

# Policy Sentiment in the UK

James A. Stimson  
University of North Carolina at Chapel Hill

March 9, 2006

Prepared for the Department of Government, Essex University, March 9, 2006. This is a large scale revision of an earlier paper prepared for the Elections, Public Opinion and Parties 2005 Annual Conference, September 9-11, Wivenhoe House, University of Essex. I am grateful to Chris Wlezien and Stuart N. Soroka for providing me their data on attitudes toward government spending in the UK.

**Abstract**

In this paper I lay out a concept of policy sentiment as an underlying attitude explaining observed respondent preferences on policy preference questions. The latent concept, left vs. right sentiment, is extracted from policy preference marginals. The extracted time series characterizes the last 45 years of UK political experience. Two analyses focus on the causes of policy sentiment—thermostatic response to eras of party government—and its consequences for aggregate election outcomes in the UK.

# 1 Policy Sentiment: The Concept

Mass publics get to choose, in surveys directly and in politics indirectly, what they wish government to do. Should the Blair Government improve NHS benefits (at some considerable cost) or rein in spending to permit tax reductions. Should there be more or less state ownership of industry. Should People in Britain be more tolerant of those who lead unconventional lives? Hundreds of such queries are posed.

If we wish to explain the answers that emerge from survey respondents, it is useful to have a model of the survey response. Zaller (1992) and Zaller and Feldman (1992) provide one. Confronted by a question that most will not have pondered before it is posed in a survey, ordinary people sample from some of the “competing considerations” they carry with them and that can be easily recalled.

The question then becomes “what considerations?” and then “how many?” Perhaps people think about policy domains separately, bringing unique considerations to each. What the government ought to do about environmental issues, in this view, might summon the respondent’s recent nature walk and whether or not it produced pleasure. Much of the literature on policy studies is based on such a presumption, normally untested, that people respond to questions about policy  $x$  in terms of considerations specific to the  $x$  domain.

The alternative view is that a very small number of general considerations dictates responses to all issues. Since politicians of the left and right manage to bring remarkable unity to views on diverse issues—most of the time unity based on considerations of left and right—perhaps citizens too bring general considerations in response to diverse issues, thinking always about how much they want government to do in response to what appear to be diverse policy queries, most of which call for government action or inaction. If that were the case, then response to policy queries would be dictated in large part by quite general considerations of whether government is a desirable or undesirable solver of social problems.

Such a general consideration would be a latent variable underlying the observed responses to survey questions. It—call it policy sentiment—would be a general response to government, not domain specific. If such a thing exists,

an empirical question, then we ought to be able to extract it from the survey responses it produces. For the technically inclined, this general consideration would be a principal component—or something analogous—to marginal responses to survey questions over time. It would explain why the more left response is dominant at some times and then gives way to a movement to the right at others.

A lot of intellectual energy is wasted in my view on the question of how many considerations underly mass response to policy issues. Theory, after all, tells us that either view—domain specific or generic—is plausible, might be the way the world works. The debate is wasted energy because the “should” does not matter. This is an empirical issue. If one, or a very small number of, dimensions explain apparently diverse attitudes, then the parsimony principle dictates that we should not complicate life by requiring diverse causation where simple causation succeeds.

For the American case the evidence is decisive. After policy sentiment along left vs. right lines has explained what it can, then a second dimension emerges which is rarely very strong and rarely interpretable.<sup>1</sup> For the UK, the issue has yet to be confronted. British politics might be high dimensional—that is require many specific attitudes to explain what emerges—or low dimensional, aligned like the party system on considerations of left vs. right, Labour vs. Tory. We will let the data speak.

**Organization** The paper consists first in Section 1 of development of a concept of policy sentiment as a latent explanation for observed specific policy preferences. Then in Section 2 I get into the business of estimating the concept by extracting a latent variable from policy preference marginals. In Section 3 I undertake a simple test of Wlezien’s thermostatic model of opinion. I look at the US and UK side by side in Section 4. And finally, Section 5 deals with predicting aggregate election outcomes.

---

<sup>1</sup>I argue in (Stimson 2004) that this second dimension is not interpretable because it is the residue of once orthogonal issues that are in the process of aligning with the left-right dimension of politics, but are not yet fully aligned.

## 2 Developing A Policy Sentiment Measure

The minimum requisite for extracting a measure of left vs. right policy sentiment from national surveys is either (1) a small number of items, each of which is asked over and over again over a long span of time, or (2) a very large number of items from which we can piece together patterns of longitudinal covariance over time. The first requisite is not met in any nation to my knowledge. Survey research, and particularly commercial survey research, has heavily focused on the issue of the moment and not on enduring controversies which could be tapped over and over again.

For the UK combing the British Election Studies, the British Social Attitudes series, and the less regular policy queries of the commercial survey houses begins to satisfy the need for a large number of items. And some of those items are posed with good frequency and continuity. Thus extracting a latent dimension is possible.

What is needed—and what I have found some of—are queries that are repeated in multiple years (at least two), ask for pure policy preferences, and do not have endogeneity complications associated with the names of politicians or parties. Thus a question of what the Blair Government should do—or how well it is doing what it chooses to do—is hopelessly contaminated by attitudes toward Blair and Labour. We simply can't know preference when it is entangled with attitude toward person or party. The designers of surveys—or perhaps their customers—seem most fond of questions buried in layers of endogenous causality.

What I have are some 114 series of policy preference questions—844 individual administrations—from which I wish to estimate the latent underlying considerations for the period 1961–2005.<sup>2</sup> That is a little less than 20 items per year, which would be ample were they evenly spread across time. But they are not. The British Election Studies, which provide the largest batch, are of course fielded only in election years. They give good long-term cover-

---

<sup>2</sup>This represents a substantial improvement—essentially a doubling—over the data used in the 2005 EPOP paper. The consequence is a very substantially improved measure of policy sentiment in this paper, one that bears limited resemblance to the earlier effort. I am grateful to John Bartle, Katarina Thomson, David Sanders, and Harold Clarke for helping me locate a treasure trove of data.

age, but cannot help much with year to year variation. Some of the commercial surveys have opposite vices and virtues. They might give, for example, annual readings for a decade or so, but not before or after. Combining the two can profit from the best of both.

The challenge is to extract one or more latent dimensions of attitude which best account for the observed marginal (aka “top line”) responses over time. The method, akin to (iterative) principal components analysis, assumes that input items are unequally valid indicators of the underlying dimension. The statistical problem is to solve for validity estimates, the square roots of which are loadings—correlations between items and the latent dimension.<sup>3</sup> The result is a set of loadings, from which we infer the meaning of the latent dimension and an empirical estimate of the hypothesized dimension itself. I begin with the loadings.

**What Does It Measure?** We learn what an extracted dimension means from the pattern of input variables which are most highly associated with it. The raw data are too messy to be presented. Some 114 variables were used in the analysis, some of which were present for only two or three years, many with trivial or even negative correlations with the derived scale. What I present below in Table 1 is a small subset that combine contributing to relatively many years (five or more) with high positive loadings. They best tell what a dimension means.

What I see over and over in Table 1 is the dualism which defines the regular controversies between left and right, whether or not government should attempt to plan the economy, whether or not to spend more on a variety of government services, whether or not to permit abortions, and whether to extend opportunity to women. Less central but still important are business vs. labor, rich vs. poor, nationalization vs. privatization.

It seems safe to conclude that what lies beneath these survey items is the left-right dimension that defines the roles and positions of Tories and Labour. That is what would be expected and, indeed, anything else would be disconcerting.

---

<sup>3</sup>See (Stimson 1991) for more detail on the estimation method.

Table 1: Selected Items and Loadings for the First Dimension Solution: Items Are Ranked by Loadings

Source	Variable	Brief Text	Years	Loading
Gallup	EDUC	spend on education	20	0.92
BSA	ABORT2	abortion allowed couple do not want child	9	0.92
BES	TOLERANT	People more tolerant of unconventional lives	5	0.90
Gallup	PENSIONS	spend on pensions	20	0.88
BSA	TAXSPEND	reduce taxes and spend less on health, etc.	20	0.86
ICM	SOCPLAN	socialist planning solve economic problems	6	0.85
BSA	SOHELP	people get social security don't deserve help	13	0.82
BSA	ABORT1	abortion allowed woman does not want child	9	0.81
BSA	UNEMPJOB	unemployed people find job if they wanted	13	0.80
BSA	STATEOWN	more or less state ownership of industry	7	0.80
Gallup	NHS	spend on NHS	20	0.80
ICM	ENVCAR	higher priority to environment penalising cars	5	0.80
BES	EQOPP2FR	give equal opportunities to women	5	0.80
BSA	ABORT4	abortion allowed cannot afford children	9	0.78
BSA	WELFEET	welfare weren't generous stand on own feet	13	0.77

This first dimension of policy sentiment accounts for about 40% of the variance in the 114 items, and is clearly the dominant dimension of policy controversy.<sup>4</sup>

A second, orthogonal, dimension may also be estimated. The result, accounting for an additional 21% of the item variance appears to be an out of phase variation on the first dimension. Its strongest loading issues are attitudes toward big business, socialist planning, defense spending, state ownership, trade unions, welfare recipients, and closeness to Europe. Orthogonal second dimensions in longitudinal analyses rarely present straightforward interpretations. If I were to hazard a guess what this one means—with handsome probability of being wrong—it would be the this issue set is the old capitalist vs. socialist issue divide between Conservatives and Labour. The evidence is both the pattern of item loadings and also that the two dimensions appear

<sup>4</sup>This is slightly better than a comparable exercise in the U.S., where the first (left vs. right) dimension accounts for about 38% of all longitudinal item variance.

to diverge fairly sharply as Tony Blair comes on the scene in 1997 with a new image of Labour. (See below Figure 2.)

The estimated first dimension is presented in Figure 1. All individual items are scored by the relative percent taking the left position. The scale then acquires the means and standard deviations of the individual items, weighted by validity. Thus the numbers show a persistent, if slight, preference for left over right alternatives in the underlying data.

Longitudinally, the series in Figure 1 points to left sentiment on policy in the 1960s that gradually gave way to more neutral views in the 1970s, a strong rebound in left sentiment and during the years of Conservative Government and a steady movement back toward conservatism during the years of the Blair Labour Government.




Figure 1: The First Dimension of Policy Sentiment in the United Kingdom: 1961 to 2005

Policy sentiment is evidently cyclical, at least within the limits of what can be learned from 45 years of data. That arises from its relativistic and thermostatic nature. If it reflects what the public wants from government *relative*

to *what government is currently doing*, then alternation in party control of government and the alternation of policy drift that entails should produce a public sentiment that cycles as the parties do. Empirical experience with the US counterpart is consistent with that cyclical alternation thesis. Since parliamentary government gives parties the ability to enact their policy views that the US mixed control system often does not, there is even more reason to expect cycles in the UK.

**The Second Dimension** Figure 2 superimposes the estimated second dimension on the first. Dimensionality in longitudinal analyses captures phase relationships. For what we consider “orthogonal” relationships in cross sectional work, in time series we have out-of-phase relationships. In either case it requires additional dimensions to account for all systematic variation. In the time series case, making sense of orthogonal dimensions is rarely straightforward.

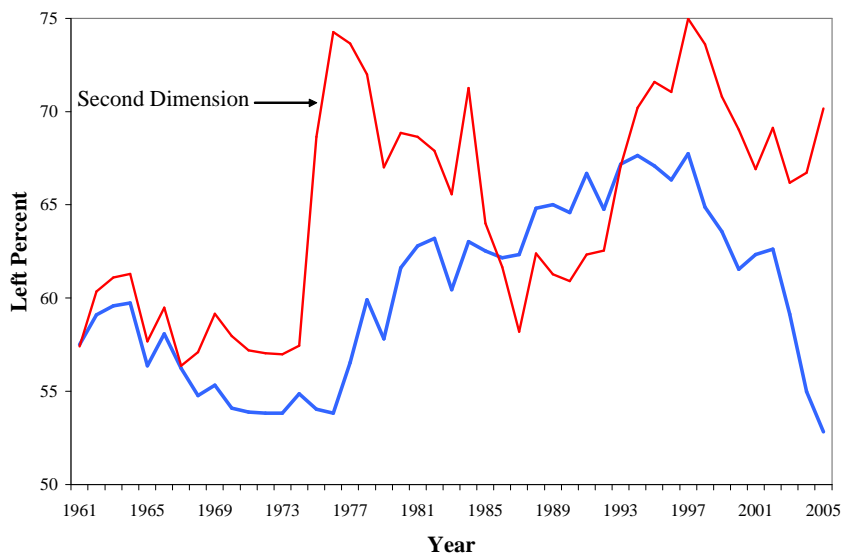


Figure 2: Two Dimensions of Policy Sentiment in the United Kingdom

Here we can say that the second dimension captures something systematic that does not move in phase with the first. Saying *what* is a much harder

matter.

### 3 The Thermostatic Theory of Policy Sentiment

Like the run-up of left sentiment that appears associated with the Thatcher-Major Governments, since Chris Wlezien's (1995) development of the thermostatic conception of public preferences we (at least in the US) have come to expect public opinion movements counter to the direction of the government in office. The theory, as applied to the UK, is that a relatively moderate central tendency in public opinion will always find the Tories too far to the right and Labour too far to the left. That being the case, the aggregate public's natural impulse is always to ask for less of what it is getting from the current government. So Thatcherism produces *relative* movement to the left while Labour tenure produces *relative* movement to the right.

Since we have in hand a measure of left and right, it should be a simple matter to observe whether it moves as the theory predicts. It is a simple matter to observe the public sentiment of the year when a party first wins a governing majority, the sentiment of the year when it is finally turned out by the opposition, and see whether movement over that span is toward or away from the party's position.

With one minor exception the predicted pattern emerges in Figure 3. All periods of Conservative government produce movements toward the left over the span of time that the government serves. All but one periods of Labour government produce countervailing movements to the right. The exception is the Labour government of 1974–1979 which left office with a modest gain in left sentiment.

We see a summary picture in Figure 4 which nets out the three Labour and three Conservative governments to show the total effect. The net pattern shows that the one small exception does not break up the general pattern; each party's span in government produces support for its opposition's ideological position. Twenty-seven years of Conservative government produce



Figure 3: Movements in Left Policy Sentiment During Periods of Party Control

a net shift to the left of about 13 points. And nineteen years of Labour produce an 18 point shift toward the right. Policy sentiment in the UK is thermostatic.

## 4 The US and the UK: Parallel Universes?

In (Erikson, MacKuen & Stimson 2002) my colleagues and I demonstrated that policy sentiment in the US was also in part a function of current economic experience. We showed in particular that high unemployment regularly adds to left sentiment while high inflation has the reverse effect. This fits nicely with the class structure of politics and as well with the typical policy dispositions of the two American parties.

If such a generic response is true more generally and if the now globalized economy is producing relatively similar longitudinal patterns of employment

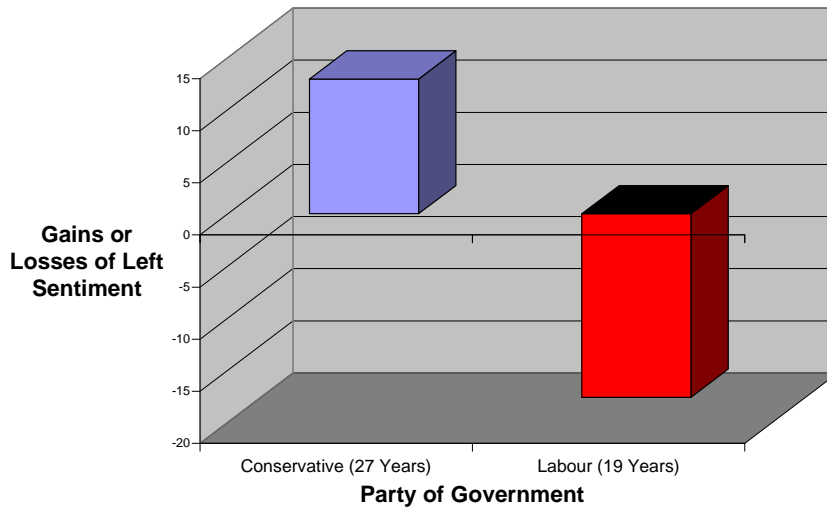


Figure 4: Net Movements in Left Policy Sentiment During Periods of Party Control

and inflation across nation-states, it follows that there should be observable parallelism of policy sentiment. Far from complete, because other factors such as party electoral success and failure will not be parallel, there should nonetheless be more than random correlation between the two tracks of policy sentiment.

Figure 5 is unpersuasive evidence for parallelism. Although much of the discrepancy between the two nations is due to party control differences—most dramatically in the recent period of conservative government in the US while Labour rules Great Britain—the net association ( $r = .09$ ) is nonsignificant. Perhaps parallelism could be teased out by a fuller specification, but perhaps too it takes imagination.




Figure 5: Policy Sentiment in the United States and the United Kingdom: 1961 to 2005

## 5 Predicting Elections

So we have a measure of public sentiment over time. Central to democratic theory is the idea that election outcomes should reflect shifts in public sentiment. A UK electorate that alters its preference in the direction of greater public spending, for example, should be more inclined to vote for the party that consistently favors greater public spending. And what should happen in a theoretical sense leads to the empirical expectation that elections should be predictable from knowledge of public sentiment.

**A Warning:** That much is straightforward. What isn't straightforward is that the task is multivariate and the number of cases—11 after losing one to lagging—isn't large enough to permit a sensitive multivariate model. The consequence of this painful absence of degrees of freedom is that estimated models are highly specification dependent. And that will be true of what I am presenting. I am going to present a model in which the key preference

variable “works” as we conventionally understand that, correctly signed and statistically significant. But it needs to be said that there are others I have tried where it does not. The key predictor variable is at least always correctly signed. It is not always significant.<sup>5</sup>

**Baseline Model:** I begin by estimating a baseline model of the Labour share in UK elections. It features one performance variable, change in Gross Domestic Product under the current government (reflected to predict Labour rather than incumbent share), the Labour share in the previous election, and a dummy for incumbency (Labour = 1). The baseline model is seen in the second column of Table 2.

Then we add the policy preference variable in the third column, asking whether it can add to predictability. It does. With a coefficient of .72 and standard error of .30, it makes the expected positive contribution—left sentiment, that is, benefiting Labour and right sentiment benefiting the Conservatives. And this is a large effect. With a maximum swing of about 15 points (from 1997 high to 2005 low) in policy sentiment, the .72 coefficient translates to an almost 11 point maximum effect on election outcomes.

As is usual with small degrees of freedom estimation, the model is less than wholly satisfactory. Most objectionable is an economic performance indicator, change in GDP under the incumbent government, which is wrong-signed. The estimated negative coefficient tells us that Labour benefits from bad performance managing the economy when it is in power or good performance under the Tories. This seems to reflect real negative covariance rather than some exotic specification wrinkle. GDP, well behaved in the American context, evidently is not in the UK.

Predicted and actual Labour shares are plotted in Figure 6 where we can see that the model underpredicts the Labour showing in both the first (1966) and last (2005) elections in the series, but is pretty much on target in between.

---

<sup>5</sup>For example, a bivariate dynamic specification produces  $\beta = 0.51$ ,  $\sigma_\beta = 0.39$ , and  $t = 1.31$  for policy sentiment predicting Labour share in election outcomes. But an attempt to add policy sentiment to the forecasting model of Bélanger, Lewis-Beck, and Nadeau (2005) does not work at all.

Table 2: Predicting the Labour Share in UK Elections: 1966 to 2005

Variables	Baseline Model	Preference Added
Previous Labour Share	1.36* (0.63)	1.65* (0.50)
Average Quarterly GDP Growth	-15.33 <sup>w</sup> (7.83)	-18.56 <sup>w</sup> (6.18)
Labour Incumbent Dummy	2.31 (5.60)	5.55 (4.52)
Public Preference		0.72* (0.30)
Intercept	-18.40 (29.64)	-77.01 (33.36)
R <sup>2</sup>	.51	.75
Adjusted R <sup>2</sup>	.31	.59

Standard Errors in parentheses

\* p < .05

w wrong signed coefficient

**A Two Stage Conception** An alternate way to think about modelling the election outcome is go ahead and use the endogenous party identification concept and then to attempt to explain its variation over time. I entertain the possibility of a causal sequence such as :

Policy Preferences  $\longrightarrow$  Party Identification  $\longrightarrow$  Election Outcome

which dictates a two-stage analysis. First we ask “does party identification move election outcomes?” and then “do preferences explain some of the longitudinal variation in party identification?”




Figure 6: The Labour Party Share, Actual and Predicted: 1966 to 2005

For the former question I set up the ultimately simple specification of just party identification  $\rightarrow$  election outcome. The identification measure, taken from the BES, is the Labour share of Labour plus Conservative identifiers. The bivariate setup is shown in Table 3. It is no surprise to discover that how voters feel about the parties affects the party for which they choose to vote. The relationship is quite strong with a good fit (Adjusted  $R^2 = .76$ ) suggesting that a bivariate specification isn't missing a lot.<sup>6</sup>

Then we ask whether policy preferences flow into identification: are voters in the aggregate more apt to think of themselves as Labour supporters when their aggregate policy views come down on Labour's side. The answer, from the regression of Table 4, is yes. The specification assumes that partisanship has a memory—it is either unit root or fractionally integrated in the U.S. case—and that changes in the current election period are due to changes in policy preference.

<sup>6</sup>Entertaining a dynamic specification, letting Labour share be predicted by previous Labour share and partisanship, produces an even stronger effect for partisanship,  $\beta = 0.85$  with  $\sigma_\beta = 0.17$ , but at the expense of a negative and nonsignificant dynamic term.

Table 3: Predicting the Labour Share in UK Elections From Party identification Alone: 1964 to 2005

Independent Variable	$\beta$	Standard Error
Labour Share of Identifiers	0.71	0.12
Intercept	11.51	6.46
Adjusted R <sup>2</sup>	.76	
N	12	

Notes: The Labour Share is  $(\%Labour)/(\%Labour + \%Conservative)$

The result is consistent with a unit root interpretation— $\beta > 1.0$  for previous partisanship—but an N of 12 doesn’t permit a unit root (statistical) inference. But the key message of the table is that current partisan views—a step away from voting intentions—appear significantly to respond to current policy preferences.

Table 4: Predicting Labour Identification from Policy preferences: 1964 to 2005

Independent Variable	$\beta$	Standard Error
Previous Labour Identification	1.08*	0.31
Policy Preferences	0.85*	0.42
Intercept	-53.45	37.12
R <sup>2</sup>	.61	
Adjusted R <sup>2</sup>	.51	
N	12	

Thus we have suggestive evidence that policy sentiment matters in the way democratic theory says that it should. But we can’t learn very much from 12 cases.

## 6 Concluding Notes

What then do we know about policy sentiment in the end? We know from the evidence internal to the solution, e.g., an estimated Eigenvalue, that a

coherent first dimension emerges. We know from the pattern of its item loadings that it appears to measure the standard left vs. right issue domain. We have strong evidence about its causation, that it moves contrary to the direction of policy-making at the time—what we expected to see. We know that it can be employed to predict aggregate election outcomes, although the evidence from a multivariate model of 12 cases can never be powerful. And we know that it predicts aggregate partisanship.

**Where To Go From Here?** Keeping in mind Lord Keynes’ aphorism, we don’t wish to wait decades for a richer set of elections data. So discovering more about this concept and measure must move to where there are more and better data. This analysis wastes some of the power we should have, using only 12 cases from a 45 year time series. There is progress to me made working on a finer time scale.

Political theory has a prime role for public opinion. And for longitudinal public opinion this policy sentiment concept should be a leading candidate. Many questions are yet to be asked or answered, most probably having in common that the data necessary for the answer might—but do not yet—exist. Among them are public policy, both as effect of opinion and as cause of opinion feedback. The role of partisan movements on a finer time scale, annual or even quarterly, ought to reflect changes in policy sentiment (although this fairly obvious link proved vexing for my colleagues and myself in (Erikson, MacKuen & Stimson 2002)). We will want to know how this actual sentiment is related to perceived sentiment, citizens’ self identification on a left–right dimension. And movements in policy sentiment require more and better explanation. The political economy link seems promising, with much remaining to be done.

## 7 Appendix: Data

The first column (after year) is the measure employed in this paper. The second and third columns are from the two dimensional solution of Figure 2. These data differ slightly from those reported in the paper because of slight assumption shifts necessary for estimating two dimensions. The two first dimensions *appear* quite different, due to a metric shift, but correlate at .98.

year	Dim. 1	Dim. 1	Dim. 2
1961	57.49	62.21	57.40
1962	59.10	64.08	60.36
1963	59.58	64.55	61.11
1964	59.72	64.67	61.30
1965	56.37	60.50	57.68
1966	58.07	62.84	59.48
1967	56.23	60.50	56.36
1968	54.77	58.65	57.10
1969	55.32	59.00	59.16
1970	54.09	56.89	57.97
1971	53.88	56.73	57.20
1972	53.83	56.68	57.03
1973	53.82	56.67	56.99
1974	54.86	58.33	57.45
1975	54.03	57.52	68.64
1976	53.83	57.39	74.27
1977	56.53	61.03	73.66
1978	59.91	65.39	72.00
1979	57.81	62.21	67.00
1980	61.63	67.50	68.86
1981	62.80	68.96	68.64
1982	63.19	69.42	67.90
1983	60.44	65.37	65.57
1984	63.03	68.71	71.28
1985	62.53	68.41	64.00
1986	62.16	68.47	61.67

1987	62.33	68.27	58.19
1988	64.82	71.62	62.40
1989	65.01	71.46	61.27
1990	64.59	71.67	60.90
1991	66.69	74.44	62.33
1992	64.76	72.48	62.54
1993	67.18	74.00	67.09
1994	67.65	73.51	70.20
1995	67.09	72.35	71.59
1996	66.34	71.04	71.06
1997	67.74	71.80	74.98
1998	64.88	68.25	73.62
1999	63.56	66.54	70.81
2000	61.55	64.13	69.03
2001	62.33	65.68	66.91
2002	62.61	65.63	69.14
2003	59.13	60.94	66.19
2004	54.97	56.75	66.72
2005	52.82	54.21	70.17

## References

- Bélanger, Éric, Michael S. Lewis-Beck & Richard Nadeau. 2005. "A Political Economy Forecast for the 2005 British General Election." *British Journal of Politics and International Relations* 7:191–198.
- Erikson, Robert S., Michael B. MacKuen & James A. Stimson. 2002. *The Macro Polity*. New York: Cambridge University Press.
- Soroka, Stuart N. & Christopher Wlezien. 2005. "Opinion-Policy Dynamics: Public Preferences and Public Expenditure in the United Kingdom." *British Journal of Political Science* .
- Soroka, Stuart N. & Elvin T. Lim. 2003. "Issue definition and the opinion-policy link: public preferences and health care spending in the US and UK." *The British Journal of Politics and International Relations* 5(4):576–593.
- Stimson, James A. 1991. *Public Opinion in America: Moods, Cycles, and Swings*. Boulder, CO: Westview Press, 1991.
- Stimson, James A. 2004. *Tides of Consent: How Public Opinion Shapes American Politics*. New York and London: Cambridge University Press.
- Wlezien, Christopher. 1995. "The Public as Thermostat: Dynamics of Preferences for Spending." *American Journal of Political Science* 39(4):981–1000.
- Wlezien, Christopher & Stuart N. Soroka. 2003. "Measures and Models of Budgetary Policy." *Policy Studies Journal* 31(2):273–286.
- Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.
- Zaller, John & Stanley Feldman. 1992. "A Simple Theory of the Survey Response: Answering Questions and Revealing Preferences." *American Journal of Political Science* 36:579–616.