

FORECASTING THE 2015 BRITISH GENERAL ELECTION:
THE 1992 DEBACLE ALL OVER AGAIN?

Stephen D. Fisher
stephen.fisher@trinity.ox.ac.uk
Trinity College, Oxford, OX1 3BH, UK
Department of Sociology
University of Oxford

Michael S. Lewis-Beck
Department of Political Science
University of Iowa

ABSTRACT

This article introduces and reviews a set of twelve academic forecasts of the 2015 British general election. Along with the vast majority of others including journalists and betting markets, they failed by a big margin to predict that the Conservatives would emerge with an overall majority of seats. Several suffered from the 1992 scale inaccuracies of the vote-intention opinion polls. Forecasts based on other data sources typically did a bit better, but also fell short. Nonetheless, this was not 1992 all over again. The dramatic collapse of the Liberal Democrats and rise of the SNP, UKIP and Greens were successfully anticipated. Also this collection includes numerous methodological advancements, with several new methods and developments to established approaches.

Keywords: Election forecasting; UK elections; Hung parliament; Prediction models; opinion polls

Highlights:

- Twelve forecasts of the 2015 British general election are introduced and reviewed.
- Some of the reasons for failure to forecast a Conservative majority are elucidated.
- Issues for future election forecasting, in Britain and elsewhere, are discussed.

INTRODUCTION

In Britain, scientific election forecasting has been a serious enterprise for some time, at least since the 2001 general election. Its zenith was reached with respect to the 2010 general election, when academic modellers unanimously foresaw a “hung

parliament” for that contest, something that many pundits scoffed at. That triumph of the scholars was documented in a special issue of *Electoral Studies*, which published a collection of papers on these successful ex ante forecasts, by six different teams. [See the critique and summary of these papers in Gibson and Lewis-Beck (2011).] The journal, *Electoral Studies*, has again graciously agreed to host the publication of papers forecasting, before-the-fact, the outcome of the 2015 general election in Britain. That collection is at hand, and records the predictions of twelve different teams, double the number for 2010.

At first blush 2015 appears to be the nadir of the forecasting enterprise, for each of the teams again foresaw a “hung parliament”. Of course, the election itself proved this forecast wrong, giving the Conservatives a solid majority of seats, so enabling them to rule alone.

That this same prediction – a “hung parliament” – served the forecasting community so well in 2010, and so badly in 2015, holds considerable irony. The forecasters were not alone in making this big mistake, however, as the pollsters and the media also chimed in with their own false predictions. In their February survey, Hanretty and Jennings (2015) found academics, pollsters and journalists had very similar views on the likely outcome in both votes and seats, with just 6% suggesting one party would emerge with a majority. Throughout the year the betting markets had a hung parliament as overwhelmingly likely, with just a 6% implied probability of a Conservative majority in the final week (Tapper 2015).

The situation is reminiscent of 1992. That was hitherto the last time the campaign polls suggested a close race and it also was the last time the Conservatives had won a

majority. Inevitably many of the forecasting teams fell foul of the error in the Conservative and Labour vote intention polls. But there is much more to election forecasting models now than just extrapolation of opinion polls. And there was much more to the 2015 election than the contest between the two main parties. The Scottish National Party (SNP) took nearly all the seats in Scotland on just under 50% of the vote (up 30 points on 2010). Meanwhile both UK Independence Party (UKIP) and the Greens increased their vote four fold and the Liberal Democrats lost nearly two-thirds of theirs. These developments were anticipated by most opinion polls and forecasters alike. Summing up the pre-election consensus view is a lead header from *The Economist* (May 2nd – 8th, 2015, p.48): “Next week Britain goes to the polls in its strangest, closest general election for many years.” They were right about the strangeness if not about the closeness.

Below, we first characterize the different papers and their approaches, highlighting just some of the numerous methodological developments. Then, we assess their vote share estimates, before we turn to assessment of the seat share estimates. We conclude with reflections on some of the potential lessons for the future of election forecasting in Britain.

THE PAPERS

These dozen papers were submitted to the editors *before* the May 7th election day. Therefore, the forecasts were on record in advance of the contest. In other words, they are true forecasts, studied *ex-ante* predictions. While the science, not to say the art, of

election forecasting stands well-advanced in Britain, it is not unique to those islands. The American community of election forecasters has also been quite active, most notably in the 2012 United States presidential election, which saw the return of President Obama. After that contest, *PS: Political Science and Politics*, published comments from sixteen different forecasting teams, each assessing their performance. [See Lewis-Beck and Stegmaier (2014), for an introduction to that symposium.] The approaches employed there varied, but could be grouped according to their reliance on vote intention polls, explanatory structural modelling, political stock markets, or subjective expert judgments. The British work here draws on the vote intention and modelling approaches, as well as adding several other strategies.

Four papers, primarily using vote-intention polling and survey data at national, sub-national and constituency levels, attempt both to forecast vote share at the national level and estimate variation in party performance between constituencies. In general, these efforts move forward in analysis steps. For example, Fisher looks at average vote intentions, then at a regression models that includes vote intention and vote history. He then simulates hypothetical election results which are adjusted using estimates from constituency and individual-level polling data to eventually arrive at seat predictions with uncertainty estimates. Using similar ingredients but a much more formal and thorough method, Hanretty et al. present very thoughtful, creative and technically impressive solutions to a range of modelling problems, even before tackling the most problematic issue of reconciling constituency and national level polling information. In another methodologically impressive approach, Ford et al. build on their polls-to-votes-to-seats model from 2010 in various ways, most notably creating a smoother evolution of the

polls-to-votes part and by a complex elaboration of the votes-to-seats model incorporating more parties' various constituency level factors. Mellon and Fieldhouse exploit the unprecedented availability of large-scale individual-level British Election Study (BES) data with full constituency coverage in order to estimate constituency-level patterns in the flow of the vote, correcting for turnout and aggregate characteristics.

Another four papers offer structural models at the national level, using time series data. The Lebo and Norpoth pendulum model has worked well for predicting Conservative and Labour seat totals over decades. The idea is that voters swing between choosing Conservative and Labour governments roughly every 2.5 terms. This time they made adjustments for some of the Prime Ministerial approval coming not from the Conservatives but their coalition partners, and for a different pattern of party performance in Scotland. Stegmaier and Williams develop a model for the pattern of change in party popularity over time based on the Continuous Monitoring Survey (CMS) from 2004 to 2015. For independent variables they include prior party vote, and national economic and political evaluations. Among other interesting findings, they show that economic evaluations affected Tory but not Liberal Democrat support since 2010. Lewis-Beck et al. synthesize economic, political, and vote intention variables into an aggregate dynamic model predicting incumbent vote share, showing how vote intention improves and economic growth declines in predictive power as the election approaches. Whiteley et al. build on a long tradition of estimating the votes-seats relationship in British politics and develop a novel approach to the problem, taking into account the implications of the massive drop in Liberal Democrats support following the coalition formation.

Local election results are the sole basis for prediction in two of the models. Prosser develops a new model for forecasting from the national equivalent/projected local-elections vote share, with intriguingly similar forecasts from both the 2014 and 2013 rounds of local elections. Rallings et al. develop their local by-elections model, that has been successful since 1997, to include UKIP, allow for different developments in Scotland, and identify constituency variation based on their component ward-level local election results.

Finally, two approaches form classes of their own. Murr shows the efficacy of citizen forecasting (asking people who will win in their constituency) for the past seven British elections, and develops a new method for predicting constituency vote shares from categorical forecasts by citizens of who would win in their constituency. Burnap et al. forecast the election outcome using Twitter. This is a first for Britain. But they take lessons from previous attempts in Germany and Italy, using a sentiment analysis of Twitter mentions of parties and leaders up to two months in advance.

These forecasts are considered below. We begin with a look at the vote share forecasts, then turn to the decisive seat totals forecasts.

THE VOTES AND SEATS FORECASTS: AN OVERVIEW

Table 1 presents the forecasts for the seven parties that won seats in Great Britain, together with the actual vote shares and seat totals for these parties. Forecasters varied in which outcomes they chose to forecast. Sometimes these were matters of choice, influenced by what the teams thought they could adequately forecast with their

methods. Sometimes there are structural reasons for absences, e.g. when methods by-pass share forecasts to predict seats directly. With some of the minor or newer parties forecasts were sometimes made, if at all, on a different basis from larger parties, not least because of the lack of adequate or any historical track record of polls or other pre-election indicators. Details on these choices are in the individual papers.

In order to facilitate comparison, we provide in Table 1 the actual vote and seat results by party. As a baseline, note that the Conservative party, leader of the government coalition, received, including that of the Speaker, 331 seats (+24 since 2010). Their coalition partner, the Liberal Democrats, received only 8 (-49) seats, a sum not needed by the Conservatives in any case, as they obtained their own seat majority. The chief rival, Labour, received only 232 (-26) seats, thus trailing the winning Conservatives by just short of 100 seats.

Also included is the average of the final vote-intention polls that were published before the election, and a uniform change seats projection from those polls. This projection incorporated the average of the Scottish polls and adjusted the GB polls accordingly. Although opinion polls should not be judged on the basis of any seats projection that may be generated from them, this nonetheless provides a helpful point of comparison.

[TABLE 1 ABOUT HERE]

Look at column 1, which contains the vote share forecasts. Conservative vote share was consistently underestimated, regardless of the forecaster. (Closest were three

teams, each within 2 percentage points, with an estimate of 35%). The average error was large and about the same as that for the opinion polls, minus 4 percentage points. In general, something seems clearly wrong with the Conservative forecast here. For the other two parties, in contrast, things look more reasonable. Four models called the Labour vote share correctly, within one percentage point. And both the mean and the median were just 0.8 points out; less than half the error of the opinion polls. With regard to the Liberal Democrats, we see the same general picture: four models are within one percentage point, and the average error, although bigger than that for the opinion polls, is under 2 percentage points. This is remarkably successful given the party lost two thirds of their vote share between 2010 and 2015.

Relatively few teams attempted forecasts of the other parties but those that did by and large did not do too badly in identifying the big changes that were to occur. Figures for the SNP are difficult to interpret at the GB level, especially after rounding. The party rose from 20% to 50% of the Scottish vote. One team overestimated and two underestimated this with the middle one coming close on 48%; just as the Scottish opinion polls. The biggest change in GB vote share was for UKIP, who went from 3% in 2010 to 13% this year. Again, given the difficulties of estimating such a dramatic and unprecedented development, it is remarkable that three teams and the opinion polls came within a percentage point. Also the Green "surge" from 1% to 4% was slightly overestimated by the opinion polls but well captured by one of the two teams that forecast it.

Now turn to the seats forecasts, in column 2. With respect to the Conservative seats forecast, the closest is 296 seats, by Prosser, but that falls well short of a majority.

Indeed, the average Conservative seat forecast, of just 279 seats, is 52 seats shy of the actual result. The Labour forecast is barely any better, with an average error of 46 and even the closest still 30 seats too high. When we inspect the Liberal Democrat forecasts, things look better in absolute terms: the average error was just 17 seats. But this means that the Liberal Democrats were expected to win over three times as many seats as they actually took. All the models (with the exception of one) are grossly off in the positive direction. Conversely, all but one of the models correctly predicted a dramatic rise in the number of SNP seats, with three forecasts within three seats of the eventual result. Predicting precisely the small number of seats for the remaining three parties was generally well done by those that attempted it. In predicting that UKIP would win more than 10% of the vote but take less than a handful of seats the forecasters rightly identified an important feature of the operation of the electoral system at this election.

In sum, despite some positive points, the models taken together did poorly. Focusing just on seat totals, after all the crucial measure, these models on average saw a parliament where the Conservatives and Labour were in virtually a dead heat (279 seats v. 278 seats), with no majority in sight. We turn to the question of whether any of the individual models, when put under the microscope, do much better.

THE VOTES AND SEATS FORECASTS: WHY DID SOME DO BETTER THAN OTHERS?

For all of the methods, apart from Murr, the seats based forecast is highly dependent on a forecast of the overall share of the vote. The most important aspect of the share forecast is the difference between the Conservatives and Labour: seats totals are

more sensitive to this aspect of the vote share than any other. In general, those with a larger Conservative lead over Labour typically did better on their seats forecast, effectively or actually because they forecast fewer net Labour gains from the Tories.

So overwhelmingly the main problem for the large majority of the forecasting teams was the failure to even come close to accurately predicting the eventual 6 point Conservative lead over Labour. Across the 10 teams who presented a forecast of the Tory lead over Labour on votes, the average was just a one-point lead and the best was just 4 points; still well short.

For those primarily using vote-intention opinion polls the immediate problem was the failure of those polls. This affected similar academic survey data (BES and CMS) too. There were differences between teams in how they thought that polls would translate into future vote shares, which led to noticeable differences in forecasts of the Tory lead. Those expecting swing-back to 2010 shares or historical averages did a bit better in this respect, but correspondingly less well for those parties where the polls predicted big changes accurately. Only the Twitter forecast overestimated the big changes for LD, SNP and UKIP. By underestimating change in general the forecasters did slightly better on average than the final polls at predicting the eventual vote shares for the two main parties, but less well than the polls for the other parties. Indeed the opinion polls did such a good job for the smaller parties (within two percentage points for each) that they were hard to beat.

But it was not just the polls that had problems anticipating a big win for the Tories. The problem was also present for those using local election results, Twitter data or structural models based on economic growth and/or other data. Apparently none of

these data types provided enough prior evidence of a large Conservative lead. Or, when the different sources were combined, none were sufficient to overwhelm the problems of vote intention. Although vote intention based methods did not always do worse at predicting the Tories than those using other kinds of data, they typically did, and the three best forecasts for the Conservative seat total avoided vote-intention data altogether.

The one indication in these papers that the big Tory lead was predictable comes from Lebo and Norpoth who had a January forecast which predicted a 7.1 point Tory lead from an unadjusted model (close to the actual figure of 6.5). Their attempt to account for the effects of the Liberal Democrats being in government on PM approval made matters worse. This speaks well for their pendulum model and badly for *ad hoc* adjustments (made by several teams) to cope with the apparent uniqueness of the current election.

But even with a 7 point Conservative lead, Lebo and Norpoth were predicting the Conservatives to be short of a majority. Other forecasters have essentially argued that, even with the right shares of the vote, the Tories outperformed uniform change (and other data driven prior estimates of constituency variation) in key marginals (Hanretty et al. 2015). This looks to have been worth about ten seats to the Tories and is related to first-term incumbency (sophomore surge) effects.

The second main source of error for most of the forecasts was overestimation of the number of Liberal Democrat seats. Lewis-Beck and his colleagues got much closer to the actual outcome than others by having an accurate estimate of the Lib Dem share and a relatively high (and so better) Conservative share forecast, and by ignoring the constituency polls that pointed to disproportionate success of Lib Dem incumbents. The main problem for most other forecasters was underestimation of the Conservative lead

over the Liberal Democrats. Had this been done correctly, forecasts on the basis of uniform change or constituency polls (appropriately adjusted) would have pointed to similar single digit seat tallies for the Lib Dems. So again it seems that inaccurate overall share forecasting was the main problem.

However, while uniform change from the correct vote shares pointed to practically the right seat totals, it was the wrong model. There were 152 constituencies where the Liberal Democrats could not have lost 15.5 points implied by uniform change because they had not won as much as 15.5 points in 2010. So there is an important question as to how well forecasters did in estimating the particular pattern of Liberal Democrat decline as opposed to benefiting by chance from an oversimplified model of constituency level change that seems to have worked for the wrong reasons on this occasion.

One of the main sources of evidence for estimating constituency variation were constituency polls, which were conducted on a large scale in Britain for the first time, principally by Lord Ashcroft. A key feature of these polls was their constituency vote choice question, which prompted respondents to think, "specifically about your own PARLIAMENTARY constituency at the next General Election and the candidates who are likely to stand FOR ELECTION TO WESTMINSTER there." (Emphasis in original.) This question was also fielded by the BES and used by Mellon and Fieldhouse. While experience of the 2009 Politics Home poll and subsequent 2010 election provided some limited support for the efficacy of this survey question, at first glance the 2015 experience seems to undermine it. For those that used them, constituency polls and the constituency specific question seemed to have made their forecast seat totals worse

overall by pushing up the Liberal Democrat figure. But this is conditional on their mistaken estimates of overall vote shares. Hanretty et al (2015 and in this collection) argue that the survey question and the constituency polls more generally picked up broadly the right pattern of constituency variation once properly adjusted for changes in national levels of support. Much to their credit, they were the only ones to publish constituency forecasts from a model incorporating constituency polls. Certainly they came a lot closer to the true pattern of constituency change for the Liberal Democrats than the default alternative uniform change, or similar alternative such as Martin Baxter's (2007) strong transition model.

Also the constituency polls provided pretty much accurate information showing that the SNP would beat uniform change by taking more from Labour in their former heartlands, that UKIP would struggle to win more than one seat and that the Greens were unlikely to make any gains. Understanding how to use such information, and other constituency-level data will perhaps be increasingly important for future forecasts as the number of marginal seats has declined (Curtice 2015) and so accurately predicting what happens in them becomes especially important.

The one method, from Andreas Murr, that forecast seat outcomes directly without appeal to any national vote share produced the best forecast for the two main parties. However, his experience is not necessarily positive in the lessons it teaches us. That forecast appears to have benefited from a status quo bias in citizen forecasting that exists in the US (Murr 2015). His British survey respondents predicted less change in Conservative and Labour held seats than our other academic forecasters expected from their methods. Moreover Table 2 from Murr's paper shows a general tendency for citizens

to under-estimate change in seat totals at all elections since 1997. This helped Murr relative to other forecasters on the two main parties, but also led to serious under-estimation of the Lib Dem collapse and the SNP rise. Nonetheless his strategy has full transparency, a straightforward application, a commendable theoretical base (in the Jury Theorem of Condorcet), and strong track record in Britain and elsewhere.

Finally, returning to the three best forecasts for the Conservative seat totals (Prosser, Murr, and Lebo and Norpoth); not only did avoiding vote intention polls help them, but also they do not seem to have been hurt by substantial lead times (12 months, 2 months and 2 months, respectively). Indeed in general those forecasts closer to the election were not systematically better than those further from it. Avoiding the use of vote intention measures, plus employment of a decent lead time, seems to have worked well, too, in the forecasting of the 2014 American congressional elections (Lewis-Beck and Tien, 2015).

CONCLUSION

The triumph of the Conservatives, returned to government with a majority of seats, came as quite a surprise. Should it have been?

The irony of forecasting, as John Curtice has noted, is that we assume the future will be like the past. Historically unprecedented or rare events are thus difficult to forecast. This election looked strange and difficult to predict beforehand. Afterwards it looks equally strange and difficult to explain, for some of the same and some different reasons. The most extraordinary and genuinely new developments, the collapse of the Liberal Democrats and rise of the SNP, UKIP and Greens, were remarkably accurately

estimated by the opinion polls. So much so that the forecasters' main trouble with these parties was having methods which, in effect, assumed that the polls would be exaggerating change when they were not.

But the main trouble was still with the traditional competition between the two main parties. Just as in 1992, that was supposed to be too close to call. Now the puzzle is not only how the polls and forecasters could have been so far out, but how could the election outcome have been so decisive. How did David Cameron increase his party's share of the vote when no other Prime Minister since 1900 that was in power for more than 18 months has managed it, and after five years of practically no economic growth? How did Labour manage to win even fewer seats than in their dire post financial crisis 2010 result? These are substantive questions, but they highlight how the actual outcome was historically surprising, and so from some approaches difficult to predict.

Nonetheless, polling data on 2015 vote intentions, which ought to come close to any election result, no matter how strange, led to serious forecasting error for the Conservatives and Labour. One obvious question is whether the problems of the polls could or should have been accurately predicted ex-ante. Whiteley et al in their postscript show how adjusting survey data to reflect likelihood of voting improves estimates (see also Whiteley and Clarke 2015). Many pollsters made similar adjustments but still fell short. This and other issues, such as the possibility of late swing, are being investigated by the British Polling Council/Market Research Society Polling Inquiry (<http://www.ncrm.ac.uk/polling/>) which will be making recommendations which hopefully help improve both polling and academic practice.

After the 1992 debacle there was also a polling inquiry. Their report identified methodological problems were partly but not completely to blame. Despite efforts to improve, the polls still managed to overestimate Labour and underestimate the Conservatives at the 1997, 2001 and 2005 elections, and underestimated the Tories again in 2010. The errors were smaller than in 1992 (or 2015) but might still have been consequential in a close election.

So there will remain a question at the time of the next election as to whether it is possible to improve estimates of the likely errors in future vote-intention. Matt Singh (2015) argued shortly before the 2015 election that there was considerable evidence, from local elections, leadership ratings and evaluations of economic context, pointing to the possibility of a large polling bias in vote intention of the kind that occurred. Similarly, Lewis-Beck and his colleagues rightly argue in their postscript that future forecasting models would do well to incorporate, and synthesize, more of this kind of information.

How best to do so will be a considerable challenge. As noted above, several forecasters were already drawing on other sources of information in addition to vote intention, but their analysis of the historical experience suggested that vote intention and similar party support data should be weighted more highly. It is not enough after the fact to point to the indicators that suggested a Conservative win; it is also necessary to show that there was a good basis before the election for preferring such evidence to vote-intention data. There was no shortage of arguments made before the election to explain why Labour were ahead in the polls despite Ed Miliband's poor leadership ratings and why the Conservatives had suffered a voteless economic recovery (Dale 2014). A key

question for future forecasting is whether, in light of this and previous elections, we now have good justification for placing less weight on vote intention data.

Perhaps we also need to look further afield for yet more and different kinds of data. After all, none of the forecasters came within 30 seats of the eventual tallies for either the Conservatives or Labour. None of the forecasters here provides any indication of how they could have predicted a Conservative majority in 2015 without better or different data.

Hanretty et al (2015, and in their postscript here) argue that polling industry bias is a moving target and so they see more potential to better estimate the uncertainty in rather than the direction of future polling errors. This too would help. One of the most positive developments in election forecasting in Britain this year was the increased use and publication of uncertainty estimates, including forecast probabilities for key outcomes such as majorities for particular combinations of parties. However, these estimates of prediction uncertainty were clearly overconfident, not least because they often did not include the actual result. Only Prosser put more than a small probability on the event of a Conservative majority. But this was largely because of high levels of uncertainty; he also suggested a sizeable chance of a Labour majority too.

There is room for improvement in uncertainty estimation and it would be worthwhile. Although much of the estimation of probabilities of different kinds of governments forming seems now to have been misguided, done well in the future it could make very important positive contributions for voter decision-making. This does of course imply forecasting of party behaviour as well as voter behaviour.

Another challenge to relish is that of lead time. Looking at the forecasts of these dozen teams, we see that making a call closer to election day did not really seem to improve accuracy. On the one hand, that could be lamented. On the other hand, it could be taken as a spur to recalling that, after all, forecasting is about distance, not closeness. The farther an accurate forecast can be made from election day, the more impressive. We might consider, then, putting more energy into figuring out the optimal trade-off between distance and accuracy.

ACKNOWLEDGEMENTS

Thanks for comments to Eilidh Macfarlane and to participants at a conference in March 2015 at the London School of Economics and Political Science (LSE) where earlier versions of the papers in this collection were presented. Thanks also to Simon Hix, Tony Travers and others in the LSE Government Department for funding and supporting the conference. Most especially thanks to the anonymous reviewer for their excellent comments on all the papers in this collection as well as on this overview.

BIBLIOGRAPHY

Curtice, J., 2015. A Return to Normality? How the Electoral System Operated.

Parliamentary Affairs 68(suppl 1): 25–40.

Baxter, M., 2007. *Strong Transition Model*.

<http://www.electoralcalculus.co.uk/strongmodel.html> (January 14, 2015).

Dale, I. (Ed.), 2014. *The Politicos Guide to the 2015 General Election*. Biteback Publishing, London.

Gibson, R., Lewis-Beck, M.S., 2011. Methodologies of Election Forecasting: Calling the UK “Hung Parliament”. *Electoral Studies* 30 (2), 247-249.

Hanretty, C., Lauderdale, B.E., Vivyan, N., 2015. Evaluating our 2015 General Election Forecast. <http://electionforecast.co.uk/resources/election-forecast-evaluation-report.pdf> (Accessed July 3, 2015).

Hanretty, C., Jennings, W., 2015. Expert Predictions of the 2015 General Election. <http://www.psa.ac.uk/sites/default/files/PSA%20GE%20Election%20Predictions%20Report.pdf> (Accessed March 3, 2015).

Lewis-Beck, M.S., Stegmaier, M., 2014. United States Presidential Election Forecasting: An Introduction. *PS: Political Science and Politics* 47(2), 284-288.

Lewis-Beck, M.S., Tien, C., 2015. Forecasting Lessons from 2014: Long Lead, Strong Theory, No Polls. *PS: Political Science and Politics*, 297-299.

Murr, A.E., 2015. The wisdom of crowds: Applying Condorcet’s jury theorem to forecasting US presidential elections. *International Journal of Forecasting* 31(3): 916–929.

Tapper, A., 2015. Predicting GE 2015: which group of forecasters should we rely upon: 'the expert crowd', 'the betting crowd' or the 'mass public crowd'?. http://alberttapper.blogspot.co.uk/2015/05/predicting-ge-2015-which-group-should_7.html (Accessed July 3, 2015).

Singh, M., 2015. Is there a shy Tory factor in 2015?. <http://www.ncpolitics.uk/2015/05/shy-tory-factor-2015.html/> (Accessed July 3, 2015)

Whiteley, P., Clarke, H., 2015. In defence of pollsters: they never said they could predict the UK election. In: *The Conversation*. May 15. <https://theconversation.com/in->

defence-of-pollsters-they-never-said-they-could-predict-the-uk-election-41693 (Accessed June 4, 2015).

TABLE 1. FORECASTS OF THE 2015 BRITISH GENERAL ELECTION +

Team ^a	Vote Share (%) ^b	Seat Share (#) ^b	Rank ^c	Lead ^d
	(C, L, D, S, P, U, G, C-L)	(C, L, D, S, P, U, G)	(Seats Accuracy)	(months)
<i>Actual</i>	(38, 31, 8, 5, 1, 13, 4, +6)	(331, 232, 8, 56, 3, 1, 1)		
<i>Final polls ^e</i>	(34, 34, 9, 5, 1, 13, 5, +0)	(270, 285, 19, 55, 2, 0, 1)		
BGSSW	(30, 29, 5, 9, 0, 24, 2, +1)	(285, 306, 21, 9, 3, 5, 1)	(5, 12, 2, 8, 1, 6, 1)	2
F	(35, 32, 10, 4, -, 12, -, +3)	(285, 262, 25, 53, 3, 3, 1)	(5, 1, 6, 1, 1, 5, 1)	.03
FJPW	(35, 33, 9, -, -, -, +2)	(274, 272, 24, 53, -, 2, -)	(9, 5, 5, 1, -, 2, -)	.23
HLV	(34, 33, 12, 4, 1, 11, 4, +2)	(278, 267, 27, 53, 4, 1, 1)	(7, 4, 8, 1, 4, 1, 1)	.03
LN	(-, -, -, -, -, -, +3)	(287, 263, -, 41, -, -, -)	(3, 3, -, 6, -, -, -)	2
LNB	(34, 32, 8, -, -, -, +3)	(286, 274, 10, -, -, -, -)	(4, 6, 1, -, -, -, -)	4
MF	(33, 31, 10, -, -, 14, -, +2) ^f	(274, 278, 23, 52, 2, 0, -)	(9, 7, 4, 4, 4, 2, -)	.25
M	(-, -, -, -, -, -, -)	(292, 262, 28, 40, 4, 5, 2)	(2, 1, 9, 7, 4, 6, 5)	2
P	(35, 31, 16, -, -, 13, -, +4)	(296, 287, 38, -, -, -, -)	(1, 10, 11, -, -, -, -)	12
RTB	(33, 33, 10, -, -, 13, -, +0)	(276, 280, 22, 48, 3, 2, 1)	(8, 8, 3, 5, 1, 2, 1)	3
SW	(31, 35, 8, -, -, -, -4)	(245, 299, 26, -, -, -, -)	(12, 11, 7, -, -, -, -)	3
WCSS	(-, -, -, -, -, -, -)	(271, 281, 34, -, -, -, -)	(11, 9, 10, -, -, -, -)	6

Average (34, 32, 10, 6, 0.4, 15, 3, +2) (279, 278, 25, 43, 3, 3, 1)

+ If an entry is blank (-), that signifies no forecast was reported in that instance. All figures are for Great Britain rather than the UK, i.e. excluding Northern Ireland.

^a = initials (first letter of each last name) for each forecasting team: Burnap, Gibson, Sloan, Southern and Williams (BGSSW); Fisher (F); , Ford, Jennings, Pickup and Wlezien (FJPW); Hanretty, Lauderdale and Vivyan (HLV); Lebo and Norpoth (LN); Lewis-Beck, Nadeau, and Bélanger (LNB); Mellon and Fieldhouse (MF); Murr (M); Prosser (P); Rallings, Thrasher and Borisjuk (RTB) ; Stegmaier and Williams (SW); Whiteley, Clarke, Sanders and Stewart (WCSS); ^b = the national share (votes or seats) by party (C = Conservative, L = Labour, D = Liberal Democrat, S = Scottish National Party, P = Plaid Cymru, U = UK Independence Party, G = Green) ; ^c = the accuracy ranking of the seats forecast by the team, from most accurate (rank = 1) to least accurate (rank = 12); ^d = lead time before the election in months (or fraction of a month); ^e = average of the final vote-intention polls that were published before the election, and a uniform seats projection which also uses the average of Scottish vote-intention polls for change in Scotland. ^f = these figures were not strictly forecasts but what Mellon and Fieldhouse weighted their shares of the vote to and so presented for information.