

WORK IN PROGRESS – COMMENTS WELCOME

25th October 2013

Predictable and unpredictable changes in party support: A method for long-range daily election forecasting from opinion polls

Stephen D. Fisher, University of Oxford
stephen.fisher@trinity.ox.ac.uk

Abstract

This paper outlines a method for forecasting the next British general election on a daily basis from twenty months prior using opinion polls. The historical relationship between the polls and the votes suggest that government support rises substantially in the run up to elections, that Conservatives outperform but Labour underperform relative to the polls and parties generally recover from low points or decline from high ones. So despite the Conservatives trailing in the polls by seven points in early October 2013, the method forecasts a healthy Conservative lead at the 2015 election. Approximate prediction intervals and probabilities for key events are also generated. The Conservatives are estimated to have an 85% chance of being the largest party and 54% chance of an overall majority. The estimated probability of a hung parliament (currently 28%) is instructive for understanding the operation of the electoral system.

Introduction

There is considerable public interest in the likely outcome of elections even well in advance. If, as Harold Wilson famously claimed, a week is a long time in politics, then it would seem like a hopeless task to try to say what might happen in an election more than a year away. But we know that in Britain opinion polls are informative of future election results quite some time in advance of the election (Wlezien et al. 2013). Not least this may be because most voters are relatively stable in their party preference, the parties do not change much within electoral cycles and they are unlikely to change their leaders or policies within the last couple of years before the election. So some of the key factors that affect election outcomes are in place well before polling day. Thus it is reasonable to consider what the opinion polls tell us about the likely outcome of an election even over a year in advance.

Running regressions of election results on opinion polls at different times before the election enables us to predict in what direction public opinion is likely to change. Most importantly the two main parties are more likely to increase support in the run up to the election when in government rather than in opposition. In addition the Conservatives typically outperform and Labour under perform the opinion polls. Also if a party is doing unusually well in the polls some way before the election their support is likely to drop and *vice versa*. These three factors all help explain why the Conservatives are expected to do a lot better at the election and Labour worse than the polls currently suggest.

Public opinion changes in unpredictable as well as more systematic ways. Although there will be events that change public opinion between now and the election, past election cycles are informative about how much opinion is likely to change. Having

some sense of the uncertainty in the outcome is perhaps more interesting than the point estimate. It tells us the extent to which the outcome looks like a foregone conclusion or whether history suggests there is everything to play for. So while it is most likely, according to current polls and the method here, that the Conservatives will be the largest party at the next election, there is still a decent chance it could be Labour. The probabilities also are revealing about the operation of the electoral system, showing a 28% chance of a hung parliament and only a 54% chance of a Conservative majority even those Labour would more clearly be set for a majority if the Conservative and Labour forecasts were reversed because of the well known bias (Johnston et al. 2012).

So this paper develops a methodology for forecasting election outcomes from opinion polls on a day-to-day basis up to twenty months¹ in advance of an election. The paper first considers the case for long range forecasting from polls relative to previous election forecasting technologies and then the general lessons from looking at how polls twenty months out corresponded with election results in previous cycles. I then describe the votes-on-polls regressions that are the main bases for the forecasts and consider the out of sample prediction properties. This leads to exclusion of pre-1974 cycles for Conservative and Labour forecasting. The paper then turns to the problem of day-to-day forecasting and the need to smooth forecasts so that if public opinion in the polls is stable the forecast does not change direction as the election approaches. I propose a method of ironing out daily and longer-term variation to get linear trends in forecasts and prediction intervals if public opinion is static. Given forecasts for the shares of the vote, predicting parliamentary seat total outcomes is done with the Curtice and Firth (2008) probabilistic seat prediction method. Holding the Liberal Democrats at their point estimate for the share of the vote, an approximate prediction interval for seats focusing on variation in Conservative and Labour performance is generated from the prediction intervals for shares. In turn this is used to define an approximate uncertainty distribution over seat outcomes from which approximate probabilities of key events can be calculated, including probabilities of overall majorities for the Conservatives and for Labour, the probability of a hung parliament outcome and the probability that the Conservatives will be the largest party instead of Labour.

Why forecast from polls rather than other methods?

A variety of different methods based on different kinds of data were used to forecast the last British general election. These include approval of the prime minister, local election results, citizen forecasting as well as forecasts from vote intention opinion polls (for a summary introduction to a special symposium see Gibson and Lewis-Beck (2011)). Of these, forecasting from properly adjusted vote intention polls was the most successful even though the final polls suffered some serious industry bias², especially with respect to the Liberal Democrats (Pickup et al. 2011).

In addition to reasonable accuracy vote intention polls also have the virtue of ubiquity. With the advent of internet polling they are now daily in Britain and this raises the prospect of a continually updated forecast some way before the election. There are few

¹ This time reflects the time of writing relative to the 2015 election.

² This is a technical term is used just to refer to the difference between the average poll and the election result. It is not necessarily the product of any methodological or other choice by pollsters and may not be possible to predict, control or correct for and certainly there is not accusation here of politically motivated interference.

alternatives to vote intention polls for forecasting from more than a year before the election, especially if frequent updates are desired. Economic indicators more than a year before the election are probably premature for classic forecasting from ‘fundamentals’. This ‘historical’ approach to prediction is usually conducted with data ‘three to six months before the election, sometimes a full year’ (Belanger and Soroka 2012). There are other opinion poll questions that are available well before the election, such as prime-ministerial approval but these are only available on a monthly basis for previous cycles (Lebo and Norpoth 2011).

While it is interesting to consider ‘nowcasting’ to tell us what election outcomes might be if they were held now or in a few months time (Lewis-Beck et al. 2011), there is also a virtue in long range forecasting. It is especially important for political commentators to understand not just what the likely outcome would be if the election were held today or in three months time, but when it will actually or most likely be held.

Experience of previous election cycles

The left hand side of Table 1 shows how much change there was between the polls in the twentieth month before the election and the election result itself. In the 15 electoral cycles from 1950 to 2010 that exceeded twenty months public opinion rarely changed by more than ten points and on average by only six points in either direction. The Tories have tended to improve by an average of three points and Labour drop by two and half points in the final twenty months. Some of this seems to be due to discrepancies between the final week polls and the eventual election-day result shown in the right hand side of Table 1. Either there is a regular late swing from Labour to the Conservatives, or the polls have had a tendency to over estimate Labour at a cost to the Tories. But notice while the figures are potentially important substantively the mean differences in the first row of summary statistics are not statistically significant.

The overall means mask big differences for the Conservatives and Labour according to prior government status. The right most column shows which party was in government prior to the election and a tendency for governments to recover in the final twenty months is clear from scanning the figures. The final three rows of summary statistics show mean changes by prior government and the difference in these means (a difference-in-difference estimator). Conservative performance over the last twenty months of the cycle is more than eight points better when in government than opposition. Heath was the only Tory prime minister to have gone down in the polls in the last twenty months. Similarly Labour government recovery relative to opposition performance is more than six points.

Despite the small sample size the differences are clearly statistically significant. Moreover government status explains half of the variation in the final twenty-month change across cycles for the Conservatives, and thirty per cent of that for Labour.

The same pattern is not evident for the final week polls relative to the election results, so there is no evidence that governments generally benefit from late swings or poll bias. Instead, as one would expect, the expected scale of government recovery declines as the election approaches. But note that the tendency of Conservatives to outperform and Labour to underperform final week polls is still there regardless of whether they are in or out of government. This adds strength to the observation that there is likely to be a

small systematic error in the polls even in the last week, albeit one that could end up being overwhelmed (or enhanced) by something specific to the next election.

Table 1 Difference between polls and votes a week before and 20 months before the election

Election	%Vote – Ave poll share in 20 th month before election			%Vote – Ave poll share in week before election			Prior gov
	Con	Lab	Lib	Con	Lab	Lib	
1950	-5.0	7.0	0.3	-1.1	5.3	-3.2	Lab
1951	3.7	5.2	-7.2	-4.2	8.4	-3.9	Lab
1955	4.2	-0.1	-4.2	-1.8	0.0	1.2	Con
1959	12.4	-0.3	-11.7	1.9	-1.9	-0.3	Con
1964	9.4	-5.0	-4.4	-2.4	-1.7	3.6	Con
1966				0.7	-2.6	1.1	Lab
1970	-4.2	8.4	-3.2	2.8	-4.7	0.9	Lab
Feb-74	-0.8	-11.0	10.4	-0.2	1.0	-1.8	Con
Oct-74				2.6	-3.0	-0.2	Lab
1979	-2.6	-5.0	7.9	1.2	-2.8	1.1	Lab
1983	15.0	-2.8	-12.1	-2.1	2.2	-0.2	Con
1987	8.6	-4.4	-4.4	0.3	-2.6	1.6	Con
1992	6.4	-13.0	9.8	5.5	-3.6	-1.9	Con
1997	3.8	-7.5	1.1	0.7	-3.1	1.9	Con
2001	2.3	-8.4	5.4	2.0	-3.6	1.4	Lab
2005	1.9	0.6	-3.6	1.2	-1.6	0.2	Lab
2010	-5.8	-0.1	7.2	1.7	2.2	-3.9	Lab
Mean	3.3	-2.4	-0.6	0.5	-0.7	-0.1	
S.D.	6.1	6.1	7.1	2.3	3.4	2.1	
Mean abs. dev.	5.7*	5.2*	6.2*	1.9*	3.0*	1.7*	
Max abs. dev.	15.0	13.0	12.1	5.5	8.4	3.9	
Con gov mean	7.4*	-5.5*	-1.9	0.2	-1.2	0.5	
Lab gov mean	-1.4	1.1	1.0	0.8	-0.3	-0.7	
Gov diff	8.7*	-6.6*	2.9	-0.5	-0.9	1.2	
N	15	15	15	17	17	17	

Note: Figures are differences in the GB share of the vote (Actual election share of the vote – mean share in the published polls). *Con gov mean* is mean for years when Con was in government prior to election. *Gov diff* is *Con gov mean* – *Lab gov mean*. * p<0.5

So, the main observations from Table 1 for this paper are that there is useful information in polls as much as twenty months from the election both in telling us about the likely direction of change and in the likely magnitude of changes for each party, and also that government status needs to be taken into account.

Votes-on-polls regressions

The aim of this paper is to develop a methodology that can be used to forecast the GB share of the vote for the three main parties at the next general election from current vote

intention polls, up to twenty months ahead, together with prediction intervals. In particular to have a methodology which will allow for daily forecasts that can be published in a consistent manner, for which fluctuations in the forecast will be driven either by fluctuations in the polls or otherwise easily explicable.

First it is necessary to specify the future election date. In previous election cycles it was possible to make intelligent guesses about when the next election would be. With the passage of the Fixed-term Parliaments Act 2011, the time of the next general election is no longer simply the prerogative of the prime minister but scheduled for 7th May 2015. Under the current legislation the election can only be brought forward with a two-thirds majority vote in the Commons for an early election or a motion of no confidence in the government without a new government within two weeks. So there is a high degree of confidence in the scheduled date and this paper proceeds as though that date is fixed. There were just under 600 days to the election at the start of the time of writing and so I consider forecasting over this period.

Data on opinion polls are those used in Wlezien et al. (2013) which include an estimate of the average vote intention share for each of the three main parties for every day of each election cycle, starting with the first opinion poll after 1945. Figures are averages of headline published shares, indexed by the midpoint of their fieldwork period. Where pollsters have published daily ‘tracking’ three-day moving average polls, these are used only every three days to avoid double counting of respondents. Linear interpolation is used to estimate voting intention for days without any polls. For further details of the data and discussion of their properties see Wlezien et al. (2013).

For a given number of days before the next election, call this t , it is possible to consider how, across previous election cycles, the opinion polls t days before the election corresponded with the eventual election outcome. This can be done by considering votes-on-polls regressions of the form,

$$\text{ELECTION_DAY_VOTE}_{pti} = \alpha_{pt} + \beta_{pt}\text{POLL_SHARE}_{pt} + \gamma_{pt}\text{GOV_STATUS}_{pt} + \varepsilon_{pti} \quad (1)$$

for each party p , where t is the number of days before the election and i is an election cycle indicator. The POLL_SHARE for a given day is the average of the polls for that day as described above. GOV_STATUS is a dummy variable for being the government prior to the election. These regressions can be used to forecast the 2015 election outcome with estimates of public opinion from t days before the election for each party. Despite the Liberal Democrats now being in coalition they are coded zero for GOV_STATUS and so the regressions for this party effectively omit that term.³

Wlezien et al. (2013) showed that regressions of this form, but without government status, are informative and increasingly so approaching the election. Fisher et al. (2011) showed how they can be used to forecast the share of the vote a month before the election.⁴ Such regressions allow for ‘reversion to the mean’ when parties are

³ Since this is the first time the Liberals have been in government since the war there is no evidence as to how this affects the likelihood of recovery before the election. One could look at the experience of the 1979 minority administration supported by the Liberals to argue that a recovery is likely, but this is too limited an evidence base to estimate the magnitude of the effect.

⁴ That paper provided two forecasts - one from just before the campaign and one from the final week of the campaign. The pre-campaign forecast was adjusted for the historical polls-vote relationship (stage 2

performing usually well or badly in mid term. They also average over historic polling industry bias, which has tended to be pro-Labour and adjust for that. This paper expands that approach to consider government status and argue that the prediction intervals associated with forecasts from such regressions reflect the scale of the historic experience of random poll error and unpredictable changes in public opinion before polling day.

Out of sample forecasting

While it is the case that polls are informative and there is evidence of a government recovery effect, it is necessary to consider the extent to which votes-on-polls regressions actually improve prediction over simple projection of current polling. This section considers out-of-sample prediction, i.e. running regressions of the form in equation (1) excluding each election cycle in turn and producing forecasts for the share of the vote for the omitted election. This is done separately for polls 20 months, 12 months and 6 months prior to the election. Also, because there are good reasons to think that older election cycles, especially before 1970, may be a less good guide to more recent ones, I consider weighting more recent cycles more heavily than older ones.

Also there was a distinct change in British electoral history in 1974 as Britain moved most decisively from a two to a three party system. The Liberals were a tiny party in the 1950 to 1970 elections only once winning more than 10% of the GB vote and usually contesting around half or fewer seats. Since their break through in February 1974 (with 19.8%) the Liberals have consistently contested the large majority of seats and only once secured less than 17%. Moreover, polls in the pre-1974 period were much fewer and further between, creating more uncertainty about trends in public opinion. In light of these considerations, it is worth testing whether forecasting for elections since 1974 is best done with information from this period only.

The mean overall errors from out-of-sample forecasts are shown in Table 2. Here I consider polls averaged over the 20th, 12th and 6th month prior to each election with separate results for each in different columns. The first row shows the mean overall error from using the straight polling average as a prediction of the election result. As with all the figures in the table the averaging is over all available cycles. For any given election the overall error is half the sum of the absolute differences between the prediction and the actual share of the vote for all three main parties and other parties as a group.⁵ This can be thought of as the proportion of voters who would need to be reallocated to correct the prediction of the four-party share of the vote. With this in mind the figures in the first row are substantial but not huge.

of the method) but the final forecast wasn't (given the nature of the regressions run being close to identity). Note that the pre-campaign polls-vote relationship was estimated using polls aggregated over a week before the campaign (or equivalent) for each cycle.

⁵ I.e. $\sum \text{abs}(\text{predicted}_p - \text{actual}_p) / 2$, with summation over $p = \text{Con, Lab, Lib, Oth}$.

Table 2 Mean overall forecast errors for elections since 1974 from out of sample prediction

	<i>Time prior to the election:</i>		
	20 th month	12 th month	6 th month
Mean overall error in polls	9.0	6.8	6.8
N (cycles for averaging)	8	8	9
<i>Forecasts with regressions from 1950-2010</i>			
Unweighted	7.5	7.2	6.1
Weighted	5.8	5.2	4.2
N (cycles for regressions)	14	15	16
<i>Forecasts with regressions from 1974-2010</i>			
Unweighted	5.7	4.6	4.0
Weighted	6.2	4.5	4.0
N (cycles for regressions)	8	8	9

Notes: See text for explanation. Ns reflect the exclusion of a single election for out-of-sample prediction.

The middle block of rows concern regressions based on all available cycles from 1950 to 2010, first unweighted and secondly weighted as a linear function of time since 1949. Even though all available cycles are used, the mean overall error is calculated just for elections since 1974 for comparability. Strikingly, on this measure, the unweighted regressions provide modest improvement if any over raw poll projection. The weighted regressions do better, suggesting that more recent elections are indeed more relevant.

The final block of rows shows mean overall errors for regressions using cycles from 1974 onwards, unweighted and then weighted for time since 1973. With one exception these provide improvement in out-of-sample prediction compared with regressions using all available cycles. There is not much to choose between the weighted and unweighted versions but the unweighted versions are marginally better and have the virtue of simplicity.

Table 2 suggests that using post-1974 unweighted regressions will improve prediction error by about thirty to forty per cent over treating current polls as a forecast. This is some but not a tremendous benefit. The modest improvement reflects the scale of unpredictable variation. It is also worth noting that in addition to reducing absolute prediction error somewhat, the regression almost entirely eliminates the expected bias in the prediction for out-of-sample prediction. This is to say that the out-of-sample predictions averaged out to 0.0 plus or minus 1.0 and no party systematically over or under performed in the out of sample regression based predictions.

In addition to understanding how well the forecasting method performs for past elections it is also important to consider whether there are factors that might affect the performance for this election. One issue is the rise of UKIP for which there is insufficient historical electoral or polling data to be able to forecast shares for the party with this kind of technology, but allowance can be made in the seats calculation. More important for three-party share forecasting is the advent of the coalition government,


after which support for the Liberal Democrats in the polls dropped dramatically in the second half of 2010, most likely as a result of more left-wing former Liberal Democrats switching support to Labour (Curtice 2013). Since the start of 2011 the Liberal Democrats have been consistently polling close to 10%, less than half of their 2010 vote. So while the Conservatives and Labour have been polling within their interquartile range for polls between 1974 and 2010, the Liberals are currently standing near the bottom of their distribution for this period.

This not to say that the current cycle is completely out of the range for the 1974-2010 period for the Liberals. But it does have a very important impact on the trajectory for the Liberal forecast if they persist close to 10% in the polls. Essentially, from such a low base, votes-on-polls regressions from 1974-2010 period suggest a recovery to over 15% right up to the final week when the forecast would rapidly move in line with the polls as the historical correlation between polls and vote dramatically strengthens. So for most of the run up there would be too strong a reversion-to-the-mean effect to want to apply given that it is very plausible, under the new political circumstances, that the Liberal Democrats may not recover much at all by the election. Incorporation of pre-1974 cycles into the Liberal forecasting regressions removes this problem because it increases the range of Liberal polls and votes closer to zero making current polling levels closer to the historical average. For this reason I believe it is better to use all prior cycles from 1950 onwards for forecasting the Liberal Democrat share of the vote in 2015, even though the Conservative and Labour forecasts are best done based on regressions with 1974-2010 cycles. This is the basis on which the rest of the paper proceeds.

Smoothing for day-to-day forecasting

A significant problem with daily forecasting from votes-on-polls regressions (equation 1) is that predictions have considerable temporal variation even for stable public opinion as measured by the polls. For example, consider what would happen if the Liberal Democrats continue to flat line in the polls on 10% as they have been doing. The following figure shows a scatter plot of the raw forecasts from our daily regressions for the 600 days up to the election. Two things are apparent. First, there is lots of day-to-day variation in the forecast. Second, over the first year the forecast trends downwards and then it goes back up again in the second before falling in the final month. More generally, equivalent graphs for show various different monthly or longer trends and waves.

Figure 1 Forecasts of Lib Dem election-day shares with Lib Dem poll share of 10%, using votes-on-polls regressions



It is possible to smooth over day-to-day variation easily enough with moving averages or other approaches, but the weekly/monthly and longer terms waves are much more problematic for both presentation and more substantive reasons. For presentation is it hard to explain to a public (or any) audience why a forecast should change direction if public opinion does not change. It is possible to argue that a party should be expected to over or under perform its current poll share but not both at different times with the same share of the vote. Again using the example above, if the Liberal Democrats stay on 10% and the forecast drops from above 10 to below 10 over the first year and then rises above 10 again in the second year, the question would be why does the forecaster think the Liberal Democrats will do better than the polls twenty months before the election but worse than the polls a year out from the election. There is no good explanation for this.

The related more substantive reason why such waves are problematic is that they have no origins in theoretical or consistent empirical phenomena. There is nothing in the literature on public opinion or election campaigns that lead us to expect such waves over the two years before an election. They are idiosyncratic and accordingly the patterns vary from party to party and somewhat according to the level of support. Thus such waves need to be smoothed out because there is no basis on which to think they are informative about what will happen in this or future election cycles.

Linear smoothing

What is ideally required is some way of identifying the best smooth monotone (i.e. always either rising or dropping, maybe sometimes flat but never changing direction) sequence of forecasts given constant public opinion for a given party at a given level of

support. It is reasonable to suppose that the way in which public opinion approaches the eventual outcome on average follows an arbitrary smooth monotone function.

Exploratory work for this paper considered quadratic and other approaches to smoothing but did not identify a method that would ensure a smooth monotone function beyond a simple linear trend. It also revealed that for a reasonable range of possible levels of support for each party, a linear trend was a decent approximation most of the time especially when the trend was forced to go through the final forecast one day before the election.

Fitting a linear trend to sequential forecasts from a given stable level of support that goes through the final forecast can be done as follows. For a given party at a given level of support take the sequence of forecasts from the votes-on-polls regressions for each t from 1 to say 600. Centre the data on the final forecast (from the day before election day)⁶, fit a no-constant OLS regression model to the forecast, and then use the fitted values (after decentering) from those regressions to provide the point-estimate forecasts for election-day shares of the vote. Formally the regression is of the following form where $RAW_FORECAST_{pxt}$ are the forecasts from votes-on-polls regressions (in equation 1) for party p , with poll share x at time t days before the election.

$$RAW_FORECAST_{pxt} - RAW_FORECAST_{pxl} = b_{px} * t. \quad (2)$$

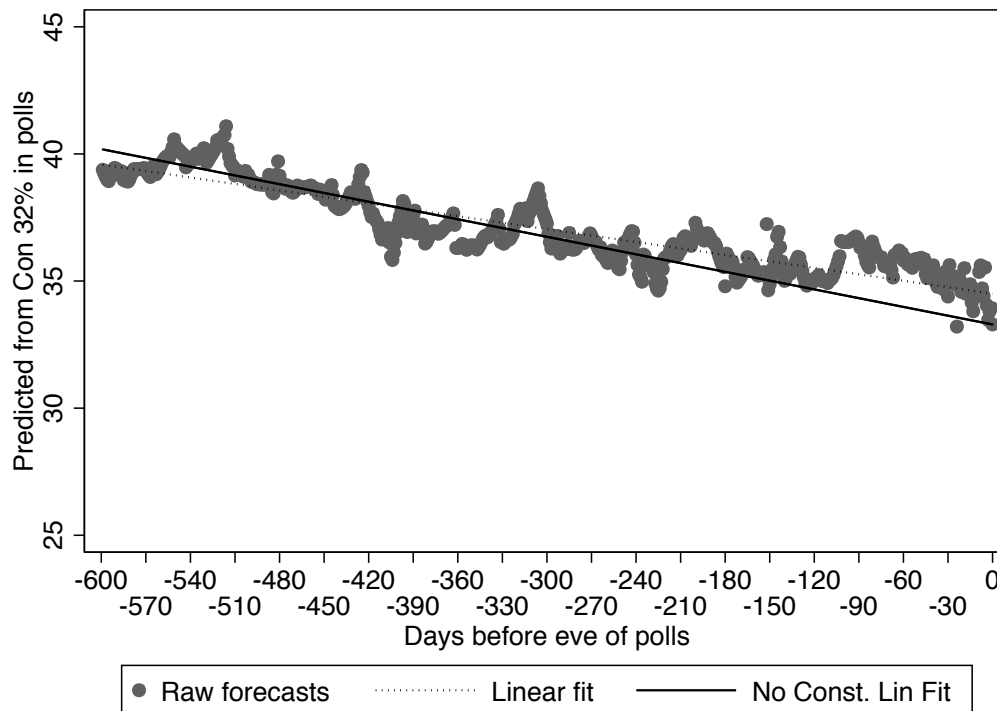
This linear regression forces the forecast to progress steadily in a linear way as you approach the election. Since this process forces the sequence of forecasts onto a straight line, it is more like ironing than smoothing. So the fitted values will be referred to as the ironed forecasts, which for party p on share x at t days before the election are given by,

$$IRONED_FORECAST_{pxt} = b_{px} * t + RAW_FORECAST_{pxl}. \quad (3)$$

Figure 1 shows both a regular OLS fit and the result of the no-constant regression on centred data. In this case there is very little difference between the no-constant regression over the regular OLS which ends very close to the final prediction. However Figure 2 below shows an example with the Conservatives on 32% (reflecting current public opinion) with a starker difference. The fact that the no-constant regression line is below the forecasts for the final four months is an example of how the no-constant regression typically attenuates any reversion-to-the-mean effect from the votes-on-polls regressions late in the cycle and brings the smoothed forecasts closer to the polls. This is probably a good thing and so an added bonus from smoothing.

⁶ Note I realise that the words centred and centring are odd given the location is far from the centroid of the data, but my understanding is that this terminology is used regardless of the chosen point on which to 'centre' the data.

Figure 2 Forecasts of Conservative election-day shares stable poll share of 32% from votes-on-polls regressions



The steep slope of the line reflects the expected scale of the Tory recovery. Over 500 days out they are expected to gain over eight points by the election, but if they continue to poll at 32% their expected election share of the vote declines to 35%, but no lower. This final three-point difference reflects historical polling industry bias.

Prediction Intervals for Shares of the Vote

Prediction/forecast standard errors from each votes-on-polls regression (equation 1) can be generated for given levels of party support straightforwardly (Greene 1991, p.166-8). The inexplicable variation in these can then be ironed out using the same process described above for the ironing the trajectory of the point estimates. Prediction intervals using ironed forecast standard errors naturally narrow in a linear fashion as the election approaches. This is important for both presentation and substantive reasons as for the point estimates. People expect to become more and more confident in the forecast the closer the election. While it would be fair to say that the rate of narrowing might not be linear, there remains the same problem of how to do monotonic smoothing and again ironing out with a linear trend is a good approximation most of the time.⁷

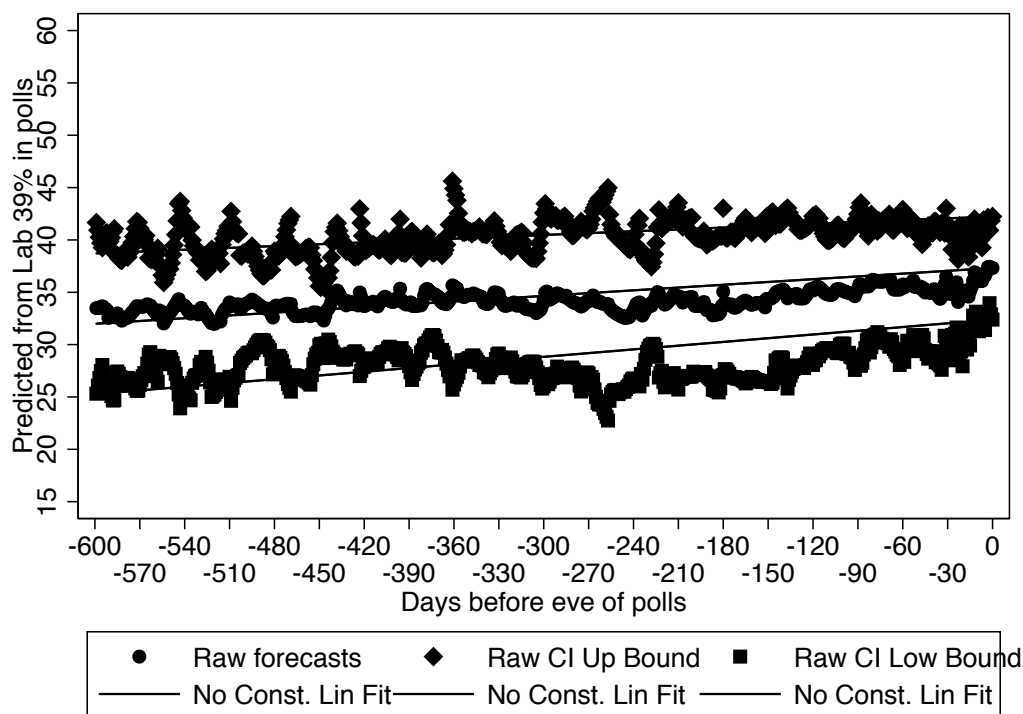
Figure 3 shows an example of how the upper and lower bounds of a 95% forecast confidence interval for the Labour election-day share can be ironed as well as the point estimate. The width of the prediction interval for Labour narrows a little but not much.

⁷ In truth there is perhaps the possibility of developing a theory of when opinion turning events are most likely to occur, e.g. at party conferences and Budget speeches, and which periods are likely to be disproportionately quiet, e.g. holidays. Developing this idea would require too much groundwork for this paper and various informal analyses suggest that there is limited systematic lasting legacy of budgets and conferences.

The other two parties have wider intervals to start with and they attenuate more over the period. These party difference primarily reflect the better fit of the Labour votes-on-polls regressions than that for the other parties, but why the Labour models should fit better is not clear.

While most of the time the ironing seems to be a good job there are some sustained periods where ironed prediction intervals are noticeably different from the raw intervals. For Labour the smoothing arguably narrows the prediction interval too much between 90 and 300 days out. For the Tories the smoothing widens the prediction interval too much between 30 and 240 days out. But conversely it is not clear why the raw prediction intervals should be unusually wide/narrow in these periods and so maybe the smoothing is doing the right thing to iron these differences out. Most importantly, for all parties there is much more uncertainty than you would expect just from sampling error in the polls. This is true even the day before the election because of the varying nature of polling industry bias over time.

Figure 3 Forecasts of Labour election-day shares stable poll share of 39% from votes-on-polls regressions, and upper and lower bounds for 95% prediction intervals



Predicted seat totals and approximate prediction intervals

Predictions for seat totals for each party can be generated using the method first outlined in Fisher et al 2011 which involves applying a uniform change in the share of the vote implied by the predicted share on to each constituency and then using the formula below developed by Curtice and Firth (2008) estimating probabilities of each party winning each seat. Here the probability for party j to win in constituency i from the predicted shares of the vote s_{ij} as $r_{ij} / \sum_j r_{ij}$ where

$$r_{ij} = \exp \left[- \left(\frac{\max_j \{s_{ij}\} - s_{ij}}{\sigma} \right)^\lambda \right], \quad (4)$$

where σ and λ are tuning constants. For the purposes of election-night prediction in 2005 the values were $\sigma=4$ and $\lambda=1.5$. These were picked using experimentation on 2001 results and they produced well-calibrated probabilities for 2005 (Curtice and Firth, 2008) and 2010 (Curtice et al. 2011) and so these values have been used again here. The predicted seat total for a party is then just the sum of the predicted probabilities for the party across constituencies.

This method can straightforwardly be applied to the point estimate for the share of the vote with some assumptions about the nationalist and minor parties and special seats.⁸

The idea of prediction intervals for seats for each party is complicated conceptually. To estimate seats predicted shares for all three parties are needed. But since the prediction intervals for each party share were generated independently it is unclear where to set the performance of the other two parties when considering one of the bounds for another party. E.g. it makes no sense to ask what the seats would be if all three parties were at their upper bounds. Instead I consider the range of outcomes for the Conservatives and Labour, fixing the Liberal Democrat and Other shares at their point estimates. Thus I construct a prediction interval for the Conservative lead over Labour assuming a constant sum for the two-party vote by pooling the forecast variances for the Conservative and Labour prediction intervals discussed above. The upper and lower bounds of this prediction interval imply two distinct scenarios:

A: Conservative forecast share close to its upper bound, Labour close to its lower bound and Liberal Democrats at their central point estimate;

B: Labour forecast share close to its upper bound, Labour close to its lower bound and Liberal Democrats at their central point estimate.

These scenarios are the most politically interesting since seat totals for the two main parties are the most important for government formation and they depend primarily on their performances and are relatively insensitive to the performance of the Liberal Democrats. So there is less need to consider possible variation in the Liberal Democrat share than in that for the two main parties.

Scenario A provides a best possible outcome for the Conservatives given the current polls and historical relationship between the polls and the vote, and similarly scenario B is the best possible outcome for Labour. Having identified the forecast shares of the vote for the three main parties in scenarios A and B, seat totals for all parties under these scenarios can be estimated with the formula above. The predicted seats for scenarios A and B provide approximate prediction intervals for seats.

⁸ I assume no change in vote share for the SNP and Plaid Cymru. Otherwise I assume that the largest minor party in the constituency in 2010 will get all the increase or decrease in the Other share of the vote. In many cases this will be UKIP. I further assume that the Green's will hold their seat, that Wyre Forest will stay Conservative (after they narrowly defeated an independent incumbent in 2010) and that the speaker's seat will remain Conservative.

Predicted probabilities of overall majorities and hung parliaments

To estimate the probabilities of a Conservative or Labour majority or a hung parliament some estimate of the uncertainty distribution over seat totals is needed. This is a multidimensional problem if all parties are to be taken fully into consideration, but given only the top two parties have any prospect of gaining an overall majority and their seat totals are relatively insensitive to variation in the levels of support for the Liberal Democrats and others, it is reasonable to take an approximation approach as with the prediction interval for seats. Moreover there are benefits for communication in providing probabilities of key events that are consistent with the seats prediction intervals.

The approach I take here is to assume that the approximate prediction intervals for seats are in effect 95% confidence intervals from normal distributions describing the probabilities of different seat totals for each party. For each party the centre of the distribution is the best point estimate for seats as described above and a quarter of the width of the prediction interval provides an estimate for the standard deviation for the number of seats for a party.⁹ Given the normal distributions for each party thus defined, the probability of an overall majority for the Conservatives or Labour can be inferred from the normal distribution function at the predicted minus threshold (326) number of seats. The probability of a hung parliament is then given by

$$\Pr(\text{Hung parliament}) = 1 - \Pr(\text{Conservative majority}) - \Pr(\text{Labour majority}). \quad (5)$$

The probability that the Conservatives will be the largest party is $\Pr(\text{Con seats} - \text{Lab seats}) > 0$. Again an approximate solution is to assume that the uncertainty in the predicted difference of seats for the two parties follows a normal distribution with a mean given by the difference in the point estimates and a standard deviation as the square root of the average variance in seats for the Conservatives and Labour, as calculated for the approximate prediction intervals above.

Example forecast with discussion

Table 2 shows a sample forecast calculated with the above methodology. The imputed current average poll shares are the www.ukpollingreport.co.uk polling average. The first of the two most striking features of the forecast are that the Conservatives are forecast to do much better than they are currently polling, mainly because they are in government but also partly because they are historically at a relatively low point (from which regression suggest they should recover) and as a party they have tended to outperform polls at elections. Labour are predicted to do correspondingly worse and so the forecast for the top two is almost symmetrically opposite from the current polls: 40:32 instead of 32:39.

⁹ In practice the prediction intervals for seats are slightly skewed and doubtless a full simulation exercise would not produce a density which was perfectly normal, but the approximation seems good enough for the purpose.

Table 2 Example forecast from 8th October 2013

Date of forecast: 8.10.2013

Days till the election: 576

Inputed current average poll shares

Con : 32

Lab : 39

LD : 10

Forecast Election Day Shares and 95% Prediction intervals

Con : 39.9 plus or minus 12 i.e. between 28 and 52

Lab : 32.2 plus or minus 6.6 i.e. between 26 and 39

LD : 11 plus or minus 14.8 i.e. between 0 and 26

Forecast Election Day Seats

Con : 333

Lab : 270

LD : 20

Con majority of 16

Forecast Election Day Seat approx. 95% Prediction intervals

Assuming LD share at 11 and allowing Con and Lab to vary as per interval above

Con between 216 and 468

Lab between 143 and 381

LD between 13 and 28

Probabilities of key outcomes

Pr(Con majority) = 54%

Pr(Lab majority) = 17%

Pr(Hung parliament) = 28%

Pr(Con largest party) = 85%

Pr(Lab largest party) = 15%

The second of the two most striking features is that the prediction intervals for shares of the vote are enormous. For the main parties it is clear that there is more uncertainty in the Conservative vote but even the forecast Labour vote could be out by as much as 6.6 points: a huge political difference. At first glance these prediction intervals may seem to encompass all foreseeable outcomes and more. For the Liberal Democrats the interval from nothing (a hard boundary that had to be invoked!) to 26 seems ridiculously large. While these may seem hilarious at first sight, remember that not all points within the intervals are equally likely to occur. Also as 95% forecast confidence intervals they reflect the historical variation in the votes for these parties and there should be only a 5% chance of a result outside the interval. So they are bound to be very broad to be credible. Even so, the lower bound for the Conservative forecast, at 28, tells us that it is very unlikely the Conservatives will do much worse than they currently stand in the polls, which is informative. Similarly, the Labour prediction interval suggests it is extremely unlikely that Ed Miliband could do as well as Tony Blair in 1997 or 2001 but could do worse than Gordon Brown or Michael Foot.

The forecast election-day seats are as you would expect them given the forecast shares. A classic uniform change prediction would produce slightly different figures but not by much, especially given the large prediction intervals for seats that follow from the large prediction intervals for votes.

The estimated probabilities for key outcomes are perhaps the most helpful feature of the forecast so far from an election. These show that the Tories have an 85% chance of being the largest party but only 54% chance of an overall majority. At this stage in the cycle the main reason why these probabilities are not more one sided is because there is plenty of scope for change in party support before the election. But even immediately before the election there is considerable uncertainty given the historic record of the polls as shown in Table 1.

Having some estimate for the probability of a hung parliament (28% above) is particularly helpful given that it is otherwise hard to get a sense of this. Traditionally there has been a tendency for commentators to underestimate the chances of a hung parliament (Gibson and Lewis-Beck 2011).¹⁰ If the parties follow the expected trajectories or similar, the methodology here suggests the probability of a hung parliament is likely to increase up to the final month to around 40% as the prediction intervals for seats narrow towards hung parliament territory, but then diminish again in the final month if the Conservatives do have the predicted strong lead. A Liberal Democrat recovery would also increase the chances of a hung parliament.

Methodological discussion

There are two key methodological decisions within the approach taken here. The first was to incorporate government status as a predictor in the votes-on-polls regressions. This has a big effect on long range forecasting, boosting estimates for the Conservatives and lowering those for Labour for 2015 by about four points over a year out and by around a couple of points even as close as three months out. The second was to use the full range of cycles for votes-on-polls regressions for the Liberal Democrats rather than just the 1974-2010 period. Had the post 74 cycles only been used the Liberal Democrat forecast in Table 2 would have been lifted to as much as 18%, as the product of a strong regression-to-the-mean effect. Current point estimates for the Conservatives and Labour are fairly insensitive to the base set of cycles for the votes-on-polls regressions, but ironically the prediction intervals for Labour are much narrower as result of the smaller more appropriate post-74 sample. The Conservatives and Labour have not regained their pre-74 dominance with consistent 40+ shares of the vote, but it is safer to hedge bets on whether the Liberal Democrats will continue their post-74 success or return to pre-74 levels of performance.

It might be thought that there is a danger that some election cycles are driving the results here more than others. Exploratory data analysis suggests not. Even dramatic cycles like 1983 with the advent of the SDP and the Falklands war are not clear outliers in the votes-on-polls regressions. The famous failure of the polls in predicting the Conservative victory in 1992 means that that cycle is the furthest from the regression line for the Conservative share, but it is not clearly more influential than others.

One of the most undesirable features of the method of smoothing/ironing forecasts outlined above is that the outcome of the exercise depends on the range of days over which you fit the regression in equation (2). I have considered the period up to 600 days before because that reflects the time of writing. Extending back much before this period

¹⁰ Approaching the last election a hung parliament was more likely than not, but this was not reflected in the media debate. Things may be different now that there is hung parliament with a coalition, but it could still be seen as an aberration or people may become too inclined to think it will happen again.

would involve incorporation of much more unstable patterns of correspondence between polls and votes, so it is potentially problematic to apply this approach to longer range forecasting. Narrowing the period over which smoothing is done would make some but not considerable difference to the trajectories of point estimates. The implications for prediction intervals are more dramatic since in practice they tend to narrow more dramatically in the last few months than on a steady basis. There might be some junctures at which it makes sense to switch to a new basis for smoothing (e.g. the start of the campaign). Further work on the functional form might also be fruitful, but overall, the linear ironing seems to be a reasonable approach given approximate nature of the methods overall.

Further elaboration of the votes-on-polls regression models is possible. One could incorporate a cost of governing effect (Paldam 1986) or similarly an electoral pendulum effect (Lebo and Norpoth 2007). Especially closer to the election it may be worth adding leading economic indicators (Erikson and Wlezien 2012) or considering whether PM approval data add value (Lebo and Norpoth 2007). But it is also worth bearing in mind that with just 17 prior electoral cycles there is a danger of over fitting the data. There are also interesting questions of model selection when evaluating the same model fitted to hundreds of times (one for each day before the election). This would be a good topic for future research.

For the example forecast I have used just one possible way of computing a polling average, but there are others. To be most consistent with the historical data a straight average of the polls on the day with interpolation would be most natural. But it is probably best to use the best estimate of current vote intention on the basis that better information in means better information out. More generally the technology can be used to generate forecasts for different scenarios so different poll averages can be compared.

A conceptual difference in the approach of this paper compared with the two-step approach by Nadeau et al. (2009) is to treat current opinion poll vote intention, even three months or closer to the election, not as an ‘empirical proxy’ (Nadeau et al. 2009) but as a something that can and will change in systematic and unpredictable ways. This idea becomes increasingly questionable as the election gets closer and certainly by the final week the historical votes-on-polls regressions may well be picking up primarily historical polling errors. Whether or not these are relevant for short range forecasting needs further consideration.

Conclusion

The development of a method for long range forecasting of the next British general election from opinion polls has been useful for illustrating a number of important points about the historical relationship between opinion polls and election results and about the operation of the electoral system. These include the following observations. The polls are informative about the likely outcome even a long time beforehand. Historical experience suggests that the Conservatives are likely to do better at the polls than they do in the polls, and vice versa for Labour. This is true, but to a diminishing extent, right up to the final week. Relative to this phenomenon, there is a very strong tendency for the two main parties to do better over the final twenty months if they are in government rather than in opposition. Although the tendency for Conservatives to outperform polls

and any party to recover from a weak position are important, their position in government is the main driving force behind a forecast of a decent Conservative lead at the election despite the party currently trailing in the polls by a similar amount.

In addition to suggesting systematic changes between now and the election, the historical data also show that substantial unpredictable changes have happened over the final twenty months before an election in the, but nonetheless really big changes are still relatively rare. This helps to explain how forecast confidence intervals are very broad, although not uninformative. While the Conservatives, at the time of writing, are clearly expected to be the largest party at the election, there is still a 15% chance that Labour will be.

Approximate methods for estimating the probabilities of key outcomes such as single party majorities and hung parliaments have also been developed here. These are important for understanding the most politically salient features of the forecast point estimates and uncertainty. They also help illuminate the operation of the electoral system by showing how the probability of a hung parliament, currently at 28%, is relatively high and may well grow.

The paper also outlined a method for producing daily forecasts that will not bounce around dramatically or change direction if there is stable public opinion in the polls, which would otherwise happen with raw forecasts from daily votes-on-polls regressions. These are important features for public communication and to retain the most salient and theoretically well-founded features of the historical data. So the methodology should be helpful for understanding the implications of current public opinion for future political outcomes in a way that also informs about the operation of electoral politics in Britain.

Acknowledgements

I am very grateful to Jouni Kuha and Mark Pickup for comments on previous drafts and to Rob Ford, Will Jennings, Mark Pickup and Chris Wlezien for early discussions about this work. Most especially I am indebted to Will Jennings for compiling and preparing the polling data, including interpolation and correction for overlaps, as used in Wlezien et al. (2013).

References

- Bélanger, Éric, and Stuart Soroka. 2012. "Campaigns and the prediction of election outcomes: Can historical and campaign-period prediction models be combined?." *Electoral Studies* 31(4): 702–714.
- Curtice, John. 2013. "The Political Landscape", in PSA eds. *PSA local elections media briefing pack 2013*. <https://www.psa.ac.uk/sites/default/files/Local%20Elections%202013-briefingpack.pdf> (Accessed May 2, 2013).
- Curtice, John, Stephen D Fisher, and Jouni Kuha. 2011. "Confounding the Commentators: How the 2010 Exit Poll Got it (More or Less) Right." *Journal of Elections, Public Opinion & Parties* 21(2): 211–235.
- Curtice, John, and David Firth. 2008. "Exit polling in a cold climate: the BBC–ITV experience in Britain in 2005." *Journal of the Royal Statistical Society. Series A:*

- Statistics in Society* 171(3): 509–539.
- Erikson, Robert S, and Christopher Wlezien. 2012. “The Objective and Subjective Economy and the Presidential Vote.” *PS: Political Science and Politics* 45(04): 620–624.
- Fisher, Stephen D, Robert Ford, Will Jennings, Mark Pickup and Christopher Wlezien. 2011. “From Polls to Votes to Seats: Forecasting the 2010 British Election.” *Electoral Studies*.
- Gibson, Rachel, and Michael S Lewis-Beck. 2011. “Methodologies of election forecasting: Calling the 2010 UK ‘hung parliament’.” *Electoral Studies* 30(2): 247–249.
- Greene, William H. 1991. *Econometric Analysis*. Macmillan: New York.
- Johnston, Ron, Galina Borisyyuk, Michael Thrasher, and Colin Rallings. 2012. “Unequal and Unequally Distributed Votes: The Sources of Electoral Bias at Recent British General Elections.” *Political Studies* 60(4): 877–898.
- Lebo, Matthew J, and Helmut Norpoth. 2007. “The PM and the pendulum: dynamic forecasting of British elections.” *British Journal of Political Science* 37(01): 71–87.
- Lewis-Beck, Michael S, Richard Nadeau, and Éric Bélanger. 2011. “Nowcasting V. polling: The 2010 UK election trials.” *Electoral Studies* 30(2): 284–287.
- Nadeau, Richard, Michael S Lewis-Beck, and Éric Bélanger. 2009. “Election Forecasting in the United Kingdom: A Two-Step Model.” *Journal of Elections, Public Opinion & Parties* 19(3): 333–358.
- Paldam, Martin. 1986. “The distribution of election results and the two explanations of the cost of ruling.” *European Journal of Political Economy* 2(1): 5–24.
- Pickup, Mark, J Scott Matthews, Will Jennings, Robert Ford, and Stephen D Fisher. 2011. “Why Did the Polls Overestimate Liberal Democrat Support? Sources of Polling Error in the 2010 British General Election.” *Journal of Elections, Public Opinion & Parties* 21(2): 179–209.
- Wlezien, Christopher, Will Jennings, Stephen D Fisher, Robert Ford and Mark Pickup. 2013. “Polls and the Vote in Britain.” *Political Studies* 61: 66–91.