

How the Exit Poll got it more or less right

Stephen Fisher Jouni Kuha

University of Oxford, and the LSE

National Centre for Research Methods & the British Polling Council
'The Polls in 2010: Learning the Lessons'
22nd November 2010

Exit poll prediction: 10pm on May 6th

	UK Seats			
	CON	LAB	LD	OTH
Prediction	307	255	59	29
Final result	307	258	57	28

How did we get this close?

Part I: General description of the methods

2010 Exit poll and analysis team

- Commissioned by BBC, ITN and Sky
- Fieldwork by GfKNOP and Ipsos MORI.
- Statistical analysis of exit poll data by John Curtice (Strathclyde), Stephen Fisher (Oxford) and Jouni Kuha (LSE)
- Also in the psephology team: Rob Ford, Will Jennings (Manchester), Clive Payne (Oxford)
- Statistical analysis and prediction of elections since 1970: Phil Brown, Clive Payne, David Firth
- The methods of analysis in 2010 were almost entirely, and computer code very largely, the same as in 2005, and due to David Firth
 - see Curtice and Firth (2008), *Journal of the Royal Statistical Society A*, 171, pp. 509-539

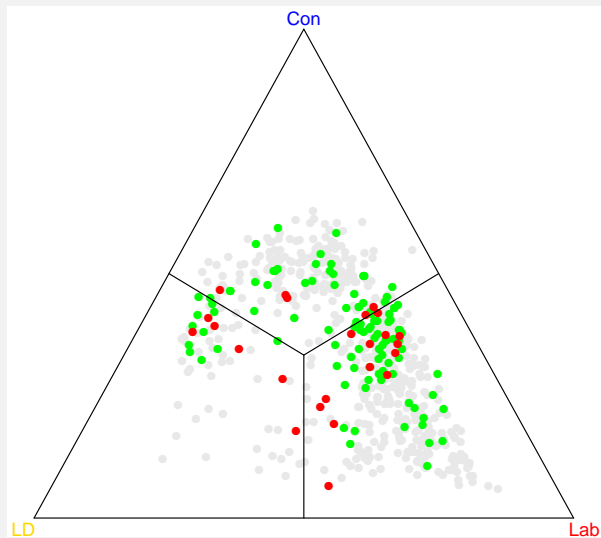
Basic Principles

- Variance between constituencies greater than variance in change between constituencies.
- There may be bias in exit polling that is fairly stable over time.
- So revisit same exit poll locations as last time for efficiency and hope that any bias cancels.
- Also, avoid uniform change assumptions by using a statistical model to estimate the pattern of change.
- Finally, allow for random fluctuations and so produce predicted *probabilities* of each party winning, not (deterministic) assignments of each seat to a single predicted winner.

Data structure

- 130 polling stations (up from 120 in 2005)
 - 104 were also used in 2005
 - a top-up of 26 new ones, with 2005 represented by *constituency* results
 - Choice of top-up locations motivated by:
 - rebalancing previous sample to restore strategic and regional profile after attrition from boundary changes and other feasibility problems with old locations
 - increasing coverage of Lab-LD contests (newly prevalent post 2005)
- At each location, a systematic sample of voters throughout the day
- 18,688 interviews: an average of 144 per location, so large variance in estimates of change at any particular location.

Exit poll locations relative to 2005 results



Main-party shares in the constituency in 2005 (red=new location in 2010)

Summary of the modelling procedure

- 1 Model for 2005-2010 changes in response variables Y_{party} , given constituency-level explanatory variables
- 2 Predicted values $\hat{Y}_{party,seat}^{2010}$
- 3 Predicted shares $\hat{P}_{party,seat}^{2010}$
- 4 Predicted probabilities of winning $\hat{\pi}_{party,seat}^{2010}$
- 5 Predicted number of seats for a party is

$$\hat{Seats}_{party}^{2010} = \sum_{seat} \hat{\pi}_{party,seat}^{2010}$$

Other methodological considerations

- **Postal Voting:** No adjustment since limited change in its usage since 2005 and expensive and difficult to estimate the effects of change.
- **Refusals:** Interviewers guessed the vote choice of all those approached so that the relationship between the guesses and votes of those who did respond could be used to impute votes for the 18% who did not respond.
- **Non-contacts:** There was a weighting method to adjust for lower contact rates at busy times of day.
- Neither imputation nor weighting made much difference so the published predictions were based on raw unweighted data.

Part II: Evaluation of the Exit Poll

Main questions for the evaluation

- How good was the data?
 - How good were the exit poll estimates of change both overall and in comparison with the actual results in the sampled seats?
 - How good were the estimates for the new locations and should they have been used?
- How good was the modelling?
 - Did the models identify patterns of change that really existed, either overall or in the sampled seats?
 - Were there any variables we missed?
- Was it better to use probabilistic rather than deterministic prediction?
- How did each of the sources contribute to the prediction?

Changes in the share of the vote

	CON	LAB	LD
Exit poll	+5.3	-7.5	+1.0
Overall Result	+3.8	-6.5	+1.0
Seats in Sample	+4.3	-7.1	+1.0

Spot on for the Lib Dems, but too large a Tory rise and Labour fall. This is partly due to the nature of the selected locations.

Changes in the share of the vote: New and Old Locations

	CON	LAB	LD
Old Locations:			
Exit Poll	+5.6	-8.5	+1.3
Seats in sample	+4.2	-7.5	+1.3
New Locations:			
Exit Poll	+4.1	-3.0	-0.2
Seats in sample	+4.8	-5.6	-0.2

- Both old and new were spot on for the Lib Dems.
- New seats had too small a Conservative rise and Labour fall, opposite for old cases: differences not significant.
- Nonetheless, compensating biases helped improve the overall figures.
- Existence of a difference suggests we shouldn't have used new locations, but this was not observed on the day.
 - within the exit poll the difference between new and old was not statistically significant for Con and LD.

Exit poll models: LD/(LD+LAB+CON)

Predictor	Exit poll		Actual (All Seats)	
	Coeff.	<i>p</i> -value	Coeff.	<i>p</i> -value
England	1.8	0.004	1.9	< 0.001
Scotland/Wales	-5.0	0.030	-2.3	< 0.001

- Exit poll model very close to reality.
- Hard to improve on the model given we restricted ourselves to binary predictor variables with no sharper than a 80:20 split.

Exit poll models: $LAB/(LAB+CON)$

Predictor	Exit poll		Actual (All seats)	
	Coeff.	<i>p</i> -value	Coeff.	<i>p</i> -value
England	-13.0	< 0.001	-10.4	< 0.001
Scotland/Wales	-2.1	0.52	-5.1	< 0.001
<90% white population	5.2	0.012	3.5	< 0.001
Labour incumbent standing	3.9	0.024	2.3	< 0.001

Regression coefficients remarkably similar, and same again for selected 130 seats.

What predictors did we miss?

- Most notably,
 - Labour did worse where unemployment went up.
 - Labour did better where there were more public sector workers.
- But these patterns were not apparent in the actual results for the seats in the sample.
- Strong preference for parsimonious models, so we may not have wanted both these and previous variables all in the model.

Calibration of the predicted probabilities

Predicted probability (%)	Wins/possible	% Won
(0,10]	13/109	12
(10,20]	7/39	18
(20,30]	6/27	22
(30,40]	10/25	40
(40,50]	19/46	41
(50,60]	23/38	61
(60,70]	14/21	67
(70,80]	18/24	75
(80,90]	26/32	81
(90,100)	88/100	88

Given the predicted shares and the actual results, it appears that the predicted probabilities could not have been improved on easily.

Results are sensitive to the modelling choices

Real and predicted UK seats:

	CON	LAB	LD	OTH
Real result	307	258	57	28
Uniform change from final polls	282	256	81	31
Deterministic & uniform change	321	234	65	30
Probabilistic & uniform change	319	236	65	30
Deterministic & regression models	303	254	63	30
Probabilistic & regression models (<i>as published at 10pm</i>)	307	255	59	29

Conclusion

- Successful prediction was aided by all parts of the methodology: data, modelling and probabilistic prediction.
- Key shock of Lib Dem failure by comparison with the polls, was clear from the very good data: Thanks to Ipsos MORI and GfKNOP!
 - But we were lucky to have compensating biases from new and old locations for overall Labour and Conservative changes.
 - The models also helped redress the error.
- The regression models did identify important sources of variation between polling locations that was present in the actual constituency results.
 - But other factors, such as unemployment and employment sector were not identifiable because of the choice of locations.
- As expected, probabilistic prediction made only a small improvement in the prediction on this occasion.
- While not perfect, the methodology seems to be broadly working well.