

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 several methods for forecasting the next British general election on a daily basis from twenty months prior using opinion polls. It discusses their performance for previous electoral cycles and shows that the two models with the best historical record lead to substantially different predictions for 2015, but they can be averaged. The historical relationship between the polls and the vote suggests 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 models suggest a substantial Conservative lead at the 2015 election. Approximate prediction intervals and probabilities for key events are also generated. According to the average of the two best models and polls in October 2013, the Conservatives were estimated to have a 64% chance of being the largest party and 42% chance of an overall majority. The estimated probability of a hung parliament (40%) 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 40% chance of a hung parliament and only a 42% chance of a Conservative majority in October 2013. Labour would clearly be set for a majority if the Conservative and Labour forecast shares of the vote were reversed because of the well known bias in the electoral system (Johnston et al. 2012).

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 three main types of votes-on-polls regressions that can be used as bases for forecasts. Out-of-sample prediction properties identify two formulations as being better at predicting previous elections: a votes-on-polls-and-government-status model and a relative change since the last election model. Although these two models have typically been very close to the final result or wrong in the same direction in previous cycles, there are important substantial differences between them for a twenty-months out forecast of the 2015 election, especially with respect to the Liberal Democrat share. So an average of the two models is considered.

The paper also discusses 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

¹ This time reflects the time this work was started 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

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 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

Table 1 shows three sets of differences between polls and election results. The left hand side gives the changes from the previous election to the polls twenty months before the next election, the middle set shows change from the polls twenty months prior to the next election, and the right hand side shows change from the average of the final week polls to the election result. The final column shows which party was in government before the election. Various patterns can be discerned that constitute important tendencies in the operation of electoral cycles in Britain that might be repeated in the future.

First, governments tend to lose support up to the twentieth month prior to the next election, while the main opposition party tends to gain over the same period. The middle set of columns shows us that more often than not the losses are partly recouped and the gains are partly undone in the final twenty months, but sometimes there is overcorrection and sometimes there is no reversal of fortunes. The means under Conservative and Labour governments separately show that on average there are strong cycles for governments and principal oppositions that run in opposite directions. By twenty months before the next election Tory government support dropped by 10.6 points on average, but 7.4 points were recovered on average. Heath was the only Tory prime minister to have gone down in the polls in the last twenty months. Similarly

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.

Labour oppositions have in the final twenty months on average lost ground that was gained earlier in the cycle. The ‘Gov diff’ row gives the difference between a party’s performance according to the colour of the government. This difference-in-difference estimator is a measure of the strength of the electoral cycle effect and the figures are statistically and substantively significant for the two main parties for both the first and second parts. If previous experience of different government and opposition electoral cycles is a reasonable guide to the current election cycle then we can expect some Conservative gains and Labour set back before 2015 relative to the polls in October 2013.

Table 1 Differences between polls twenty months before the election and both previous and eventual elections, and between polls one week out and the election

Election	Ave poll share in 20 th month before election - %Vote in previous 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	Con	Lab	Lib	
1950	8.6	-9.0	-0.2	-5.0	7.0	0.3	-1.6	2.2	-1	Lab
1951	1.2	-2.6	0.5	3.7	5.2	-7.2	-2	3.3	-1.1	Lab
1955	-2.8	-1.9	4.4	4.2	-0.1	-4.2	-1.5	0.1	0.9	Con
1959	-12.8	-2.5	14.9	12.4	-0.3	-11.7	0.6	-1.6	0.9	Con
1964	-15.3	5.2	9.8	9.4	-5.0	-4.4	-1.4	-1.6	2.6	Con
1966							1.1	-2.8	1.0	Lab
1970	7.5	-9.4	-0.6	-4.2	8.4	-3.2	1.9	-3.9	1.0	Lab
Feb-74	-6.8	5.2	1.8	-0.8	-11.0	10.4	0.5	2.1	-3.3	Con
Oct-74							2.6	-3.0	-0.2	Lab
1979	8.9	4.7	-13.6	-2.6	-5.0	7.9	1.2	-2.8	1.1	Lab
1983	-16.4	-6.6	24.0	15.0	-2.8	-12.1	-2.1	2.2	-0.2	Con
1987	-8.9	7.6	1.5	8.6	-4.4	-4.4	0.3	-2.6	1.6	Con
1992	-6.8	16.7	-14.6	6.4	-13.0	9.8	5.5	-3.6	-1.9	Con
1997	-15.1	16.6	-2.2	3.8	-7.5	1.1	0.7	-3.1	1.9	Con
2001	-1.1	6.1	-3.8	2.3	-8.4	5.4	2.0	-3.6	1.4	Lab
2005	-1.4	-6.5	7.4	1.9	0.6	-3.6	1.2	-1.6	0.2	Lab
2010	9.5	-6.3	-6.2	-5.8	-0.1	7.2	1.7	2.2	-3.9	Lab
Mean	-3.4	1.2	1.5	3.3	-2.4	-0.6	0.6	-1.1	0.1	
S.D.	8.9	8.2	9.6	6.1	6.1	7.1	1.9	2.4	1.7	
Mean abs. ch.	7.2	6.3	6.2	5.7	5.2	6.2	1.6	2.5	1.4	
Max abs. ch.	16.4	16.7	24.0	15.0	13.0	12.1	5.5	3.9	3.9	
Con gov mean	-10.6*	5.0	5.0	7.4*	-5.5*	-1.9	0.4	-1.0	0.3	
Lab gov mean	4.7*	-3.3	-2.4	-1.4	1.1	1.0	0.9	-1.1	-0.2	
Gov diff	15.3*	-8.3 ^a	-7.3	8.7*	-6.6*	2.9	-0.5	-0.1	0.1	
N	15	15	15	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<0.05. ^ap=0.053 so only marginally insignificant.

Second, the tendencies for governments to fall then rise but oppositions to rise then fall

over the cycle is part of a more general tendency for initial changes in party support to be reversed. While it looks from the averages (overall and by governing party) that the election cycles for Liberals are like Labour's, with the party going up and then down under Tory governments but down and then up under Labour, the Liberal averages are heavily influenced by the large swings during the 1979 and 1983 cycles. More generally though it is still true that for the Liberals, along with the two main parties, mid term gains or losses tend to be reversed by the time of the election. If repeated for the 2015 election, the Liberals should recover somewhat.

Third, the consistency and strength of the cycle effect are greater for the Conservatives. While the average change and variation in changes up to and after the twentieth month prior to an election are similar in magnitude for the Tories and Labour (see the Standard Deviation, Mean absolute change and Maximum absolute change rows), the Conservative figures more frequently conform to an electoral cycle model of mid term changes partly reversed by the next election. This suggests that for the current electoral cycle a Conservative recovery should be more likely and larger than any Labour set back in the final twenty months. But note that one reason why the large average Conservative government mid term losses then recovery are not matched by equally large Labour poll gains and then reversion is that on average Liberals have also benefited from Conservative mid-term slumps. The reverse has happened for the Liberals since 2010, so perhaps as a consequence of the Liberals in government we should expect a stronger Labour opposition cycle effect than previously.

The fourth and final main pattern to note from Table 1 focuses just on the end of the cycle. The Conservatives have on average over performed in elections relative to the polls twenty months prior, and correspondingly Labour under performed on average. This is true of all but two of the elections since 1974. This is true both overall and specifically under Conservative governments. Whereas Tory governments have typically recovered most of their early cycle losses in the last twenty months, in opposition Labour have on average lost slightly more in the final twenty months than they gained in the polls up to the twenty months out. Some of this asymmetry 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. Again this is true of all but two of the elections since 1974, and it is true on average for Conservative and Labour governments. While it is true that 2010 was one of the exceptions, in itself this is not sufficient evidence that the tendency has gone away for good. A more important caveat to note is that the figures are small and statistically insignificantly different from zero. Still, they are potentially important substantively and so these observations suggest that the Conservatives may do a little bit better at the election, and Labour a bit worse, relative to the polls because of either late swing or polling industry bias.

Data

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).³

Votes-on-polls regression methodology

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.⁴ 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.

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 (for example see Erikson and Wlezien (1996) and Brown and Chappell (1999) for applications in the US, Selb and Munzert (2013) for Germany and Fisher et al. (2011)).

I extend this approach to consider three related models: vote shares regressed on poll shares and a government status dummy variable; vote shares regressed on poll shares and the share of the vote at the previous election (a lagged dependent variable or LDV); and the change in vote share since the last election regressed on the difference poll shares and the last election result. These are described in turn in this section and evaluated according to out-of-sample prediction quality in the next.

Votes on Polls and Government Status

These regressions have the following form:

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

for each party p , where t is the number of days before the election and i is an election

³ I have added 96 polls that were missing from the Wlezien et al. (2013) dataset, all of them from the final three months of election cycles up to February 1974. Also note that allocations of fieldwork midpoints involved rounding down not rounding up as was mistakenly reported in Wlezien et al. (2013).

⁴ The 2015 general election is expected regardless of the outcome of the referendum on Scottish independence.

cycle indicator. The VOTE is the share of the vote for the party at the election. POLL_SHARE is the average of the polls for a given day as described above. GOV_STATUS is a dummy variable for being in government prior to the election. Essentially, equation (1) describes a series of t separate regressions, one for each day before the election. Each regression has as many observations as there are election cycles since 1945 that were at least t days long. In effect they tell us how public opinion changed in the last t days of the election cycle, taking into account government status.

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.⁵

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.⁶ Such regressions allow for ‘reversion to the mean’ when parties are performing usually well or badly in mid term compared to their historical average. They also average over historic polling industry bias, which has tended to be pro-Labour and adjust for that.

Votes on Polls and Prior Election (LDV)

Similar regressions but with a lagged dependent variable (LDV) are possible and sensible. One such model has the following form:

$$\text{VOTE}_{pi} = \alpha_{pt} + \beta_{pt}\text{POLL_SHARE}_{pti} + \gamma_{pt}\text{PREV_VOTE}_{pi} + \varepsilon_{pti} \quad (2)$$

again for each party p , where t is the number of days before the election and i is an election cycle indicator. The difference between equations (1) and (2) is that GOV_STATUS is replaced by the share of the vote for the party at the previous election (PREV_VOTE). Models of this form essentially say that the expected election result is a weighted sum of the previous result and the current polls for a party. Naturally enough different estimates of this model for different days in the cycle show that current polls become increasingly important and the partial predictive power of the previous election weakens as the election approaches. But the previous election does not become close to being irrelevant until the final week. So a tendency for some reversal of fortunes remains till near the election.

It is possible to add a government status dummy variable to this model, but it proved to have a consistently small and statistically insignificant coefficient. This does not tell us that government status does not matter, only that with a small number of electoral

⁵ 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.

⁶ 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 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.

cycles it is hard to identify a government status effect above and beyond the electoral cycle effect that this model picks up. Government status is highly correlated with whether or polls are higher or lower than the previous election result. This model effectively tells us that knowing by how much a party has gone up or down relative to the last election matters more for understanding performance in the remainder of the cycle than knowing who was in government.

Change in Votes on Change in Polls since the last election

A more parsimonious way to allow the previous election to influence expectations for the next election is to measure both votes and polls relative to the previous election. Thus a relative change model would have the form:

$$\text{VOTE-PREV_VOTE}_{pti} = \alpha_{pti} + \beta_{pti}\text{POLL_SHARE-PREV_VOTE}_{pti} + \varepsilon_{pti}. \quad (3)$$

Whereas in equations (1) and (2) what is being predicted is the absolute share of the vote for a party, here the change since the last election predicted as a function of change in the polls relative to the previous election. Beta coefficients from models of this kind are typically close to 0.5 between twenty months to a year before the election. This tells us that parties are expected to lose roughly half their gains (or recoup half their losses) in the polls since the last election. As the election approaches the coefficient approaches one, but remains somewhat below one until the final week. Again this tells us that some of the reversal of fortunes is likely to occur late in the cycle.

Model (3) is actually a special case of model (2). If PREV_VOTE is subtracted from both sides then the right-hand side of the equation (3) becomes $\beta_{pti}\text{POLL_SHARE} + (1 - \beta_{pti})\text{PREV_VOTE}$, which is the same form as for model (2) but with a constraint that the coefficients of POLL_SHARE and PREV_VOTE sum to one. So in practice the LDV and relative change models are rather similar to each other for forecasting purposes and both have important differences from the first model with government status. As with the LDV model adding a government status dummy variable would not improve the model because government status is extremely strongly correlated with whether parties are polling more or less than at the last election.

This model is also formally equivalent to a regression of change in the final t days of on the cycle $\text{VOTE-POLL_SHARE}_{pti}$ on the change from the previous election up to t days out $\text{POLL_SHARE-PREV_VOTE}_{pti}$. So model (3) can be thought of as the most direct model of the electoral cycle effect whereby change in the remainder of the electoral cycle is modelled and forecast as a function of change in the first part of the cycle.

Forecasting from regressions always involves a regression (or reversion) to the mean effect (hence the name). In the votes on polls model (1) the mean to which a party reverts is essentially the historical post-1945 average for that party (Con 41, Lab 40 and Lib 14). In the relative change model (3) reversion is to the mean changes since the previous election (Con -0.14, Lab -1.12, and Lib 0.85). Since these are all quite close to zero the relative model predicts that parties will move back towards their 2010 result at the 2015 election. The LDV model is between the two but very close to the relative model in this respect because the previous election is such a powerful predictor of the next election. The practical implications of this are discussed further later in the paper. First it is important to consider which is the best for forecasting.

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 and which model would have worked best in the past at predicting elections. This section considers out-of-sample prediction, i.e. running regressions of the forms in equations (1 to 3) 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.

A further consideration is the 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 breakthrough 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 (1) 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 since and including 1974 on the basis that there are enough of these elections to judge performance, they are the most relevant ones and critically, they allow sensible comparability with models based on just 1974 onwards data.

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.⁷ 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.

⁷ 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 predictions

		<i>Time prior to the election:</i>		
		20 th	12 th	6 th
		month	month	month
(1)	Mean overall error in polls	9.6	6.8	6.8
	N (cycles for averaging)	9	9	10
Votes on Polls with Government Status				
<i>Forecasts with regressions from 1950-2010</i>				
(2)	Unweighted	7.5	7.2	6.1
(3)	Weighted	5.8	5.2	4.2
	N (cycles for regressions)	14	15	16
<i>Forecasts with regressions from 1974-2010</i>				
(4)	Unweighted	5.7	4.6	4.0
(5)	Weighted	6.2	4.5	4.0
	N (cycles for regressions)	8	8	9
Votes on Polls and Prior Election (LDV)				
<i>Forecasts with regressions from 1950-2010</i>				
(6)	Unweighted	5.9	6.2	4.2
(7)	Weighted	5.8	5.7	3.5
	N (cycles for regressions)	14	15	16
<i>Forecasts with regressions from 1974-2010</i>				
(8)	Unweighted	6.0	5.6	3.9
(9)	Weighted	6.4	5.8	3.9
	N (cycles for regressions)	8	8	9
Change in votes on change in polls since last election (relative change)				
<i>Forecasts with regressions from 1950-2010</i>				
(10)	Unweighted	5.4	5.2	3.7
(11)	Weighted	5.2	4.9	3.2
	N (cycles for regressions)	14	15	16
<i>Forecasts with regressions from 1974-2010</i>				
(12)	Unweighted	5.8	5.1	3.4
(13)	Weighted	5.9	5.2	3.9
	N (cycles for regressions)	8	8	9
(14)	Average of row (4) and (13) forecasts	5.3	4.5	3.2

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

The remainder of the table is organised first by model type, and then by whether the sample was restricted to elections since 1974 or not, and finally according to whether

the data were weighted or not. There are substantial differences between the best and worst prediction models, with one case of no improvement over the polls at all (12 months out row 2) and one case of error reduction by more than a half (6 months out row 11). The votes-on-polls-with-government-status models are clearly better done with elections only from 1974 (rows 4 and 5, better than 2 and 3), whereas the LDV and relative change models are not so sensitive to the baseline set of cycles. In general weighting makes little difference when regressions are based on the 1974 onwards set, but it matters for the full sample models.

Looking across the whole table there are two rows that stand out as having the lowest prediction error. Votes-on-polls-with-government-status models based on 1974 onwards unweighted data (4), and the relative change model based on a weighted full sample (11). The latter has the best on average over the three columns, but the former was marginally better for the six months out data. For this reason both these models are given particular attention for the remainder of the paper.

Although models in rows (4) and (11) produce similar predictions for past elections, either being close to the outcome or wrong in the same direction, as shown later in the paper they produce rather different point estimates for 2015, especially for the Liberal Democrats. So given their similar levels of success for past elections, it is reasonable to consider the average of the forecasts from row (4) and row (11) models, as shown in row (14). This average is only very little if at all better than the component forecasts and so it is not definitely to be strongly preferred to those from either component models. But without a strong preference for either the votes-on-polls-and-government-status or the relative change models, the average of the two is a useful compromise.

Before moving on to applications note that Table 2 suggests that the best models can reduce prediction error by about a third to a half of that from 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 approach 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.

Issues with applying the models to the 2010 to 2015 cycle

In addition to understanding how well the forecasting methods perform 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 the UK Independence Party (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 specifically. But their current poll standing will have an impact on the poll shares for other minor parties collectively and for the three main parties also, which in turn will affect predictions. Any differential effect that the party has on the poll share of another party at the national level will also therefore be reflected, as any change in the level of party support would do.

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 is 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 arguably be too strong a reversion-to-the-historical-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 when using equation (1) because it increases the range of Liberal polls and votes 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 from equation (1), even though the Conservative and Labour forecasts from equation (1) are best done based on regressions with 1974-2010 cycles. This is the basis on which the votes on polls and government status model is applied in the rest of the paper. This is also a problem for the relative change model in equation (3) since the Liberals have rarely deviated so much in the polls from their previous general election result as they have currently done, but it is not clear whether there are any consistent strategies for dealing with this problem within the relative change model.


Smoothing for day-to-day forecasting

A significant problem with daily forecasting from any of these models 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 close to 10% as they have been doing. Figure 1 shows two scatter plots of the raw forecasts from daily regressions for the 600 days up to the election. The left hand side uses equation (1) votes on polls and government status regressions while the right hand side comes from the weighted relative change model in equation (3). Most strikingly the former has a much lower forecasts than the latter until the very final few days. This is because the votes-on-polls model tends to drag the prediction to the historical average of 14% while the relative change model tends to push the forecast towards the last election result of 24%.

For both sets of forecasts there is lots of day-to-day variation. The greater noise for the relative change model is an undesirable feature and a reason to be cautious about its use. Also note that for both sets, over the first year the forecast trends downwards and then it goes back up again in the second year before falling in the final months. More generally, equivalent graphs for different parties at different levels show various monthly or longer trends and waves. It is possible to smooth over day-to-day variation easily enough with moving averages or other approaches, but the weekly/monthly and

longer-term 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. Using Figure 1a as an example, 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 thinks 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.

Figure 1 Forecasts of Lib Dem election-day shares given stable poll share of 10%



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 regressions for each t from 1 to say

600. Centre the data on the final forecast (from the day before election day)⁸, 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 from any of the models in equations (1-3) for party p , with poll share x at time t days before the election.

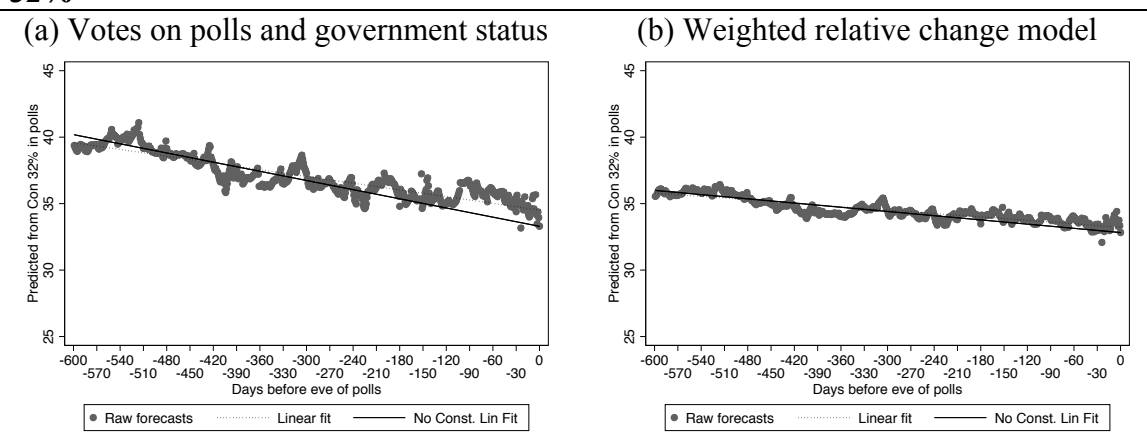
$$RAW_FORECAST_{pxt} - RAW_FORECAST_{pxl} = b_{px} * t. \quad (4)$$

This linear regression forces the forecast to progress steadily approaching 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,

$$IRONED_FORECAST_{pxt} = b_{px} * t + RAW_FORECAST_{pxl}. \quad (5)$$

Figure 1 shows both a regular OLS fit and the result of the no-constant regression on centred data. In Figure 1a this case there is very little difference between the no-constant regression over the regular OLS which ends very close to the final prediction. But there is a massive difference for Figure 1b because the regressions continue to put heavy weight on reversion to the previous election result until the final three months. Figure 2 below shows an example with the Conservatives on 32% (reflecting current public opinion) with a starker difference. There is some but very little deviation in both panels. Where there is a difference between the regular and no-constant regression, the latter 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.

Figure 2 Forecasts of Conservative election-day shares given stable poll share of 32%



The two models differ importantly in the scale of the predicted Tory recovery. Over 500 days out, the votes-on-polls-and-government-status model predicts an eight point gain by the election, but the relative change model forecast is for half that. Both models suggest that if they continue to poll at 32% their expected election share of the vote

⁸ 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.

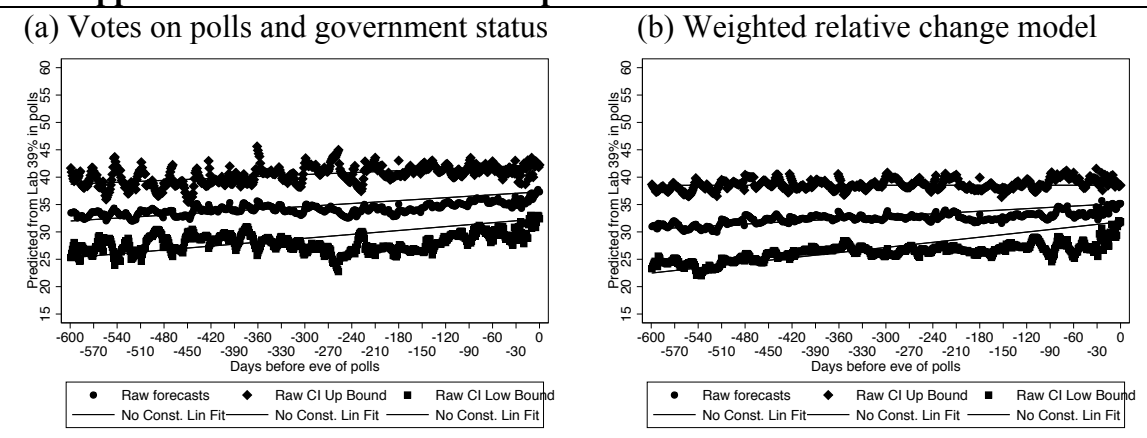
declines to around 35% or 34%. This final two or three point difference reflects their historical tendency to out perform the polls in the final twenty months, including but not mainly from polling industry bias.

Prediction Intervals for Shares of the Vote

Prediction/forecast standard errors from each votes-on-polls regression (equations 1-3) 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 ironing out with a linear trend is a good approximation most of the time.⁹

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, for both the votes-on-polls-and-government-status and the relative change models. Unlike for the Conservatives and Liberal Democrats, the long-range forecasts for these two models for Labour differ only by a point. The width of the prediction interval for Labour narrows a little but not much. With the exception of the Tories in the relative change model only, the other two parties have wider intervals to start with and they attenuate more over the period. These party differences 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.

Figure 3 Forecasts of Labour election-day shares given stable poll share of 39% with upper and lower bounds for 95% prediction intervals



While most of the time the ironing seems to do a good job there are some sustained periods where ironed prediction intervals are noticeably different from the raw intervals.

⁹ 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.

For Labour the smoothing arguably narrows the prediction interval too much between 90 and 300 days out in Figure 3a and 30 to 210 days out in Figure 3b. For the Tories the smoothing widens the prediction interval too much between 30 and 240 days out for both models. 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.

Predicted seat totals and approximate prediction intervals

Predictions for seat totals for each party can be generated using the method 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

$$r_{ij} = \exp \left[- \left(\frac{\max_j \{s_{ij}\} - s_{ij}}{\sigma} \right)^\lambda \right], \quad (6)$$

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 estimates for the shares of the vote with some assumptions about the nationalist and minor parties and special seats.¹⁰

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. For example, 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:

¹⁰ 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.

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

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

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.

Predicted probabilities of overall majorities, hung parliaments and largest party

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. To enable consistency with the probability of being the largest party calculations, these party specific standard deviations are averaged to provide a single standard deviation to reflect uncertainty in predicted seats total for both parties.¹¹ Given the normal distributions for both parties 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

$$\text{Pr(Hung parliament)} = 1 - \text{Pr(Conservative majority)} - \text{Pr(Labour majority)}. \quad (7)$$

¹¹ 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 given that with a fixed estimate for the Liberal Democrat and Other party share the variation in seats is overwhelmingly between Conservative and Labour and the probability calculations are relatively insensitive to the estimate of the standard deviation within a reasonable range.

The approximate probability that the Conservatives will be the largest party in parliament is estimated as the probability that the predicted number of Conservative seats exceeds the average of the point estimates for Conservative and Labour seats. Thus,

$$\Pr(\text{Con seats} - \text{Lab seats} > 0) \approx \Pr(\text{Con seats} - (E(\text{Con seats}) + E(\text{Lab seats}))/2 > 0), \quad (8)$$

where $E()$ is the expectation function and so refers to the point estimates for seat totals. Again this is calculated using the pooled standard deviation for seats and the cumulative normal distribution. The corresponding approximate probability that Labour will be the largest party is the symmetrical opposite and so by design is equal to one minus the approximate probability that the Conservatives have the most seats.

Example forecasts with discussion

Table 3 shows three sample forecasts calculated with the above methodology for a 7th May 2015 election from 8th October 2013, at which point the average poll shares from the www.ukpollingreport.co.uk polling average were Conservative 32%, Labour 39%, Liberal Democrats 10%, Others 19%. The first forecast is based on votes-on-polls-and-government-status regressions and the second forecast comes from on weighted relative change votes-on-polls regressions. Column (c) forecasts are based on the average of the daily forecasts and prediction intervals from the first two, after the application of the methodology for smoothing, seats predictions and probabilities. So while close, the figures in the third column are not the straight average of those in the first two. They deviate because of the smoothing process and then the implications of that for later calculations.

Both basic forecasting models (columns (a) and (b)) suggest a Conservative lead in the share of the vote over Labour despite the inputted polling average of a seven point Labour lead: an almost perfect switch in relative standing. This is perhaps the most important feature of both forecasts despite the differences between them.

The most striking difference between the two basic models is the gap in the forecast Liberal Democrat vote share, which has a point estimate of 10.9 from the votes-on-polls-and-government-status models but 20.3 from the relative change model. This is because the Liberal polls have historically been fairly weak predictors of the eventual Liberal vote this far away from an election. So both models have very strong regression to the mean effects driving their forecasts, but with regression to different means: to the historical average (14%) with votes-on-polls-and-government-status models and to the last election (24%) for the relative change model.

This difference has a big impact on the implied share of the vote for Others, which is expected to fall by just two points from current polls in the first column but by six points in column (b). Given that nearly all the variation in the Other share since the last election has been due to the rise of UKIP, these forecasts could effectively be read as suggesting that the UKIP vote will fall by about these amounts before the next election. So the relative change model is suggesting UKIP will lose about half their current support in the polls. But note that neither model is suggesting that UKIP will win any seats even though the calculation assumes that all of the predicted rise in the combined

other share of the vote will go to the largest minor party¹² from each constituency in 2010.

Table 3 Example forecasts from 8th October 2013, 576 days till the election

	a. Votes on polls and government status	b. Weighted relative change model	c. Forecast from average of a and b models
Inputted current average poll shares			
Con	32	32	32
Lab	39	39	39
LD	10	10	10
Others	19	19	19
Forecast Election Day Shares with 95% Prediction Intervals			
Con (36.9 in 2010)	39.9 +/- 12.0 i.e. between 28 and 52	35.9 +/- 7.3 i.e. between 29 and 43	37.9 +/- 9.7 i.e. between 28 and 48
Lab (29.7 in 2010)	32.2 +/- 6.6 i.e. between 26 and 39	30.7 +/- 7.8 i.e. between 23 and 39	31.5 +/- 7.2 i.e. between 24 and 39
LD (23.6 in 2010)	10.9 +/- 15.7 i.e. between 0 and 27	20.3 +/- 9.2 i.e. between 11 and 30	15.6 +/- 12.5 i.e. between 3 and 28
Implied point estimate for Others combined:	17.0	13.1	15.0
Forecast Election Day Seats			
Con (307 in 2010)	333	295	315
Lab (258 in 2010)	270	276	275
LD (57 in 2010)	19	51	33
	Con majority of 16	Con largest party, but short of a majority by 31	Con largest party, but short of a majority by 11
Forecast Election Day Seats with approximate 95% Prediction Intervals			
Con	216 to 468	203 to 402	211 to 436
Lab	143 to 381	176 to 364	162 to 372
LD	13 to 28	44 to 58	25 to 42
Approximate probabilities of key outcomes (%)			
Pr(Con majority)	54	27	42
Pr(Lab majority)	19	16	18
Pr(Hung parliament)	27	58	40
Pr(Con largest party)	69	58	64
Pr(Lab largest party)	31	42	36

Notes: Seats intervals assuming LD and Other shares fixed at point estimates while Con and Lab vary as per prediction intervals above.

The predicted seat totals depend mostly on the predicted shares for the three main parties. The higher Liberal Democrat vote share estimate and the lower predicted Tory lead from the relative change model both contribute to a prediction of a hung parliament with Conservatives the largest party, instead of a small Conservative majority from the votes-on-polls-and-government-status models.

This is certainly a big political difference between the two models. But notice that the prediction intervals for shares of the vote for all three parties are very broad and mostly overlap. At first glance the combined prediction intervals may seem to encompass all foreseeable outcomes and more. For the Liberal Democrats the interval for the average

¹² Not including SNP or PC.

forecast from 3 to 28 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, conditional on the model being appropriate. 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 intervals suggest it is extremely unlikely that Ed Miliband could do as well as Tony Blair in 1997 or 2001, but there appears to be a fair chance he could do worse than Gordon Brown or Michael Foot.

The forecast election-day seat totals 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 either a 69% or 58% chance of being the largest party but only 54% or 27% chance of an overall majority respectively from the two basic models. At this stage in the cycle part of the 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 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).¹³ The two different models help to show just how, even so far out from an election, the probability of a hung parliament is very sensitive to the scale of the Liberal Democrat recovery. With the relative change model a hung parliament is the most likely outcome with a probability of 58%, but note that this assumes that the Liberal Democrats will definitely be on 20.3% of the vote.

As discussed in the section on out-of-sample predictions, in the past the two basic models would typically have pointed in the same direction and both been either very close to right or wrong in the same direction for each party, for 2015 the models point in two very different directions for the Liberal Democrats. So it is reasonable to either pick the one you think is methodologically preferable or take the average between them, as in the final column of Table 3. The average of the two has not been substantially better as a predictor in the past, but without a strong reason to prefer one or other of the two more basic models the average of the two is a natural choice.

Methodological discussion

Various methodological issues have been discussed along the way in this paper. This section serves to revisit some of them and introduce new issues for consideration. The

¹³ 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.

discussion is here to provide further information to aid understanding of the properties of the data, models, forecasts and the thinking behind them, but also to deal with potential criticism and highlight issues for further research.

Using the full range of cycles for votes-on-polls-and-government-status 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 column (a) of Table 3 would have been lifted to 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-and-government-status 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 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 and different formulations of the votes-on-polls and similar 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.

More substantively it would also be good to see more research on the extent to which electoral cycles are linked with other political, social and economic phenomena. In

particular, whether or not the tendency for governments rise and fall over the electoral cycle may be sustained in part by the economic management in which governments are disproportionately more likely to stimulate the economy or provide tax cuts at the end of the electoral cycle.

Political, social and economic phenomena also have played a part in determining when prime ministers have chosen to call elections in the past and so the previous electoral cycles have been of different lengths for different reasons which are linked to the success or failure of the electoral cycle. Traditionally governments that are likely to secure re-election call them after four years but those more likely to lose go the full five years. Maybe this strengthens the measured electoral cycle effect and maybe then we should expect it to be weaker in the current cycle than in the past. If so this is reason to suspect that the outcome will be closer to current opinion polls than any of the models. The government status as a predictor in the votes-on-polls regressions has a big effect on long-range forecasting, boosting estimates in the first column of Table 3 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. So without a government recovery and opposition setback effect in the model the forecast would be for the two main parties to be neck-and-neck.

For the example forecasts I have used just one possible way of computing a polling average for current public opinion, but there are others. The technology can be used to generate forecasts for different scenarios so predictions from 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. Governments are likely to rise then fall in the polls over the course of the electoral cycle, and *vice versa* for the principle opposition party. More generally parties' gains or losses in the polls over the first half or so of the electoral cycle are likely to be undone somewhat by the next election. Even accounting for these features, historical experience suggests that the Conservatives are likely to do better at elections than they do in the polls twenty months prior, and *vice versa* for Labour. This is true, but to a diminishing extent, right up to the final week. A very small amount of this looks like either very late swing or polling industry bias.

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. But really big changes are relatively rare. This helps to explain how forecast confidence intervals are very broad without becoming uninformative. While the Conservatives, in October 2013, were clearly expected to be the largest party at the election, there is still a 36% chance that Labour would be according to the average of the two best models identified in this paper.

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 is relatively high, at 40% in the average of the two best models and as high as 58% in a model that emphasises reversion to the last election result.

This 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.

Of course this time might be different, and there are major features of the current political landscape to support such a view, not least the coalition and the rise of UKIP. While the methodology allows these developments to be somewhat reflected in the forecasts, there simply is no evidence base for proper modelling. Substantial though these changes are, there were exceptional events and developments in many of the other election cycles. Major party performance dynamics will only really be different this time if the outcome falls outside the prediction intervals, which we will not know till after the election.

Acknowledgements

I am very grateful for comments on previous drafts to John Curtice, Jouni Kuha, Clive Payne, Mark Pickup and Chris Prosser, and to Rob Ford, Will Jennings 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). Finally I am also very grateful for the numerous and wide ranging comments on Twitter and elsewhere online in response to previous drafts. The paper has been revised substantially in part thanks to them.

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.

- Brown, Lloyd B, and Henry W Chappell Jr. 1999. "Forecasting presidential elections using history and polls." *International Journal of Forecasting* 15(2): 127–135.
- 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.
- Erikson, Robert S, and Christopher Wlezien. 1996. "Of Time and Presidential Election Forecasts." *PS: Political Science and Politics* 29(1): 37–39.
- 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*. 30(2): 250-257.
- 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 Borisyuk, 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.
- Selb, Peter, and Simon Munzert. 2013. "Forecasting the 2013 Bundestag Election Using Data from Various Polls." : 1–13. https://www.researchgate.net/publication/255991948_Forecasting_the_2013_Bundestag_Election_Using_Data_from_Various_Polls?ch=reg&cp=re243_x_p5&pli=1&loginT=trauldVzfxjjmr6kpPy8TUegmJ278kamvq7SwvpU18YdCRlhb_XVoyXp7XIWVQCF (Accessed January 6, 2014).
- Wlezien, Christopher, Will Jennings, Stephen D Fisher, Robert Ford and Mark Pickup. 2013. "Polls and the Vote in Britain." *Political Studies* 61: 66–91.