

Election Forecasting in the UK

Clive Payne, Nuffield College, University of Oxford, UK

Email: Clive.Payne@nuf.ox.ac.uk

Paper presented at the Institute of European and American Studies, Academia Sinica,
Taipei, Taiwan

November 28, 2001

Abstract

The paper reviews the various statistical methods used by the BBC to forecast different types of election in the UK in the last thirty years. Elections in this period were of three main types:

- The Plurality system (Westminster Parliament, European Parliament until 1999, and annual elections for local government councils)
- National Referendums (Entry to the EEC, Welsh and Scottish devolution, London Assembly, Good Friday Agreement in Ulster)
- Variants of the Additional Member system used for the Welsh, Scottish and London assemblies post 1997 and the European Parliament in 1999

The forecasting methods used in three different contexts are described:

- Predictions of the result before the election takes place using opinion-poll series and other data
- Predictions of the result as soon as the polling stations close using exit polls and other techniques
- Predictions of the final outcome on election night itself using the subset of actual results declared

The paper describes the special features of the UK electoral systems which determine the statistical methods used and includes comparisons with election forecasting in other countries. The BBC's performance in election-night forecasting over the period is assessed. Media aspects of the presentation of election forecasts are also highlighted.

1. Introduction

A regular feature of the coverage of elections in the UK is forecasting the final result. This activity involves pollsters, parties and pundits and punters. Before the election the media gives extensive coverage of poll-based predictions and pundits' forecasts of the outcome. A climax is reached on the night of the election itself when the main TV and radio channels give a forecast as soon as the polling stations close and produce updated forecasts throughout the night as results are declared.

The main aim of forecasting is to predict the number of seats won by each party in the representative body and, as a subsidiary exercise, the parties' overall share of the votes cast and the turnout. Three forecasting contexts can be distinguished:

1. predicting the final result before the election takes place (the 'pre-forecast'),
2. immediately after the polling stations close (the 'prior forecast')
3. during election night itself using the subset of actual results declared (the 'results-based forecast').

Each context presents varying and challenging problems in the UK.

The UK has traditionally used the plurality system in single-member constituencies for its elections and this system presents special problems for the forecaster and is the main emphasis of the paper. The paper describes the statistical methods used by the BBC, and modifications to cope with the changing political context (particularly the increasing influence of the Liberal Democrat party and the Nationalist parties in Scotland and Wales which have challenged the Conservative-Labour hegemony in recent general elections. But since the election of a Labour government in 1997, new semi-proportional electoral systems have been used to elect representatives to new bodies in Wales, Scotland and London. The paper briefly describes the modifications to the basic methodology introduced to deal with these new types of election.

The UK has also made occasional use of referendums as a way of deciding on a major issue of national policy. The main aim of forecasting here is to predict the overall percentage of the voters voting for and against the proposition(s). Election-night forecasts were not produced for the referendum on the Good Friday agreement in Northern Ireland in May 1998 as was thought that publishing forecasts and detailed analyses if the results might offend political sensitivities. Currently referendums are also being undertaken in a number of towns and cities to determine whether the local electors are in favour of having a directly elected mayor to run the city council. These have been fairly low profile elections with little media interest in forecasting the result. A key national referendum on the Euro currency is expected to be the next major vote in the UK.

2 Types of Election

Plurality	Referendum	Additional Member	Others
Parliament (Westminster) and Byelections	Entry to the EEC (Common Market) 1975	Scottish Parliament (1999)	Regional List: European Parliament (1999)
UK MEPs for European Parliament (1979-1995) *	Devolution in Scotland (1979 and 1997)	Welsh Assembly (1999)	Supplementary vote: London Mayor (2000)
Local Councils *	Devolution in Wales (1979 and 1997)	Greater London Authority (2000)	
	Greater London Authority (1998)		
	Ulster: Good Friday agreement(1998)		
	City Mayors (2000-)		

* Ulster has used the Single Transferable Vote (STV) to elect its representatives in the European Parliament and for the Ulster Assembly

2.1 Plurality

Westminster Parliament: Elections to the national parliament called ‘general elections’ and are the most important elections in the UK. Currently 659 Members of Parliament (MPs) are elected by the plurality system (‘first-past-the-post’) in single-member constituencies (or ‘seats’) with electorates averaging about 60000 electors. Vacancies are filled by byelections, which sometimes have a high media profile, particularly when the incumbent governing party is unpopular. Constituency boundaries are revised about every 10 years to (approximately) equalise electorates. The maximum term of the Westminster Parliament is 5 years but elections are called at the governing party’s choosing. Each elector is assigned to a polling district and votes in a polling station. The number of polling stations in a seat ranges from 30-100, depending on the size of geographical area. A small proportion of electors (<3%) cast their vote by post or proxy.

European Parliament (until 1994): 84 MEPs representing Great Britain elected in single-member constituencies, each of which is an aggregate of between 5 and 8 Westminster constituencies, with large electorates averaging about 0.5m. The election is held every five years. A new voting system was used in 1999 (see below)

Local Government Councils: About 13000 councillors are elected by simple plurality in single (mainly) or multi-member wards for 363 councils. Local Council Elections for County, District and Borough Councils are held on a fixed day in May every year.

The cycle and coverage of the local elections varies considerably with either whole Council elections every four years or one-third Council elections each year. The main interest in local elections has been in estimating the national standing of the parties, particularly near the time of general elections, and then forecasting the notional outcome in seats in the Westminster Parliament. Problems are presented by the complexity of local elections, especially their varying coverage, the uneven patterns of candidacy of the major parties, the strength of independents, the existence of multi-member seats in some Councils, and boundary changes. Curtice, Payne and Waller (1991) describe the problems of using local election results to estimate the national standing of the parties. There is now increasing interest in forecasting the results of local elections, especially of which parties will have control in individual councils which are at risk of a change of control.

2.2 Additional Member system

This system was introduced by the incoming Labour Government in 1997 for elections to new devolved bodies in Scotland, Wales and London, as part of an agreement with the Liberal Democrat party made before the 1997 General Election to give greater proportionality of representation in these new bodies.

This system has two components: some MPs are ‘directly’ elected by plurality in single-member constituencies which are grouped into regions; the remaining additional or ‘top up’ MPs are selected from a closed party list for each region by applying a procedure described below. The proportion of directly elected MPs has varied according to the election: Scottish Parliament (57%), Welsh Assembly (67%) and London Assembly (56%).

Scottish Parliament: 129 Members of the Scottish Parliament (MSPs) with 73 elected by simple plurality in single-member constituencies (corresponding to the 72 existing Scottish Westminster constituencies with the Orkney and Shetland constituency split into two) and the remaining 56 MSPs elected by the Additional Members system (7 in each of 8 regions).

Welsh Assembly: 60 Members of the Welsh Assembly (MWAs) with 40 elected by simple plurality in single-member constituencies (corresponding exactly to the existing Westminster constituencies in Wales) and the remaining 20 MWAs elected by the Additional Members system (4 in each of 5 regions).

Greater London Authority: 25 Members of the Greater London Assembly (MGLAs) with 14 elected by simple plurality in single-member constituencies (corresponding to the existing London Boroughs) and the remaining 11 MGLAs elected by the Additional Members system (11 in one region – London as a whole).

Voters cast two votes in these elections: a *first* vote for the candidate of their choice in their constituency and a *second* vote for the party list for their region. The top-up procedure in each region is carried out by applying the d’Hondt rule as follows:

Let $V(i)$ be the total (second) votes cast for party i in the region,
 $S(i)$ be the number of constituencies won by party i in the direct (first vote)

election,
A(i) be the number of additional seats gained in the top-up procedure, initially set at zero.

Then the first additional topup seat is allocated to that party which has the maximum value of $V(i)/(S(i)+A(i)+1)$.

Increment A(i) by 1 for the successful party and repeat the previous step until all the additional seats have been allocated. The final seats for each party within the region is then $R(i)= S(i)+A(i)$

Sum R(i) over all regions to give the overall distribution of seats between the parties in the Assembly

2.3 Other Systems

Regional List

This system was introduced for the elections to the 1999 European Parliament replacing the plurality system used for these elections previously. 84 Members of the European Parliament (MEPs) were elected by a Regional List system (71 from 9 regions in England, 8 from Scotland and 5 from Wales). The regions are aggregates of Westminster constituencies (17-70?) electing between 3 and 11 MPs in proportion to their electorates (1.2m to 5.2m). In this system the voter chooses one party list (closed) for their region. Within each region seats are allocated using the d'Hondt rule as with the Additional Member System, except that as there are no directly elected MPs the initial divisor for each party is set at 1 (i.e $S(i)=0$). Thus this system can be regarded as a special case of AMS (0 directly elected MPs and 84 from the top-up).

Supplementary Vote

This system was introduced for the election of the London Mayor in May 2000. There were eleven candidates for this office. Voters marked the ballot paper with their first and second preferences. All but the top two candidates on the first preference count are then eliminated and the ballot papers for the discarded candidates are redistributed to the two remaining candidates according to the second preferences until one of the top two achieved a majority. A key feature here is that the overall 'turnover' matrix of first by second preferences is required to forecast the outcome if the election is not decided on the basis of first preferences only.

3. Problems of UK forecasting (and comparison with other countries)

There are some aspects of elections in Great Britain which make election forecasting more difficult than in most other countries (see Morton (1990) for a description of election-night forecasting in various countries). The following factors, in decreasing order of importance, determine the level of difficulty encountered.

(i) Whether election results are available at a disaggregated level particularly for polling stations (districts). If they are then the forecasting problem is much easier as predictions can be based on a suitable sample of the locations. The sample of locations can be selected either to achieve representativeness (eg Spain) or to give indicator swingometric or barometric districts (eg USA). The availability of results at polling district level also makes the problem of selecting a sample of voters to interview about their voting intentions much easier. The greater volume of district results also allows the rapid calculation of adjustments for bias in the early declared results (eg in Australia). The UK is one of the few countries where election results are not available at polling district level.

(ii) The system for converting votes into seats. Here there is a basic distinction between proportional systems and other systems. The prediction problem is much easier in proportional systems, which are used in most other European countries. Forecasting here is essentially an exercise in sampling voters (usually after they have voted) in representative districts to estimate the shares of vote for each party. The conversion to seats is then straightforward using simple algorithms such as the d'Hondt rule, usually subject to a threshold which a party has to pass to qualify for seats. For other electoral systems the conversion to seats is more difficult, particularly in those where voters rank candidates in order, such as the Single Transferable Vote used in Ireland (see Smith (1982)) although this is a proportional system and the Alternative Vote used in Australia, where preference distributions among voters must be built into the conversion procedure. Some other non-proportional systems present problems of conversion that are much more difficult to handle than in the UK. Examples are the US presidential elections where the seats are allocated on a winner-take-all basis at the State level (as evidenced by election-night forecasting problems experienced the 2001 Bush/Gore contest) and Single Non Transferable Vote in multimember constituencies system (eg in Taiwan) where this problem is particularly acute. Thus in this respect the UK's plurality system represents an intermediate level of difficulty.

(iii) How many significant parties there are and whether the parties' candidates and their patterns of support are unevenly distributed. The UK presents some difficulties in this respect as there are three major parties standing throughout GB, so that the forecasting is at least a three-dimensional problem which also has to take account of the uneven patterns of support of the Labour and Conservative parties in particular. There are two further significant parties, the Scottish and Welsh Nationalists, which have an important influence on the predictions in their respective countries. The implication is that data must be obtained from a range of types of constituency before any real confidence can be attached to the forecasts.

(iv) The quality of the socio-political-economic data available for the constituencies. In the UK constituencies there is a paucity of up-to-date data to use either as covariates in election-night prediction, or on which to base sample selection for voting intention polls. The Westminster constituencies do not correspond to local government areas where data is regularly collected for administrative purposes. The decennial Census conducted in the first year of each decade does provide data for constituencies, but such data is not available until at least a year after the Census and will be out of date for elections which take place after the beginning of the relevant decade.

(v) Whether there is significant bias in the declaration order of results on election night. In the UK counting of votes is usually done at one place in each constituency so that results from urban areas are generally more likely to be announced sooner than those in rural areas thus introducing the possibility of significant bias in the order of results declared. In recent general elections the early declarers have been disproportionately safe Labour seats which have exhibited rather different electoral behaviour from the Conservative-Labour and Conservative-Liberal Democrat marginal contests where the election outcome is decided.

(vi) Who collects the results and how they are declared. In some countries results are made available instantly in computer-readable form by an official authority (eg the Electoral Commissions in Australia and India). In the UK, the media organisations each has to employ so called 'stringers' at each count to supply the results to them; this is an error prone procedure. Another very important factor is whether results are made available on a progressive basis as they are for example in Australia and India (see Karandikar, Payne and Yadav (2001) on general election forecasting in India) rather than only when the whole constituency result is announced as in the UK.

4. Pre-election forecasting

4.1 Predicting from Polls

In the interval between national elections, voting intention polls are regularly carried out and the aim is to translate voting intentions into a pre-election forecast of the number of seats that would be won by each party if a national election were held on the date of the poll. Two pollsters, ICM (for the Guardian newspaper) and Gallup (for the Daily Telegraph newspaper) conduct a monthly poll of party support in the UK; the Gallup series started in the 1950s. During the general election campaign itself polls are published on an almost daily basis and there is special interest in the final poll-based predictions made by each polling organisation. The four main pollsters in the UK are currently Gallup, NOP, ICM and MORI.

The annual local elections in May also provide an opportunity to estimate the national standing of the parties; here local election results are aggregated to a sample of notional Westminster constituency results. There is special interest in this forecast in the period leading up to a national election, especially a general election, when the governing party is keen to assess the number of seats it is likely to win given its current standing, and so help decide when to call the election.

The data available for pre-election forecasts using the polls are very sparse. Usually they are no more than an estimate of the projected national shares of the vote for each party, based on a quota or telephone poll of no more than 1000 electors. Occasionally larger polls are carried out and disaggregated vote share estimates, usually on a regional basis, may be available. Separate polls are often carried out in Scotland. The local elections provide a rich database containing separate vote share estimates by region and type of political contest.

Thus a model is required to convert national voting shares into seats in Parliament in a plurality system. There is not a great premium on accuracy as a real election outcome is not being predicted, except in the final campaign polls, and the polling data are of course themselves subject to sampling and non-sampling errors. There are various alternative models which can be used for this purpose and their efficacy can be tested empirically by retrospective application to previous general election results.

In the 1970s, pre-election forecasts in the UK were based on 'swing', a measure of the change in the relative shares of the vote for the two major parties, Conservative and Labour. If C1 and L1 were the percentages of the vote for Conservative and Labour respectively at the previous election and C2 and L2 are the percentages forecast by the poll, the swing, S, is defined as

$$S = 0.5 * [(C2 - C1) + (L1 - L2)]$$

Then, given the assumption that all other electoral movements involving other parties and abstention cancel out, swing measures the net transfer of votes between Conservative and Labour. The forecast is obtained by applying the predicted swing S to each constituency in turn. The forecast Conservative share is C1+S, Labour's is L1-S, and the percentage shares of the all other parties are assumed to remain unchanged so that the predicted winner is the party with the highest predicted percentage. A standard presentational technique is to rank seats in order of the swing required to change hands and display this as the election 'battleground'. A 'swingometer' calibrated with the predicted seats for various levels of swing is a regular feature of the TV coverage of general elections, and some presenters take pleasure pitting the swingometer against the more sophisticated computer-based forecasts.

The Conservative-Labour swing model worked reasonably well up to the 1980s as the minor parties then won very few seats. It gave reasonable predictions of seat totals but often the outcomes in individual seats were badly predicted with sometimes as many as fifty wrongly predicted individual winners. There was no attempt to attach a probability estimate to the predicted outcome in an individual seat, or to calculate confidence intervals for seat totals.

With the rise of the third parties in the 1980s (the Liberal Democrats and its predecessor the SDP/Liberal Alliance) and subsequently, modifications were needed to cope with the threat to the hegemony of the major parties (in terms of seats). A natural extension of the basic swing model is to apply the predicted change in percentage share of the vote for each significant party uniformly across all constituencies to give a projected winner in each seat - the 'uniform change' model. This model has been used by the BBC from 1974 onwards. It has provided good prior forecasts when used on the election-night programmes and retrospective testing on actual election outcomes has confirmed that this refined model has usually been sufficiently accurate for pre-election use. But in the last two general elections this crude model had tended to overestimate the number of seats won by the Conservative party. The main explanation for this underperformance is the occurrence of anti-Conservative tactical voting in the key marginal seats.

4.2 Pollster's predictions

The final campaign polls have generally had a good track record except for the general elections in 1970 and 1992 which were both close elections and for which the pollsters' final pre-election forecasts exhibited embarrassing errors. In 1970 almost all of the final polls predicted a large Labour majority when in fact Labour won a very small majority- late swing was suggested as the main explanation for the errors. There were extensive post mortems carried out in 1992 where all the final polls predicted a Labour overall majority when in fact the Conservatives won with small overall majority (Market Research Society, 1994). The main explanations suggested for the errors were: late swing, problems with the allocation of quotas, differential refusal/lying and effects of the projected 'poll tax' where many eligible voters deliberately disenfranchised themselves. Most of the polling organisations revised their methodology after the 1992 general election with all but one adopting telephone-based polls rather than the face-to-face quota polls previously used. In recent general elections there has been persistent tendency to overestimate support for Labour and underestimate support for third parties in the final pre-election polls. Table 3 shows the final polls in the 2001 general election – all overestimate Labour support although ICM got very close (as it did in the preceding general election).

Table 3 Final campaign polls in the 2001 general election on June 9

Company	Final Result	NOP (Sunday Times)	ICM (Guardian)	MORI (Times)	GALLUP (Daily Telegraph)	BBC Weekly Poll-of-polls
Date published		June 3	June 6	June 7	June 7	June 7
Sample Size		1105	1009	1967	2399	
Field dates		May 31 - June 3	June 2-4	June 5-6		
Method		Telephone	Telephone	Face-to-face quota	Telephone With call backs	Average of June 3-7 polls
Con %	33	30	32	30	30	31
Lab %	42	47	43	45	47	46
Lib Dem %	19	16	19	18	18	18
Lab lead Over Con	9	17	11	17	17	17

Since 1992 the Royal Statistical has routinely published reviews of the performance of the polls such as Smith (1992), Worcester (1996) and O'Muirheartaigh and Lynn (1997).

4.3 Poll time series-based forecasts

An interesting approach to pre-election forecasting has been to use econometric models fitted to the time-series of poll data to predict the final result. The main exponent of this approach has been Sanders who has been producing generally good forecasts of the parties' final shares of the vote in all general elections since 1987. The following description is extracted from the YOUGOV.COM website where Sanders and his associate Kellner published their 2001 forecasts.

The predictions are based on a simple econometric model of the monthly Gallup poll series for support for the governing party. This is adjusted to compensate for bias towards Labour using a method devised by Kellner and called the Kellner Index of party support. The logic of this approach is simple. First, a statistical model is developed that provides a reasonably accurate description of movements in party support over the last few years. Second, using different assumptions about the predictor variables in this statistical model, the model is then employed to forecast future movements in support.

The core predictor variable in the model is based on the idea of "economic voting". Extensive research in several countries has shown that the electoral fortunes of governing and opposition parties are strongly linked to the performance of the economy. When the economy is going well, governments tend to get re-elected. When it is going badly, they tend to lose. There is no simple correlation between economic performance, as measured by unemployment, growth or inflation, and governing party support. However, in Britain, over the last 25 years, there has been a fairly consistent relationship between people's "economic expectations" and their tendency to support the incumbent or governing party. When expectations rise, voters tend to increase their support for the government: they want to preserve the political status quo that has created their optimism. When expectations fall, voters tend to increase their support for the main opposition party: they are more inclined to seek to change the status quo that has created their pessimism.

The first predictive component of the model, therefore, involves an aggregate measure of voters' economic expectations. This is measured using the responses taken from Gallup's regular monthly survey of economic confidence, published in the Daily Telegraph. Gallup asks a representative sample of voters the following question:

How do you think the financial situation of your household will change over the next 12 months? Will it:

- Get a lot better
- Get a little better
- Stay the same
- Get a little worse
- Get a lot worse
- Don't know/refused

The "personal expectations" index used in the model here is measured as the percentage of respondents who think things will get better minus the percentage who think things will get worse.

The second component of the model is a measure of how long the government has been in office. Research by two Danish economists, Martin Paldam and Peter Nannestad, has shown that governments in all democratic countries are subject to a "cost of ruling". Generally speaking, the longer they remain in office, the more their support tends to seep away. On average, net of all other effects, democratic governments tend to lose around 3 percentage points during the lifetime of a typical parliament. There is some debate as to the precise mechanism involved in this process. But it is generally believed to be the consequence of the fact that governments, unlike oppositions, can be held responsible for almost anything that goes wrong in the country while they are in office. This generalised attribution of blame gradually eats into the government's support - unless it can counter the phenomenon by conspicuously performing well in other ways, such as creating a huge economic boom or winning a war abroad.

The third component of the model involves terms that measure the effects of "unusual events". It is well known among observers of public opinion that opinion can change quite dramatically as a result of an unusual event. The model used here makes use of two terms that seek to measure "unusual events". The first is the dramatic surge in Labour support that took place immediately after the May 1997 general election. The model accordingly includes a term for June 1997. This is a dummy variable which takes the value zero for June 1997 and zero otherwise. On top of this "June 1997 effect", New Labour also benefited from a more prolonged "honeymoon" effect. The model accordingly includes a term that assumes this effect built up for the first six months of the New Labour government. Both of these effects are then assumed to "decay" fairly rapidly thereafter.

The model includes no measures of the "objective economy", such as unemployment, inflation or interest rates. This is because statistical tests showed that none of these variables exerts a direct effect on levels of party support. The implication of this is that people's perceptions of their own economic circumstances are more important in determining their political preferences than the performance of the economy as a whole.

The best-fitting model was an OLS (Ordinary Least Squares) regression model for the period since the 1997 general election used for projecting the 2001 general election was:

$$\text{Lab}(t) = 35.8 + .36 \text{Lab}(t-1) + .11 \text{Pexp}(t) - .16 \text{Time} + 2.1 \text{Honeymoon} + 4.3 \text{June 97}$$

where Lab(t) is Labour support in any given month as measured by the Kellner index; Pexp(t) is aggregate personal economic expectations as described above; Time is a time trend; and Honeymoon and June 97 are dummy variables. The coefficients were all statistically significant and the model passed a battery of standard diagnostic tests.

The forecasting procedure needs an estimate of the value of Pexp at the expected time of the election. Kellner and Sanders used a "stable expectations" scenario which assumed that personal expectations (the index, derived from Gallup, of the difference

between the percentage of economic optimists and percentage of pessimists) remained at the then current level of +12 through to the election.

Conservative support was calculated on the assumption that the Liberal Democrats obtain 17% and Others 5% of the popular vote.

Their final projections for 2001 were Labour: 45, Conservative: 33, Liberal Democrat: 17 which with their confidence interval of +/- 2% almost includes the final results of 42/33/19.

4.4 Other methods

In recent general elections the media have given coverage to predictions based on newer and sometimes more informal forecasting methods. These include:

1. The Reuter's panel – an average of about 50 individual psephologists' predictions of the outcome.
2. Spread betting - Bookmakers offer an even bet on a range of seats for the outcome, eg that the Labour majority would be between 250 and 260 seats. As punters place their bets this range is shifted to reflect their betting behaviour. If more money is bet on a lower figure the range is shifted downwards; if more is bet on a higher figure it is shifted upwards. In the 2001 general election the final pre-election ranges of seats for each of the major parties arrived at by this process contained the final result!
3. Using local byelection results - Rallings and Thrasher (1999) base their predictions of the outcome of the 2001 general election on the sequence of byelections for vacancies in local Councils which take place throughout the year.

5. Election night: prior forecast

5.1 Approaches

An essential part of election night media coverage is the prior forecast which is broadcast by the main TV channels as soon as the polls close at 10pm. These forecasts have a very high profile and the forecasting performance of the TV companies tends to be judged by the public primarily on what they predicted at 10.01pm before any actual results are available. A key question is what sort of information to use as the basis for the forecast. Four sources have been used, sometimes in combination, but each has cost and logistical implications.

a) Poll-of-polls

This is a weighted average of the final pre-election polls conducted by the major polling organisations giving predicted national shares. This information is essentially free, but the forecasts of shares are sensitive to a late swing (ie changes in voting

intention between the last poll and the actual vote) and there are problems of identifying respondents who will actually vote, and converting votes into seats.

b) Quota polls

A specially commissioned quota poll of large size (typically 2000-5000) voters is carried out on the day before the election and on the day itself. The poll is usually carried out to elicit voter attitudes as well as voting behaviour, but has booster samples in marginal seats. The poll gives projections of shares of the vote, usually disaggregated by region and type of political contest (categorised by the previous top two parties and the lead of winner over second party). This source is subject to the usual problems with quota polls, especially that of distinguishing people who are actually going to vote.

c) Exit polls

A national sample of polling stations is chosen, often confined to marginal seats, and a sample of voters leaving the polling stations is asked how they voted. Samples of 10000-15000 voters are typical. Estimated changes in parties' shares are applied to seats of similar type to project the winning party. However exit polls are very expensive and there is a problem over the choice of the sample of polling stations as there is very little information on previous voting (unless a portfolio of stations is established over a long period which requires a lot of effort initially and needs regular updating). The BBC has usually chosen one polling station at random within each selected constituency; ITN have attempted to select a polling station which is politically representative of the constituency. Other problems arise from the highly clustered nature of the sample, high refusal rates, dealing with postal and proxy votes, differential refusal and 'lying' rates (supporters of one party are more likely to refuse to tell the interviewer how they have just voted than supporters of another party or to lie about how they voted), and the sheer logistical problems of processing the results of interviews within the polling hours of 7am to 10pm. Exit polls have been used by ITN in every general election since October 1974 and have been regularly used by the BBC since 1992. They are now the commonly accepted methodology for producing the election-night prior forecasts.

The BBC has regularly used exit polls to predict the results of individual byelections for the Westminster Parliament – see Payne, Brown and Hannah (1986) and Moon and McGregor (1992) for details.

d) Indicator seats

Here there is an attempt to find 'swingometric' seats whose political movements and characteristics match those of the average marginal seat so that their results are taken as indicators of the national political trend. An exit poll is conducted in these seats and the projected changes are used for the whole set of marginal seats. The BBC achieved a spectacular success with the first use of exit polls in the Gravesend constituency at the 1970 election. The predicted swing in Gravesend was very close to the national average, although the use of the exit poll was rather underplayed in the election coverage itself because it was out-of-line with that indicated by the opinion

polls which, as it turned out, considerably under-predicted the swing from Labour to Conservative. Three individual exit polls were also carried out at the October 1974 and the 1979 elections, but they were not very successful. This method is reasonably cheap but has proved to be rather risky. Apart from previous votes, there is little information on which to base the choice of suitable seats. There can be large and unknown changes in the socio-economic composition of individual electorates between elections. So, despite the seductiveness of this approach, it has been abandoned.

5.2 The BBC prior forecasts

The BBC has used information from all the above sources in varying degrees throughout the last three decades. Table 1 gives the errors in the predictions of the majority of the winning party over the runner-up for the prior forecasts produced by the BBC and ITN.

TABLE 1 General Election Prior Forecasts 1974-2001

Errors in forecast of majority of winning party over second party

General Election	BBC	Main BBC source	ITN	Poll-of-Polls
1974 (Feb)	6	Exit polls in 3 seats	50	34
1974 (Oct)	137 (28*)	National exit poll	14	75
1979	4	Poll-of-polls	22	14
1983	2	Quota poll	34	60
1987	70	Quota poll	31	36
1992	62	Marginals exit poll	70	89
1997	29	Marginals exit poll	21	46
2001	6	National exit poll	10	28

Notes:

* Brown/Payne computer-based forecast. The actual forecast used was swingometer-based

1. In all cases the main BBC sources were supplemented by Scottish polls and expert opinion about particular seats for the computer-based forecasts.

2. The Poll-of-polls column gives an estimate of the error if the final poll-of-the-polls had been used, assuming uniform national swing.

The prior predictions obtained in 1983 using a Gallup quota poll were within one seat of the final outcome although post mortem analysis indicated that there were compensating errors in the regional change estimates provided by the poll. The same

methodology was used in the 1987 General Election but this led to some embarrassment with an under-estimation of seventy in the projected Conservative majority over Labour (ie thirty five seats in error). This was primarily due to a poor Gallup poll which, while giving good estimates of the minor party shares, underestimated the overall Conservative-Labour lead by seven per cent, an error which had a large effect on seat projections. Correct poll figures would have given an acceptable error of sixteen in the majority.

The 1992 general election was the pollsters' nadir. Both the ITN and BBC exit polls predicted Conservatives as the largest party, but with no overall majority when in fact Conservatives won an overall majority of 21 seats. The Sky TV exit poll predicted an overall Labour majority. All the pre-election polls predicted a Labour overall majority. There were extensive post mortems by the polling organisations (Market Research Society, 1994). The BBC's own post mortem on its exit poll (Curtice and Payne, 1995) identified the overestimate of the swing from Conservative to Labour by about 2% in its exit poll as contributing about 2/3rds of the error in the underestimate of Conservative seats – although 2% is within the bounds of sampling error expected it was obviously critical in this very close election. They suggested that the main cause was a greater refusal rate among Conservative voters. The remaining 1/3 of the error was attributed to problems with predicting the outcomes in particular Scottish seats.

Given the problems experienced in 1992, the BBC's senior management became very suspicious of exit polls. While it was eventually agreed to do an exit poll in the 1997 election, some care was taken in the design of this poll and new methods were introduced to deal with voters who refused to tell the interviewer how they had just voted. A stratified random sample of 140 constituencies in GB, chosen with probability proportional to size (ie electorate), plus a booster sample of 80 Con-Lab marginals was selected. Wards within each constituency were ranked by Conservative share in the most recent local elections and two polling stations were chosen randomly from each half of the rank order. The pollsters, NOP, developed a method for dealing with refusals where the interviewer essentially guessed the vote of a refusing voter although adjustments were made for the interviewers' performance in making his/her guesses. Another change in methodology introduced was to use the exit poll to estimate the *difference* in swing between the marginals and other types of seat rather than the overall level of swing. This modification assumes that there is a bias in the exit poll estimates of swing, perhaps due to differential lying and refusal rates, but that this bias is constant over all types of seat. Full details of the exit poll and the modifications introduced are given in Brown, Firth and Payne (1999). In the event the exit poll gave good predictions of the final result (a Labour landslide) and was regarded as a success although it should be noted that the 1997 exit poll was no more accurate than the 1992 one in terms of its predictions of the shares of the votes for the main parties. This experience demonstrates the public (and the media) view a forecasters success mainly in terms of getting the right winner!

For the 2001 general election the BBC again revised the design and analysis of its exit poll and employed a statistical modeling approach. This proved very successful and produced a prior forecast very close to the final result of another Labour landslide (as did ITN). The basic approach adopted was to use the same polling stations covered in the 1997 election to predict the change in shares of the vote for the major parties in

each polling station. Ninety of these where no significant changes in the composition of the electorate had occurred and which were considered typical of their constituency were selected for the 2001 exercise. About 180 voters were contacted in each polling station. This provided a dataset of 90 observations, one for each seat polled, which permitted the modeling of the variation of the parties' changes in share using predictors such as region, type of constituency (eg Conservative-Labour marginal, Conservative-Liberal Democrat marginal, safe Labour etc). In the event none of these predictors were significant implying a nationally uniform (and in fact negligible) shift in the parties' fortunes since 1997. So perhaps the 'no-change' 2001 election, where fewer mainland seats changed hands in any general election since the early 1900s, was very easy to forecast and it remains to be seen whether this new methodology works well in elections where more electoral change takes place.

6. Election night: results-based forecast

6.1 General Elections

The second component of election-night forecasting is the production of updated predictions as actual constituency results are received. In the UK the main problem in this respect is that the first constituency results declared do not form a representative sample of all constituencies and may give little information on the key marginal seats. Typically early results come from safe Labour seats in large urban areas, while the results from the constituencies where the election outcome is actually being decided – the Conservative –Labour marginal seats and, in recent elections, the close contests between Conservative and Liberal Democrats, tend to declare later. Furthermore, in recent elections the rise of anti-Conservative tactical voting (a Labour supporter in a Conservative- Liberal Democrat seat might vote Liberal Democrat to defeat the Conservatives and vice versa) has meant that the swings in the safe Labour seats may be untypical of those to be evidenced in the marginal seats. Statistical techniques of varying sophistication have been devised to cope with these problems.

The BBC has adopted a statistical modeling approach to its results-based forecasting on election-night which is set out in detail in Brown and Payne (1975) with revisions introduced for the 1997 general election described in Brown, Firth and Payne (1999). This approach essentially involves the fitting of ridge regression models for the change in share of the vote from the previous election for each major party using the subset of results declared. The predictors used typically include the share of the votes for the party at the previous election, and that of its major competitor, regional and type of contest dummy variables and socio-economic variables (such as % owner occupiers, or % aged over 65) from the last Census if available and current. The coefficients from the regression models are then used to predict the changes in shares of the votes for each major party in the undeclared seats. Then for each undeclared seat a probability that each party standing will win the seat is calculated taking account of the uncertainty of the coefficient estimates (the variance-covariance matrix).

In the early part of election night the results flow is very slow and, as described above, the actual results available are likely to be disproportionately from safe seats. On the other hand there is a high premium on the predictions which get a lot of attention

in the media coverage. In order to counter the possible bias in the declaration order the early forecasts are based on a weighted average of the probabilities from the exit-poll based forecast (which is specifically aimed at predicting electoral movements in the key marginal seats) and the probabilities obtained from the regression analysis of the actual results declared. The BBC has typically used a weighting scheme which gives the prior forecast a weight equivalent to between two and five actual results, depending on the confidence placed in the exit poll. For the 1997 and 2001 general elections the incorporation of the prior information from the exit poll in the results-based forecast was handled in a different and more flexible way. Here the exit poll was used to identify types of seat where the pattern of electoral changes was expected to be different (in 1997 there were four: Conservative-Labour marginals, Conservative – Liberal Democrat contests, Scotland, Rest of England and Wales). Then a pseudo observation for each such type was added to the data set for the regression analysis with each observation containing a prediction of the change in vote share for the relevant response variable. This approach is essentially a pseudo-Bayesian incorporation of prior information used to stabilise early predictions. The approach had the advantage that the forecasters could assess the overall accuracy of the exit poll dynamically and automatically eliminate any bias in the poll from the regression analyses. The statistical details are given in Brown, Firth and Payne (1999)

The other main TV channels which do election-night forecasts (primarily ITN) have never published details of their methodology. It is believed that ITN uses much simpler statistical approach than the BBC. The main elements of their approach are

1. define types of marginal seat
2. conduct an exit poll to provide estimates of changes in vote share in each type and hence a prior seat forecast
3. when a result comes in update the estimate for the relevant category by a weighted sum of the exit-poll based estimates and the average of the changes from all the results declared so far in this category.

Thus this approach puts a lot of weight on the exit poll as the prior estimates for a category are the only source used until an actual results from it is declared. The BBC method borrows strength from results from other types of seats through the regression analyses.

TABLE 2 BBC results-based forecasts for general elections 1974-2001

Errors in majority of the winning party over the second party by number of results declared

Election	0	5	10	50	100	300	Winner	Actual majority
1974 (Feb)	6	2	16	12	4	2	Lab	4
1974 (Oct)	28	13	15	7	8	4	Lab	42
1979	4	13	2	0	9	3	Con	70
1983	2	3	7	10	16	6	Con	188
1987	70	39	46	20	20	5	Con	147
1992	62	49	40	30	36	10	Con	65

1997	25	10	15	5	18	3	Lab	254
2001	6	8	5	10	6	4	Lab	253

Table 2 shows the errors in the BBC's election-night predictions at various stages in the results sequence, with special emphasis on the early stages, for all general elections since 1974. The prediction sequence for the first four elections are very good, perhaps reflecting that as these were easy to forecast as they were essentially two-party contests between Conservative and Labour with the other parties, particularly the then third party, the Liberals, and the Scottish and Welsh nationalists winning very few seats and having a minor influence on the outcome.

But the prediction performance in 1987 and 1992 was not so good and needs comment. In 1987 the BBC had a rather poor prior forecast, as explained above, which was given quite a lot of weight and whose effects took some time to purge. Similarly the 1992 sequence was also very poor (and did not predict the final outcome – a Conservative overall majority - until after 150 results had been declared). ITN had the same experience and their prediction trace was very similar to the BBC's. The main reasons for the BBC's poor performance are set out in their post mortem (Curtice and Payne, 1995) whose conclusion was that the BBC erred by changing its methodology to make it much like that used by ITN described above and abandoning important elements of its regression-based methodology!

Not surprisingly the BBC reverted to its pure regression approach in 1997 and 2001 and this together with the modifications introduced for the design and use of the exit polls resulted in prediction sequences which are perfectly acceptable, particularly in 2001. But both these elections resulted in Labour landslide victories so the real test of the methodology for the now three- and in places four-party system will only come when the election is again closely fought.

6.2 Adaptions for other Elections

(1) European elections 1979, 1984, 1989, 1994

The plurality system was used for these elections with the 81 MPs from England, Scotland and Wales elected from Euro-constituencies which were aggregates of the component Westminster seats (the three Ulster MEPs were elected using STV). The election-night forecasting method used for these elections by the BBC involved a simple modification to the method used in the Westminster Parliament prediction. The votes for the Westminster seats at the previous general election were aggregated to form notional Euro-constituency results and then the standard regression models were used to predict changes at the actual European election from this base. The modifications to the Westminster prediction models for European elections and their application to the forecasting of the 1979 election is given in Payne and Brown (1981).

For all four European elections the results-based forecasts were extremely accurate, being within two seats of the final result from the initial stages. The very large size of the Euro-constituencies, which smooth out individual variations in the component

Westminster seats, seems to bring a large degree of uniformity in their electoral behaviour.

However, prior forecasts are more problematic in European elections as there is often very little prior information available in the form of opinion polls. This is because the elections seem to be of little interest to the media, and also because of the difficulties of identifying likely voters in quota polls when only about thirty to forty per cent of electors will actually vote on election day.

(2) Additional Member System Elections

The elections of representatives to the Scottish parliament (1999), the Welsh Assembly (1999) and the London Assembly (2000) were all run under this system. Voters cast two votes: a *First* (or constituency) vote to elect a directly-elected MP in their seat and a *Second* (or topup) vote for a party from a regional list.

The basic regression modeling method developed for the Westminster Parliament prediction 1997 (Brown, Firth and Payne, 1999) was extended to deal with these elections. The predictions of the winners in the directly elected component was done using the 1997 Westminster Parliament regression methodology, but further models were developed to deal with the predictions of the second vote (including the pseudo-Bayesian incorporation of prior information) and the allocation in the topup seats process. The steps involved were:

1. Fit regression models for the change in share of the *first* vote from the previous general election (1997) for the main parties to give predicted directly-elected winners for each seat and the number of these for within each region. This step corresponds to the normal Westminster procedure.
2. Fit a regression model to predict the turnout in the undeclared seats to give a prediction of the total numbers voting in each such seat.
3. Fit regression models for the change in share of the *second* vote from the previous general election for the main parties to give a prediction of the shares of the second vote in the undeclared seats.
4. Using the predictions of turnout and vote shares for undeclared seats from steps 2 and 3 (or the actual voting figures for declared seats) calculate the predicted *numbers* of second votes for each party in each seat and sum over regions to give predicted second votes for each party in the region.
5. For each region apply the d'Hondt rule using the directly-elected seat predictions from step 1 as the initial denominator to predict the allocation of the topup seats.
6. The final seat prediction are the sum of the predicted directly-elected seats (step 1) and the topup seats (step 5)

The Scottish and Welsh elections were based on the existing Westminster seats in their country, so the 1997 general election results from which base the changes in

party shares are modeled were readily available. For the London Assembly election notional 1997 general election results for the 14 London Boroughs involved had to be constructed from the 74 component Westminster seats.

A new challenge for the pollsters with these elections was to accurately estimate the distribution of second (regional list) votes as well as the distribution of first (constituency) votes. As this was a new type of election there was some doubt that electors had fully understood the system. While many might vote the party-ticket with both votes it was expected that others might take the opportunity to cast their second vote for other candidates from minor parties, who while having little chance of success in the direct election could have good prospects in the top up where the threshold for winning is rather lower. It was also expected that some voters might cast only one of the two possible votes. In the event very few polls were actually published. The BBC commissioned a poll from ICM for its prior forecasts in each election. While the predictions of first vote shares were generally within the expected range of error these polls tended to underestimate the shares of the second vote, particularly for the minor parties standing.

For all three elections the results-based forecasts were extremely accurate, being within two seats of the overall result for each party from the initial stages. A helpful feature of this system for the forecaster is that any errors in the forecast of the directly-elected MPs (step1) tend to be compensated for in the allocation of top-up seats (step 5) with the result that the overall prediction is very good.

(3) Regional list: European election 1999

This election is essentially a special case of AMS – there are no directly-elected MPs but the d'Hondt rule is used to allocate seats on the basis of total votes for the parties within regions. So the extensions to the basic statistical modeling methodology developed for AMS set out above can be used (steps 1 and 6 are not required for this system). This is probably the easiest electoral system to predict as results are declared on election night as it combines the feature of large electorates of previous European elections where individual 'Westminster' idiosyncracies are evened out and the compensating feature of the topup seat allocation process from AMS. It is of course a system which is designed to give a high degree of proportionality so that there should be close relationship between the parties overall share of the vote and their share of the elected MEPs.

On election night results were announced for each component Westminster seat, so that change in parties' share and the turnout in each seat from the previous Westminster election (1997) could be modelled to give predictions of the number of votes cast for each party in each undeclared seat. These were then aggregated over regions to give predicted party votes in each region and the d'Hondt rule applied to allocate seats. The BBC's on-the-night forecast was very accurate being within three seats of the final figures for each party after 7 of the 659 results had been declared and within one seat after 100 results. An interesting feature of this election was that two minor parties the Greens and the United Kingdom Independence Party were successful (with two and three MEPs respectively) reflecting the proportional properties of the Regional List system. These parties would not have won any seats in

a Westminster election with the shares of the vote they obtained in the European election. On election night these successes were the most difficult to predict as the parties gained them towards the end of the topup process and accurate predictions of regional vote distributions were required to forecast their success – this happened after about 100 Westminster seats had been declared.

(4) The Referendums on devolved government in Scotland, Wales and London

Voters were asked whether or not they supported the government's proposals for devolved bodies in each area (there was an additional proposition in the Scottish referendum on tax raising powers for the proposed Scottish parliament). So the forecasting aim was to predict the percentage of voters voting 'Yes' and whether and when during election night, the critical figure of 50% would be reached. Results were declared for each Local Authority (Council): Scotland (32), Wales (22) and London (33). Notional 1997 general election results were constructed for each Local Authority from the component Westminster constituencies.

The predictions for these elections involved straightforward modifications to the standard regression 'Westminster' modeling method (Brown, Firth and Payne, 1999), using pseudo-Bayesian incorporation of prior information from specially commissioned polls. Regression models for the percent voting Yes and the turnout were fitted using as predictors selected party shares from the notional 1997 general election result (eg Plaid Cymru support in Wales), census variables (eg % young voters in Scotland) and location of the Authority (eg whether in Inner or Outer London in the London referendum). These were then used to predict the numbers voting 'Yes' in each Local Authority area which were finally summed to give the prediction of the overall numbers supporting and opposing the proposition and hence the percent voting Yes.

The traces of the prediction sequences produce on election night (available from the author) show a significant improvement over the predictions which would have been obtained if the aggregated actual results only had been used. In Scotland the predictions produced were within 0.5% of the final figure of 74.3% yes for the Scottish Parliament proposals after 6 of the 33 results had been declared. In London the percent supporting the proposals for a London assembly and Mayor was predicted to within 2% of the final figure of 72.0% after 10 results had been declared. In the Welsh referendum predictions were within 1% of the final figure of 50.3% Yes after 10 results.

However the Welsh referendum provided a salutary lesson for the presenters of forecasts. The BBC's editorial policy at the time was to show the predicted percentage voting Yes only when the actual outcome was known (ie when it was known whether the Yeses or the Noes had finally won). In this very close contest this outcome was only known when the last of the 23 results was declared! So all the viewer had was a sequence of statements that 'it is too close to call' whereas they would have been better informed if they had been presented with the actual forecasts flipping just above and below the 50% line.

(5) Supplementary vote: The London Mayor

This election took place at the same time as the election of members to the London Assembly (see above) so that voters specified their first and second preferences for Mayor as well as casting their first and second votes for Members of the London Assembly at the same time. The Mayoral results were declared for each of the 33 Local Authority areas giving the number of first preference votes and the number of second preference votes for each of the 11 candidates.

This election involves two prediction stages the second of which is only required if no candidate gets over 50% of the first preference vote: (1) a forecast of the overall shares of the vote for each candidate (2) allocation of second preference votes for eliminated candidates to the top two remaining candidates. As each major party put up a candidate the obvious modeling strategy for the first stage was to use the 'Westminster 1997 type' regression models to predict the shares of the vote for the major parties' candidates and turnout and thus the total votes for the candidates for London as a whole. However the likely (and eventual) winner, Ken Livingstone, standing as an independent, was a former Labour MP, and an additional regression model was required to predict his vote. The BBC predictions for the first preferences for the top two candidates (Livingstone (Independent) who got 39.1% and Norris (Conservative) who got 26.5%) were within 1% for each of these candidates from 5 declared results onwards. So the second stage had to be invoked. For this the BBC used a 11 by 12 matrix of first preference vote by second preference vote (with an additional column for no second preference stated) provided by the pollsters ICM who had interviewed about a 1000 London electors before the election. This matrix gave good predictions of the extra votes actually gained by the clear winner Livingstone (6.3%), and Norris (5.8%).

6.3 Discussion

For most of the elections described the BBC's results-based forecasts have settled down quickly to a position quite close to the final result. The exceptions are the general elections of 1987 and 1992 where poor prior forecasts took some time to be purged. In 1987 this was due to overconfidence in the prior, perhaps understandable given the unrivalled accuracy of the prior forecast in the previous general election. In 1992 a change in the prediction method, imposed by senior BBC management, also effectively put too much weight on to the exit poll.

But it is clear that statistical modeling provides a comprehensive and adaptable methodology bringing improved performance over cruder psephological methods on election night. The main problem has been prior forecasts where the statistical modelers are to some extent at the mercy of the pollsters and voters exiting from polling stations. It remains to be seen whether that the BBC has developed an effective method of dealing with exit poll bias through the use of the pseudo-Bayes method, first used in 1997 (and enhanced in the 2001 general election) and then in every subsequent election in Great Britain. Only time and a close election or two will tell.

7. Media Aspects

The main objective of this paper has been to give a statistical treatment of election forecasting in the UK. However, the primary use of election-night forecasting is by the media and this has implications for both the way such forecasts are used and presented. There are also constraints on what can be done because of the hectic election-night context. Some implications of the context are:

1. There is a high premium on results-based forecasts only when the election is close (only in the 1974 and 1992 general elections and the Welsh Referendum in 1997). There is always a very high premium on the prior forecast which plays a key role at the start of the election night coverage, filling the gap between the end of voting and the declaration of results. There is a strong demand for frequent predictions in the very early stages when the forecasts are likely to fluctuate a lot.
2. There is a reluctance to present the uncertainty associated with the forecasts. This could be done in various ways such as confidence intervals or probability distributions for seat totals, probabilities or odds for parties winning in each constituency and odds on particular global outcomes, such as who will be majority party. This reluctance has several explanations. One is the discrete nature of maps which are used to present predicted outcomes in individual seats and there is a requirement to use bold graphics (Labour seats all in red, Conservatives all blue..), without shading to indicate uncertain outcomes. A second explanation is a deterministic view commonly held by presenters (...if the average swing to Labour is X then all seats where Labour is less than X/2 behind *will* be won by Labour...). But most importantly there is a commonly held view that the viewer will not understand betting odds or probabilities, let alone probability distributions.
3. Forecasts are now presented in increasingly sophisticated dynamic graphical displays and the on-line prediction system is but a small component in the election-night computing system. Technical requirements impose constraints on the frequency with which forecasts can be updated and on the sophistication of the statistical method used. This is because the prediction subsystem must both compete and integrate with other components of the election-night computing system. In particular, a complex interface is required to pass individual seat predictions to the graphical display subsystem and to make sure that the two subsystems are in phase.

As a result of these media aspects, *Payne's Eleven Laws* of media election forecasts have been formulated:

1. Few people understand that UK forecasting is more difficult than in most other electoral systems
2. Prior forecasts are at the mercy of the pollsters
3. Good forecasts are soon forgotten
4. Bad forecasts are forever remembered and reminded (even by Prime Ministers)

5. You are only as good (or as bad) as your last forecast
6. As long as you get the winner right the media won't be too concerned about your accuracy.
6. Only statisticians can understand why ridge regression should always beat the swingometer
7. There is a positive correlation between the proportion of the TV audience watching your channel and the error in your forecast.
8. Hying good forecasts always results in a red face for the subsequent one.
9. Expressing uncertainty about forecasts is prohibited even though the information on which they are based is allowed to be subject to error.
10. Graphics (and their presenters) rule!!
11. Relative performance (i.e BBC versus ITN) rather than absolute performance also matters.

References and bibliography

- Brown, P. and Payne, C. (1975) Election night forecasting (with discussion), *J. R. Statist. Soc. A*, 138, 463-498.
- Brown, P. and Payne, C. (1984) Forecasting the 1983 British General Election, *Statistician*, 33, 217-228.
- Brown, P., Firth, D. and Payne, C. (1999) Forecasting on British Election Night 1997, *J. R. Statist. Soc. A*, 162, 211-226.
- Curtice, J. and Payne, C. (1995) Forecasting the 1992 election: the BBC experience in *Political Communications: the British General Election of 1992* (eds I. Crewe and B. Gosschalk). Cambridge: Cambridge University Press.
- Curtice, J., Payne, C. and Waller, R. (1991) Local elections as national referendums in Britain, *Electoral Studies*, 10, 3-17.
- Karandikar, R., Payne, C. and Yadav, Y. (2001) Predicting the 1998 Indian Parliamentary Election, *Electoral Studies*, 16, ???-???
- Kendall, M. and Stuart, A. (1950) The law of cubic proportions in election results, *Br. J. Sociology*, 183-96.
- Market Research Society (1994) *The Opinion Polls and the 1992 Election: a Report to the Market Research Society*. London: Market Research Society.
- Milledge, D. and Mills, M. (1960) Forecasting election results, *Computer Journal*, 2, 195-8.
- Moon, N. and McGregor (1992) Exit polling: developing a technique. *J. Markt Res. Soc.*, 34, 257-268.
- Morton, R. Ed. (1990) Election night forecasting, *The New Zealand Statistician*, 25, 36-84.
- O'Muircheartaigh, C., and Lynn, P. (1997). Editorial: the 1997 UK Pre-election Polls, *J. R. Statist. Soc. A*, 160, 381-385.
- Payne, C. (1992) Statistical methods for election forecasting in the United Kingdom, 1970-90 in *British Elections and Parties Handbook 1992* (eds P. Norris, I. Crewe, D. Denver and D. Broughton). Hemel Hempstead: Harvester Wheatsheaf.
- Payne, C. and Brown, P. (1981) Forecasting the British election to the European parliament, *Br. J. Polit. Sci.*, 11, 235-245.
- Payne, C., Brown, P. and Hannah, V. (1986) By-election exit polls, *Electoral Studies*, 5, 227-87.

Rallings, C. and Thrasher, M. (1999) *New Britain: New Elections*. London: Vacher Dod.

Sanders, D. (2000). Conservative incompetence, Labour responsibility and the feelgood factor: why the economy failed to save the Conservatives in 1997, *Electoral Studies*, 18, 251-270.

Smith, F. (1982) Election forecasting: transferable vote systems in *Encyclopedia of Statistical Sciences* (Eds S. Kotz and N. Johnston), New York: Wiley.

Smith, T., M., F. (1996). Public Opinion Polls: the UK general election 1992, *J. R. Statist. Soc. A*, 159, 535-545.

Worcester, R. (1996) Political Polling: 95% expertise and 5% luck. *J. R. Statist. Soc. A*, 159, 5-20.