The vast majority of seats predicted correctly were across England and Wales across all elections and while the model's performance in Scotland does lag slightly, it still gets a majority of its predictions correct. You can see its constituency-by-constituency performance in the graphics overleaf.

Figure 5: 2015 general election model performance

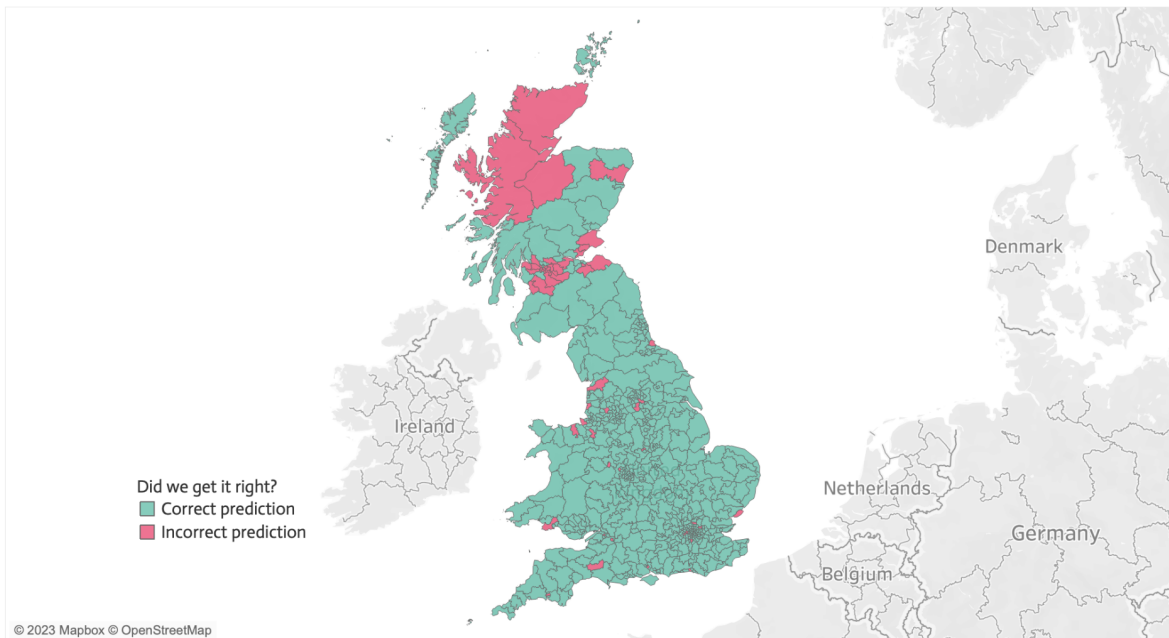


Figure 6: 2017 general election model performance

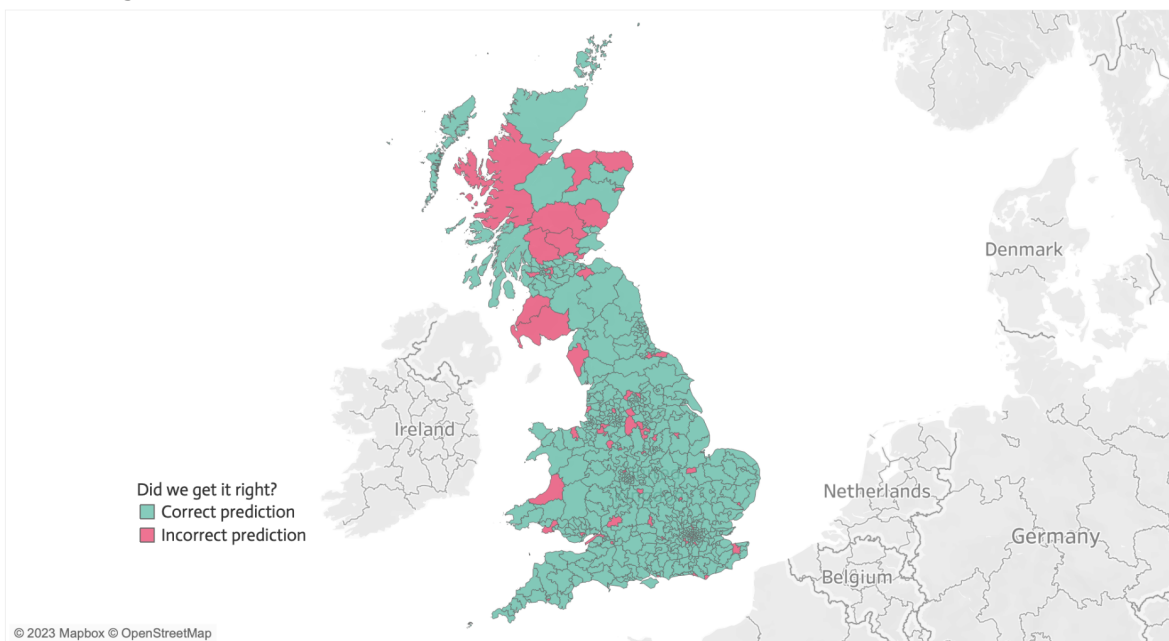
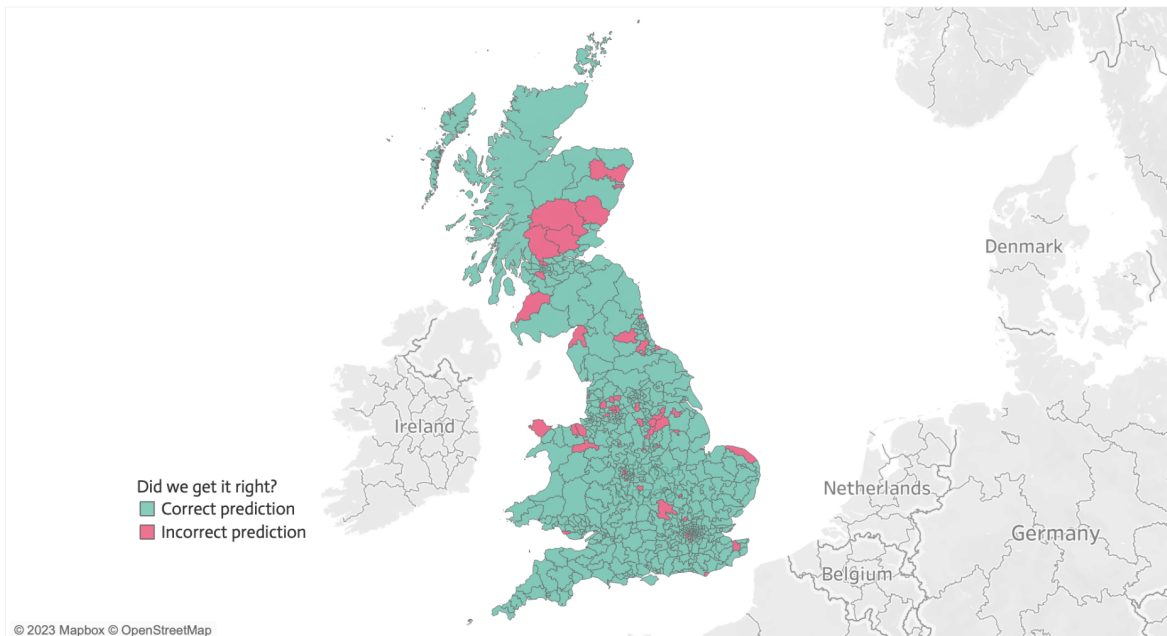


Figure 7: 2019 general election performance




So what would the model have predicted as far as overall election results go?

In 2015 we would have seen a prediction of a conservative minority government, with the Tories scooping up 312 seats, compared to 257 for Labour and 45 for the Liberal Democrats. In reality, we saw a swing to a Conservative majority of 330. While the model did correctly predict a swing in the Tories' favour vis-a-vis the 2010 election, it failed to capture its full scale, as well as failing to predict the Liberal Democrats fall from grace. They gained 57 seats in the election prior, with that number falling to just 8 in 2015. Also not captured was the SNP's meteoric rise in the independence referendum of 2014, though their significantly increased popularity is reflected in the predictions of the 2017 election.

Table 1: 2015 general election prediction vs reality

Party	Predicted seats	Actual seats
Conservatives	312	330
Labour	257	232
Liberal Democrats	45	8
SNP	7	56
Plaid Cymru	3	3
Other	8	3



Predictions for the 2017 election see almost exactly correct answers for Labour, the Liberal Democrats and Plaid Cymru. Its performance for the Conservatives and SNP remains close, with it lying only 17 seats out for the former and 16 the latter, its lack of ability to reflect the Tories' downswing in popularity generates an incorrect prediction of a Conservative majority, when a minority government was formed.


Table 2: 2017 general election prediction vs reality

Party	Predicted seats	Actual seats
Conservatives	334	317
Labour	234	232
Liberal Democrats	8	12
SNP	51	35
Plaid Cymru	3	4
Other	2	

The model correctly predicts the overall result of the 2019 general election, with the Tories gaining a majority. It does not, however, fully capture the extent of the victory, nor Labour's significant drop in performance, with them falling to only holding 202 seats.

Table 3: 2019 general election prediction vs reality

Party	Predicted seats	Actual seats
Conservatives	328	365
Labour	249	202
Liberal Democrats	14	11
SNP	33	48
Plaid Cymru	3	4
Other	4	2




Why are these results on the 'conservative' side?

On a seat-by-seat basis, the model performs well, with the results across 92% of seat changes predicted correctly. The results are, however, somewhat 'sticky.' They tend towards assuming that the results of an election will be slightly more similar to that of the prior election than they are in reality and as such underperform slightly when it comes to capturing the full electoral picture in periods when political conditions for a given party are changing significantly.

This is not particularly surprising. The model is built with significant attention paid to local conditions in constituencies. A common criticism of MRP modelling is that its approach is too 'top down,' relying too much on national trends and ignoring local peculiarities. Bradshaw's approach is more blended, taking into account both local and national variation. As noted throughout this paper, the general tendency amongst British constituencies, with most staying in the control of a single party for multiple elections.

A model weighted more closely to the behaviour of the 'average constituency' rather than to the more volatile swings in national poll results, will of course tend to behave in a way similar to those local constituencies. This points to a more fundamental statistical problem underlying this sort of modelling - those constituencies most likely to change hands tend to lie at the edges of distributions, not their centres, and as such models parameterised to perform well for average constituencies may miss some details. While an easy problem to overcome were we predicting some continuous outcome, whether or not a seat changes hands or not is a binary outcome and thus techniques such as quantile regression are unavailable to us.



Appendix: Detailed Methodology

Sources

The model runs on figures provided by Open Council Data, concerning the features and electoral behaviour of parliamentary constituencies across Great Britain, as well as the country-level (England, Wales and Scotland) results of the popular vote at each general election.

When simulating what the results of an election held today would look like, the popular vote data is replaced with polling data, gathered from: BMG, Deltapoll, Findoutnow, Ipsos Mori, Kantar, Omnisis, Opinium, Panelbase, People Polling, Redfield & Wilton, Savanta, Survation, Techne and Yougov.

Underlying model

The model underlying our estimates is a simple Probit regression with a binary dependent variable, Y_{cdpe} , which takes a value of one if constituency c , in nation d (includes England, Wales and Scotland) with an incumbent of party p , saw a change in the Westminster party controlling it in election e , and zero otherwise. Y_{cdpe} is regressed against X_{cdpe} and N_{dpe} where X_{cdpe} is a vector of independent variables which vary by constituency, nation, party and election and N_{dpe} is a vector of independent variables, varying by party and election. The model is run across data from the 1983 election to that in 2019.


$$Y_{cdpe} = \beta_x X_{cdpe} + \beta_n N_{dpe} + \epsilon_{cpe}$$

Included under X_{cdpe} are:

- MAJORITY: The difference in vote share gained by party p in constituency c in election $e-1$ compared to that of the party which came in second place.
- CHANGES: The number of times constituency c has seen its ruling party change in the 4 elections running up to election e .
- PARTY_YEARS: The number of election cycles of the past 4, where party p has been the incumbent in constituency c .
- PARTY_CHANGES: An interaction term between PARTY_YEARS and CHANGES.

Included under N_{dpe} are:

- POP_VOTE: the share of the popular vote gained by party p in nation d in election e .
- VOTE_CHANGE: the change in the share in the popular vote gained by party p in nation d in election e compared to election $e-1$.
- POWER: a binary variable denoting whether party p is currently in government in Westminster.



The model is parametrised slightly differently for Scottish constituencies, with their results predicted only using MAJORITY and VOTE_CHANGE. When tested, this highly simplified form of the model more accurately predicted Scottish election results.

Making predictions

To predict how an election will go were it to be held on any given date, the parameters derived from the model are fit to the latest data available for each constituency, while the popular vote data underlying POP_VOTE and VOTE_CHANGE are replaced, as noted above by polling data. After this change, POP_VOTE becomes the estimated proportion of the electorate, from nation d , surveyed in period t , who intend to vote for party p , while VOTE_CHANGE becomes the difference between the estimated proportion of the electorate, from nation d , surveyed in period t , who intend to vote for party p and party p 's share of the popular vote in the 2019 general election. Further details on the derivation of polling results are set out below.

Once the model is fit to the data, we are left with a set of values denoting the probability that constituency c will see a change in its ruling party during period t . In order to come to some final predictions about which seats would in fact change, we need to set a cut-off threshold for the set of probabilities, above which seats are assumed to change hands and below which the incumbent party maintains its control. A further assumption is made that if a seat does change hands, its control is shifted to the party which came in second the last time it was contested - *this may be the 2019 election or the by-elections held since then*.

To produce an estimate of the 'best' cutoff threshold, we use the model to simulate the past 3 general elections, with simulations run across various thresholds. The results of each simulation is compared to the actual results which maximises the model's accuracy in predicting the elections' true results.

The 'best guess' threshold derived from this method was 72%, which has some intuitive validity; we are essentially saying that of those constituencies which have above a 50% chance of changing hands, just over half do.

The poll of polls

To produce an estimate of each party's polling performance for each moment in time, we take a poll of the available polls taken across the UK/GB, as well as in Wales and Scotland individually. Estimates of English performance are derived from the UK/GB polls (it's very rare to find an England-only sample) and estimates for Wales and Scotland are solely based on polling within their borders.

To come up with an estimate of the polling performance of each party in each nation, at any given moment in time, we produce a weighted average of polls, with weights consisting of two components. The first is a time decay term, which decays exponentially from a value of 1 for a poll taken on our day of interest to 0.01 on the date 3 months prior. The second term is an estimator of pollster variance; the



mean (for each pollster) of the residuals of a regression of the variance in polling across pollsters, countries and parties on a set of dummy variables for each party, country and interactions between the two.

Finally, it is worth noting that in many cases, just prior to a general election, the polling between major parties has a tendency to narrow. To reflect this, we adjust all the polling slightly across nations before using it for final predictions, closing the gap between Labour and the Conservatives in England and Wales and between Labour and the SNP in Scotland by 5%.

Changing boundaries

A major challenge in undertaking this kind of modelling is the fact that constituency boundaries change dramatically over time. Were we to only work with those constituencies which existed in 2019, the resultant samples would likely be too small to provide anything workable. To overcome this problem, shapefiles of the electoral boundaries in each year were overlaid upon each other, to provide an overview of how each boundary set compares and which constituencies from 2001 would make up the area taken up by a 2019 constituency boundary.

Running under the notably strong assumptions that populations and voting preferences are uniformly distributed within constituencies, the overlaid boundaries were used to create a set of 'synthetic' election results for each year, such that all boundaries could be aligned with one another and thus a large enough sample of electoral results could be utilised.