Electoral Studies 30 (2011) 258–263

Contents lists available at ScienceDirect

Electoral Studies

journal homepage: www.elsevier.com/locate/electstud

Yes, Prime Minister: The key to forecasting British elections

Matthew Lebo*, Helmut Norpoth

Department of Political Science, SUNY at Stony Brook, Stony Brook, NY 11794-4392, USA

ARTICLE INFO

Article history:
Received 17 September 2010
Accepted 22 September 2010

Keywords:
British election
Forecasting election
Prime Minister approval

ABSTRACT

We use our "PM and Pendulum" Model to forecast the outcome of the 2010 General election. The vote function of the model, aside from a cyclical dynamic, relies on approval of the prime minister as the sole predictor. We find that PM Approval predicts the vote (and vote intention between elections) more accurately than does Government Approval. Turning to the forecasting of seats, we examine the accuracy of the autoregressive model of the vote–seat translation against the uniform-swing model, which is widely used by pollsters and the media. Testing the alternatives on election data since 1910, our autoregressive vote–seat translation model proves superior to the uniform-swing model.

Our "PM and Pendulum" Model, which got its first test in the 2005 election, relies on the approval of the prime minister as the sole predictor of the vote in British general elections, aside from the cyclical dynamic.[1] Though risky, such a minimalist choice earns points for parsimony. No other single factor, we contend, does a better job encompassing short-term forces in a British election. In this paper, we test this proposition against a close competitor, approval of the "government." Turning to the forecasting of seats in parliament, we examine the accuracy of our autoregressive model of the vote–seat translation against a popular alternative: the uniform-swing model. In so doing, we also extend the time horizon of the vote–seat function back in history to include a full century of data. We also offer separate forecasts for Labour and the Conservatives instead of just a forecast of the seat lead for the winning party.

1. PM Approval vs. Government Approval

With the infrequency of elections, monthly data of party vote intentions give us a valuable way of measuring the effects of independent variables of interest. Studying the effects of pre-election Prime Ministerial Approval data over 17 elections provides a useful, but limited, test of the variable's accuracy. Looking instead at an uninterrupted monthly version of the time series from 1979 to 2009 provides significant leverage on the question of just how good a job it does predicting vote outcomes.

We employ MORI's monthly time series data for the period of September 1979–January 2009.[2] Three variables are of particular interest: vote intentions for the incumbent party, approval of the incumbent party and approval of the Prime Minister. The first series is created by using the monthly percentage of those naming the governing party when asked: 'How would you vote if there were a general election tomorrow?' For Government Approval we use the monthly percentage of those answering satisfied to the questions: 'Are you satisfied or dissatisfied with the way the government is running the country?' and for Prime

* Corresponding author.
E-mail address: matthew.lebo@sunysb.edu (M. Lebo).

[1] Lebo, Matthew J., Norpoth, Helmut. 2006. The PM and the pendulum: dynamic forecasting of British elections. British Journal of Political Science 37, 71–87.

[2] Data are available at http://www.ipsos-mori.com/researchspecialisms/socialresearch/specareas/politics/trends.aspx. MORI data prove most useful because of their consistency in question wording and their consistent polling both before and after the end of British Gallup data in 2000. Economic data are from http://www.statistics.gov.uk/.

0261-3794/$ – see front matter © 2010 Elsevier Ltd. All rights reserved.
doi:10.1016/j.electstud.2010.09.004

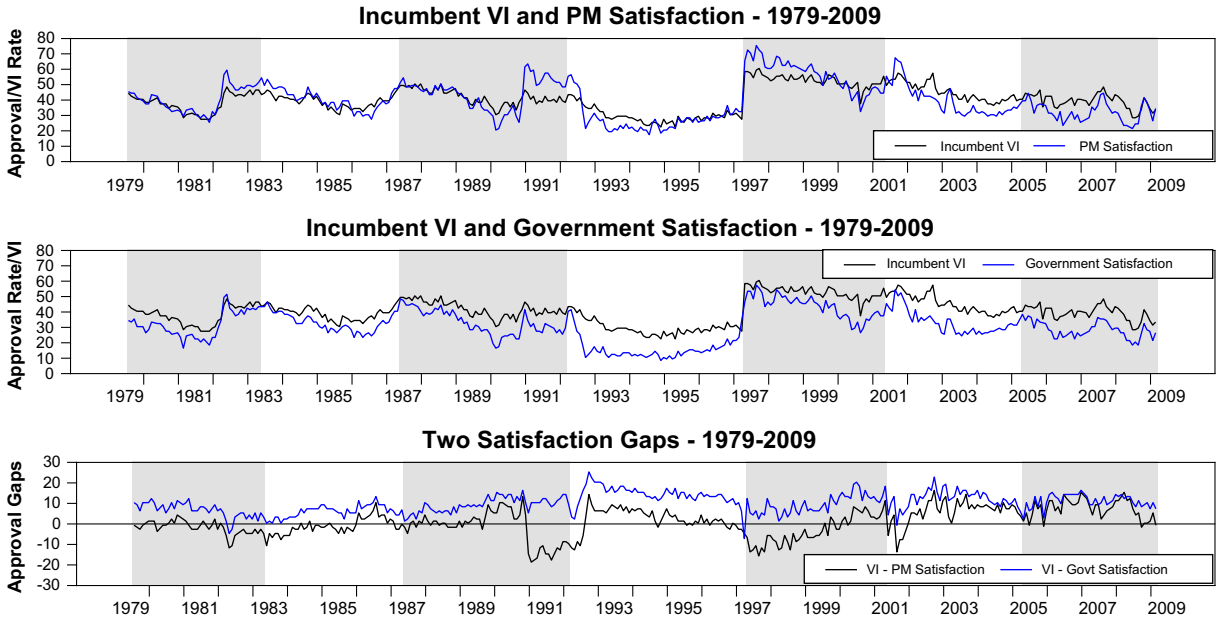

Fig. 1. Incumbent vote intentions and two approval measures.

Minister Approval we use a similar index compiled from MORI's PM satisfaction question.[3]

Fig. 1 shows the two approval measures each side-by-side with vote intentions for the incumbent party. A careful look at the top two panels of the figure will reveal what we will demonstrate empirically below, that incumbent vote intentions are more closely related with PM Approval than they are with Government Approval.

This is not an intuitive point. Elections in Britain are certainly about choosing a Prime Minister but they are thought to be first and foremost about choosing a government. It is not necessary here to outline the differences between a parliamentary and a presidential system, yet, it is an interesting practice in British politics that a party can make a change in the Prime Minister (as in the transfers of power from Mrs. Thatcher to Mr. Major and from Mr. Blair to Mr. Brown) and the new PM does not feel compelled to immediately consult the electorate.[4] This might underscore a belief that, in a parliamentary system, retrospective voting is primarily about rewarding and punishing the government and that the Prime Minister is only of secondary concern.

Yet, as the bottom panel of Fig. 1 shows, the differences between vote intentions and the two approval measures clearly favors PM Approval as the closer predictor. For nearly every month of our 30-year sample, the gap between PM Approval and Incumbent VI is the smaller of the two. Next, to test this hypothesis, we construct two popularity functions identical in every respect except for the matter of which approval measure is used.

Table 1 presents two multivariate ARFIMA models of governing party vote intentions.[5] The left-hand side model uses Prime Ministerial Approval as an independent variable and specifies a (fractional) error correction mechanism (FECM) that is created from the residuals of a regression between Incumbent Vote Intentions and the PM variable. The model on the right-hand side is identical except for the fact that is uses Government Approval as an independent variable and as a component of the FECM.

The results of the two popularity functions are very similar but it is in the comparisons of the two models that we can see the advantages of the PM variable over the Government Approval variable. At first glance, Government Approval has a higher coefficient than does PM Approval and this might lead one to think the latter is a superior predictor to the former. But, given the much higher variance for PM Approval (149.10 vs. 113.44 for Government Approval), the coefficients are not directly comparable. Thus, the bottom of Table 1 summarizes seven different test statistics, each of which points to the PM model as the stronger of the two.

Specifically, in the PM model, the *t* statistics of both the key variable and the FECM are higher than their counterparts in the Government model indicating a greater impact on vote intentions for the PM variable. Also, the model evaluation statistics all point to the PM model as the superior one:

[3] Although other variables such as the inflation rate, the unemployment rate, interest rates, and subjective measures of the economy have been used as predictors in popularity functions and in forecasting models, only Government Approval approaches the PM variable in terms of its predictive ability. Thus, we limit our comparison to these two measures.

[4] Although as our forecast would have predicted and as the MORI data demonstrate, this would have been a very good idea for Prime Minister Brown.

[5] The models and error correction procedures follow Clarke and Lebo (2003). For further explanation and a defense of these methods see: Young, E., Lebo, M.J., 2009. Long memory methods and structural breaks in public opinion time series. *Journal of Elections, Public Opinion and Parties* 19, 117–124.

Table 1
Two popularity functions compared, 1979–2009.

Independent variable	PM Approval model		Government Approval model	
	Coef. (s.e.)	*t*	Coef. (s.e.)	*t*
PM Approval	0.410 (0.026)	16.04***		
Govt. Approval			0.443 (0.030)	14.83***
(F) ECM	−0.359 (0.041)	−8.74***	−0.349 (0.042)	−8.33***
Major in	−4.450 (2.098)	−2.12*	−0.183 (2.122)	−0.09
Cameron in	−9.035 (1.799)	−5.02***	−8.563 (1.859)	−4.61***
Black wed.	5.310 (1.860)	2.86**	2.831 (1.891)	1.50
Poll tax	−2.193 (1.275)	−1.72*	−0.988 (1.319)	−0.75
Falklands	2.259 (1.058)	2.13*	2.061 (1.092)	1.89*
Berlin Wall	5.649 (1.804)	3.13***	4.295 (1.858)	2.31*
Gulf War	5.695 (0.949)	6.00***	0.931 (0.936)	0.99
Sept. 11 attack	−6.137 (1.867)	−3.29***	−4.754 (1.908)	−2.49**
Aug–Oct 2002	4.889 (1.051)	4.65***	4.352 (1.079)	4.03***
Iraq War	−6.250 (1.274)	−4.91***	−5.883 (1.315)	−4.47***
Fuel crisis	−7.371 (1.832)	−4.02***	−9.604 (1.870)	−5.14***
1997 Election	16.375 (1.99)	8.22***	21.259 (1.958)	10.86***
2001 Election	6.817 (1.282)	5.32***	6.361 (1.319)	4.82***
$\text{Unemployment}_{t-1}$	−2.456 (0.885)	−2.76**	−2.151 (0.911)	−2.36**
Inflation	−0.475 (0.208)	−2.28*	−0.617 (0.216)	−2.86**
Inflation_{t-1}	−0.396 (0.207)	−1.91*	−0.527 (0.214)	−2.46**
Inflation_{t-3}	−0.446 (0.199)	−2.24*	−0.373 (0.206)	−1.81*
Mori EOI_{t-1}	−0.010 (0.011)	−0.92	−0.006 (0.012)	−0.55
Constant	−0.426 (0.158)	−2.70**	−0.558 (0.163)	−3.41***
N	352			352
Durbin-Watson	2.02			2.02
Model comparison statistics				
R-squared	0.71			0.69
SS residuals	1059.96			1127.43
Standard error of estimate	1.789			1.846
Akaike information criterion	2494			2515
Schwartz criterion	2575			2596
Residual mean square	0.092			0.095
Davidson *J*-test *t*	6.237			4.078
T-statistic	16.04			14.83
FECM coef. (*t*)	−0.359 (−8.74)			−0.349 (−8.33)
Variance of DV	149.10			113.44

* Significant at the .05 level; ** significant at the .01 level; *** significant at the .001 level.

a lower standard error, a higher R^2, lower Akaike Information Criterion, Bayesian Criterion, and residual mean square all favor the PM model. Lastly, Davidson's *J*-test for encompassing establishes that the PM model explains all the variance explained by the Government model plus some additional variance that remains unexplained if PM Approval is excluded.[6] In sum, all of our statistical tests support the model that includes the PM Approval variable. Thus, we are confident that, in our attempt to keep our model as simple as possible, a measure of PM Approval is the best choice.

2. Forecasting the vote

Besides Prime Ministerial Approval, our vote model relies on a cyclical factor. As we have previously demonstrated, the major-party vote swings like a pendulum in some fashion (Lebo and Norpoth, 2006). The estimate of a cyclical swing of the vote comes from a second-order autoregressive process, which mimics a cyclical movement when the first parameter is positive and the second negative. The swing of the electoral pendulum that will largely define the 2010 election has long been forecast by our second-order autoregressive parameter. Indeed, the large movement away from Labour is something we would have predicted immediately after the 2005 election – the size of Labour's victory in 2001 should produce a huge swing in the opposite direction two elections later. With Labour having been in office three terms by now, the cyclical factor, all by itself, forecasts a Tory win in the next election. The vote lead over Labour would be 2.5 percent. Adding Prime Ministerial Approval as a vote predictor leaves the cyclical dynamic undisturbed (Lebo and Norpoth, 2006, Table 2). The values used for our forecast are based on MORI and YouGov's March 2010 polls that average 34.5% for PM Approval and 69% for the two-party vote. This gives us an adjusted PM Approval number of exactly 0.0 and translates into a predicted vote lead of 4.2% for the Tories over Labour.

3. Autoregressive vs. uniform-swing models for seats

In developing a forecast model for seats, there is no reason to limit ourselves to elections since 1945, the period for which we have data on PM Approval as a predictor of the vote. Votes and seats can be related as far back as 1832 in British elections. With a larger universe of elections, the

[6] This common test for model comparison is explained in: Hamilton, J. D., 1994. *Time Series Analysis*, Princeton University Press, Princeton.

Table 2
Forecasting seats from votes in British elections: the Conservatives & Labour/Liberals.

	Conservatives			Labour/Liberals (until 1918)		
	Elections 1910–2005	Excluding 1931	Uniform-swing	Elections 1910–2005	Excluding 1931	Uniform-swing
Conservative Vote Lead	9.270*** (0.329)	9.531*** (0.558)	9.429*** (0.446)	−8.661*** (0.420)	−8.879*** (0.604)	−8.713*** (0.439)
AR(1)	0.576*** (0.224)	0.579*** (0.216)	1	0.690*** (0.158)	0.674*** (0.156)	1
Constant	278.055*** (16.197)	277.701*** (16.038)	0.679 (6.778)	307.778*** (10.634)	307.992*** (10.193)	−0.072 (4.303)
Standard error (residuals)	30.679	31.112	33.925	18.596	18.783	21.457
Adjusted R^2	0.88	0.83	0.85	0.96	0.94	0.94
LBQ (k autocorrelations)	11.3 (6)	10.1 (6)	20.4 (6)	6.2 (6)	5.9 (6)	9.7 (6)
p(LBQ)	>0.08	>0.12	>0.01	>0.40	>0.43	>0.14
N	26	25	26	26	25	26

Note: The Conservative Vote Lead is the difference between the percentages of the Conservative vote and the Liberal vote until 1918 and the Labour vote afterwards. For every election, the number of Conservative seats has been adjusted to account for the changing size of the House of Commons. The adjustment factor is the ratio of 650 to the number of total seats in a given election. In this way, each election is treated as if it had the same total number of seats as the one being elected in 2010 (650). The AR(1)-parameter in the "uniform-swing" model is not estimated but set to 1, which implies that the variables are treated as first differences. Parameter estimation was done with robust standard errors. The Labour/Liberal Seat variable uses Liberal seats until 1918 and Labour seats afterwards.
***$p < 0.001$ (One-tailed).

model estimates can be expected to be more precise, robust, and reliable, at least in theory. Over a long period, of course, with changing electoral conditions, the vote/seat function might not be constant. It would be nice if we were able to use as many of the 40-plus elections since 1832, but that may not be possible given a minimal standard of fit and accuracy.

As shown in Fig. 2, the voter/seat function for elections in the 19th century, roughly speaking, differed markedly from the one since then. It became steeper, stronger, and less biased, in a partisan sense. Note that the vote axis of the figure represents the Conservative vote lead while the seat axis represents the Conservative seat lead over the other major party (Liberals until 1918, Labour afterwards). The cut-point between the two eras in this case turns out to be 1910.[7] The correlation between the vote and seat variables rises steadily as we exclude the earliest case from the full set of elections. Once we get to 1910, the (squared) correlation reaches a peak of 0.94, which remains unsurpassed with the exclusion of further elections. For elections prior to that point (1832–1906), the corresponding fit is barely half (0.49). What is more, as shown in Fig. 2, the vote/seat slope in the earlier period (7.5) is less than that (18.0) of the later period. Also note that in the earlier period the Tories would gain more seats (a positive seat lead) with a substantial deficit in the vote (a negative vote lead of up to 10 percent). This pro-Tory bias disappears in the later period; if anything, it reverts to a slight disadvantage.

These changes in the vote/seat function, we suspect, have much to do with the disappearance of uncontested constituencies, most of which were held by the Tories. So long as many MPs get elected without opposition, the vote in the remaining districts that are being contested should not be expected to be a reliable barometer of seats in the whole chamber, especially if one party controls the lion's share of such districts. While it was not uncommon in the 19th century for about half of the MPs to be elected without a contest, such MPs became a rarity in the 20th century. The change was especially abrupt between the elections in 1900 and 1910 (January), as the number of uncontested seats dropped from 243 to 75 out of a total of 670.[8] With nearly all constituencies being contested by the major parties, the national vote should be expected to provide a reliable predictor of seats.

Besides the vote in a given election, a seat-forecast model also has to reckon with the inertia of seats from the previous election. To take a simple case, assume no swing of the vote between two elections, in which case one would predict the same seat distribution as in the last election, everything else being equal. The vote matters for seats only insofar as it deviates from the previous election. As the vote pendulum swings, the favored party adds seats above its count in the previous elections while the other one sees a subtraction of its stock. A standard model, which assumes the existence of a "uniform-swing" across all constituencies,

---

[7] There were two general elections in 1910, one in January and the second in December. For convenience sake, any time we refer to a 1910 election in this paper we mean the *January* 1910 contest.

[8] Butler, D., Sloman, A., 1975. *British Political Facts*, Macmillan, London, p. 183.

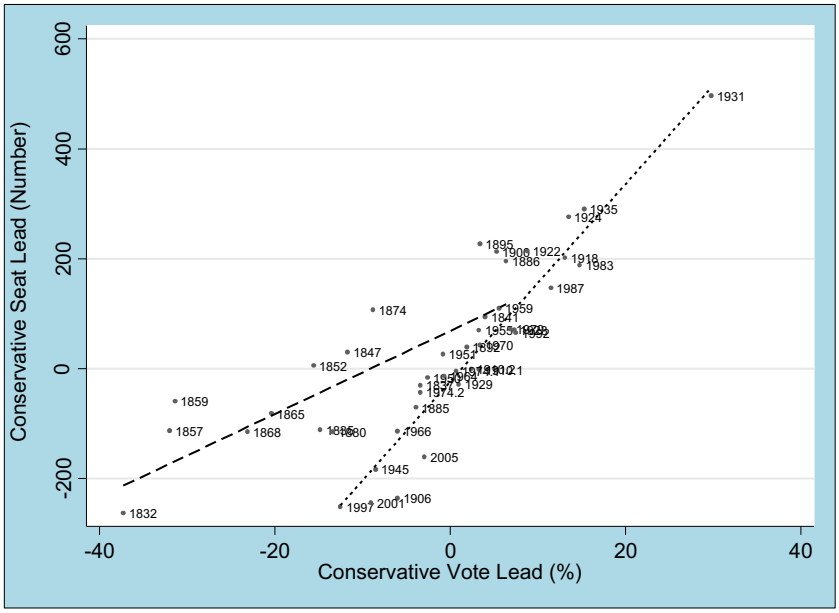

Fig. 2. Vote to seat translation 1832–2005, before and after 1910.

predicts seats with such a calculus. While this may be too oblivious to local conditions that favor a party in a constituency in a given election, the predictive power of past seats should not be denied. Our vote/seat model used for the 2005 election incorporates this predictor by means of an autoregressive process. This allows for more flexibility than the process implied by a "uniform-swing," which relates votes to seats in the form of first differences.

Covering the time horizon from 1910 to 2005, we have estimated an AR(1) vote/seat model each for the Conservatives and their major-party opponent (Liberals until 1918, Labour afterwards). Note that the number of seats of a party in a given election from 1910 to 2005 has been adjusted to account for the changing size of the House of Commons. The adjustment factor is the ratio of 650 to the number of total seats in a given election. In this way, each election is treated as if it had the same total number of seats as the one being elected in 2010 (650). Turning to the Conservatives first (left-hand side of Table 2), the findings for the 1910–2005 elections confirm a moderately strong autoregressive effect on the number of seats in a given election; beyond that a one-point increase in the Conservative vote lead converts to nearly ten more seats.

The corresponding findings for Labour (Liberals until 1918), shown in the right-hand side of Table 2, prove to be similar in size, though with the reverse sign for the vote variable (Conservative lead). Where the findings for the major contenders differ, in a politically consequential way, is in the constants of the vote/seat equation. These coefficients indicate the number of seats to be expected for a zero vote lead, which would occur if both major parties obtained the same vote share. In that event, the Tories would expect to come up with 278 seats compared to 308 for Labour (Liberals until 1918). In other words, the Tories would gain 30 fewer seats for the same vote share as their major-party opponent. This estimate is a 1910–2005 average, to be sure. Some observers contend that the partisan bias has changed in recent elections.[9] Our estimates nonetheless are consistent with their finding that the vote/seat bias favors Labour these days.

We put the autoregressive model to the test against the "uniform-swing" model. As can be seen in Table 2, the latter fares less well for both the Conservatives and Labour/Liberals than the autoregressive model. The premise that a change in the vote translates in uniform fashion into a change in seats is too rigid. The autoregressive model is more adept at capturing the random element of both votes and seats.

Looking to the 2010 election, our model forecasts 285 seats for the Tories and 287 for Labour, given our March 2010 forecast of a 4.3 percent Tory vote lead. A hung parliament. A coin toss for the largest party.

Of course, being just right well in advance would be an impressive feat. Uncertainty exists and we quantify it to estimate the likelihood of various outcomes. Uncertainty comes from several places: the PM Approval value, the impact of PM Approval on votes, the historic movement of votes and seats, and the translation of votes into seats. Allowing for uncertainty, we simulate 1,000,000 elections under the many ways the model *could* truly work and the many values our data *could* hold. The results of these simulations tell us about the probabilities of different outcomes.

The range of possibilities is, of course, very wide. But some possibilities are more possible than others. Although the Tories lead in 47.7% of the simulations, in only 4.47% of them do they win a majority (>325) of seats. That is, even

[9] See Johnston, R., Pattie, C., Rossiter, D., Dorling, D., MacAllister, I., Tunstall, H., 1999. New Labour's landslide and electoral bias: an exploration of differences between the 1997 UK General Election result and the previous thirteen. In: Fisher, J., et al. (Eds.), *British Elections and Parties Review 9*, Cass, London.

with a prediction of a 4.3% vote deficit, the model predicts a 52.3% chance that Labour will hold onto a lead in seats. Still, the chances of a Labour majority are very small (1.33%). The range of vote differentials that can produce a hung parliament is wide and our distribution is smack in the middle of it meaning that the probability of a hung parliament (94.2%) is near its maximum. Any change in Gordon Brown's popularity will drop this number substantially.

4. Conclusion

The "PM and the Pendulum" Model has fared well in tests against alternatives. Prime ministerial (PM) approval proves to be a superior predictor than Government Approval, a close competitor. This is true not only for the vote on Election Day, but also in polls between election years, which offers a far richer testing ground. PM Approval is the most accurate predictor of aggregate level trends in British vote intentions. As for forecasting the seats of the major parties in parliament, the autoregressive vote–seat translation model proves superior to the uniform-swing model, a popular alternative that is widely used by pollsters and the media. This conclusion is based on statistical tests using elections as far back as 1910. All told, our forecast of the vote in the 2010 British general election is derived from a model with the same predictors and structure as was used for the 2005 election. For the seats, the predictors are also the same as in 2005, but we have derived the estimates from a longer time series, beginning with the 1910 (January) election instead of 1945, as before. We also make forecasts for each of the two major parties separately instead of a forecast for the seat lead. This does not alter the model, but simply makes forecasts that are more useful. In particular, our 2010 seat forecasts indicate whether the winning party is able to capture a majority of seats or not. Even with the swing of the pendulum, the bias of the electoral system against the Conservatives and Mr. Brown's middling popularity make a Hung Parliament look like a near certain outcome.

Reference

Clarke Harold, D., Lebo Matthew, J., 2003. Fractional (co)integration and governing party support in Britain. British Journal of Political Science 33, 283–301.