

The information content of elections and varieties of the partisan political business cycle

Cameron A. Shelton

Received: 16 July 2009 / Accepted: 29 July 2010 / Published online: 19 August 2010
© Springer Science+Business Media, LLC 2010

Abstract This event study uses economic forecasts and opinion polls to measure the response of expectations to election surprise. Use of forecast data complements older work on partisan cycles by allowing a tighter link between election and response thereby mitigating concerns of endogeneity and omitted variables. I find that forecasters respond swiftly and significantly to election surprise. I further argue that the response ought to vary across countries with different institutional foundations. In support, I find that there exist three distinct patterns in forecasters' responses to partisan surprise corresponding to Hall and Soskice's three varieties of capitalism.

Keywords Political business cycle · Varieties of capitalism · Forecast data · Opinion polls

JEL Classification E32 · E63 · P16 · P51

1 Introduction

1.1 An example of politically induced macroeconomic volatility

The National Election Study conducted over the two months prior to the US Presidential election of November 7th, 2000 showed Democratic candidate Al Gore with a nine point lead over Republican candidate George W. Bush among intended voters. However, Election Day proved that the contest was essentially a dead heat and, due to ballot recounts in Florida, did not produce a victor. For 36 days the identity of the victor was in doubt until, late in the night of December 12th, the US Supreme Court ruled to halt the vote-recount resulting in a victory for Bush on the morning of December 13th.

During this period from early September 2000 through late December 2000 as the election slipped from Gore to Bush, projections of US GDP growth in 2001 from the Consensus Economics panel of forecasters show a major shift in expectations. As Election Day revealed

C.A. Shelton (✉)
The Robert Day School of Economics and Finance, Claremont McKenna College, Claremont, USA
e-mail: cshelton@cmc.edu

a virtual tie rather than the expected Gore victory, the average forecast of 2001 calendar year GDP growth dropped one half of one percentage point. Upon resolution of the standoff in Bush's favor by the Supreme Court, the average of GDP growth forecasts dropped another half-point. Expected bond yields also plunged, witnessing a one-time drop of six tenths of a point in the month following the Supreme Court ruling (see Fig. 1). This episode points to an intriguing connection between the information revealed during elections and forecasts of the subsequent performance of the economy.

In this paper, I examine the response of economic forecasts to changes in expected political leadership in industrialized countries. There are three main results. First, I find that forecasts respond strongly to partisan surprise, indicating that forecasters do perceive important differences between the macroeconomic policies of the left and right and expect those differences to translate into different macroeconomic outcomes. I also proceed to demonstrate that these forecasters are capable of processing political information reasonably accurately. Forecasts are ideal for measuring these partisan effects because they are jump variables and can respond to information much more quickly than macroeconomic state variables, allowing for a much tighter link between cause and effect. Use of forecast data complements older work on politically-induced macroeconomic volatility which measured responses in macroeconomic variables several years after the election and thus raised questions of endogeneity, omitted variables, and the power of the tests.

Second, I argue that the response of forecasts to election surprise ought to depend on the institutional organization of the economy in question. This is because (a) the political cleavage over economic policy is a function of the underlying organization of the economy in question and (b) the ability of a government to enact its platform is a function of the institutions of governance which are at least correlated and likely co-evolved with the economic institutions. As a result, the economic effects of political surprise should vary considerably across countries. I find strong evidence that this is the case. The perceived difference between the policies of the left and right in liberal market economies is very different than the difference between left and right in coordinated market economies which are again different from the perceived difference in Mediterranean market economies. The novel implication is that liberal market economies, coordinated market economies, and Mediterranean economies ought to experience rather different partisan political business cycles.

Finally, I argue that, in its effects on macroeconomic variables, policy is more than simply a choice of the degree of aggregate demand stimulus. In industrialized democracies, elected officials control fiscal and regulatory directly as well as possibly influencing monetary policy indirectly. Fiscal and regulatory policy are often multi-dimensional. As a result, political effects should be more complex than simple oscillation along an unemployment-inflation tradeoff. The results support this notion of a multi-dimensional policy-space. Forecasters clearly believe the difference between parties' policies to be more complex than a simple choice of position on a short-run Phillips curve. In some countries, the perceived difference between parties is a classic tradeoff between output/unemployment and inflation. But in other countries, one party is expected to increase output without any effect on unemployment or inflation. And in a third group of countries, one party is expected simultaneously to increase output, decrease unemployment, and decrease inflation compared to its (economically less competent) political opposition.

In the course of the study, I introduce two data sets: a commercially available data set on expectations which is new to this literature and a set of pre-electoral opinion polls which I have assembled for this study in order to quantify the change in expected political leadership resulting from an election.

The paper is organized as follows. The next section briefly reviews the theory and basic data on partisan political business cycles and argues that we ought to expect heterogeneous




Fig. 1 Forecasts shift strongly following the Surprise US election results of 2000

cycles. Section 3 outlines the measurement strategy. Section 4 introduces the new data sets for forecasts and opinion polls. Sections 5 and 6 detail the construction of the necessary variables from these data and the resulting specification and estimation technique. Section 7 presents and discusses the results. Section 8 investigates the rationality of the forecasts. Section 9 concludes.

2 Heterogeneous partisan cycles

As a choice between alternate policy platforms, an election reveals information about future economic policy. Because the election resolves the uncertainty over which policy will be chosen by the electorate, an election can deliver a surprise change in policy even if parties are credibly committed to policies whose effects are correctly anticipated (Alesina 1989). Thus elections are effectively shocks to future economic policy.

The partisan political business cycle is a collection of facts concerning the relation between election results and post-electoral economic performance. In his seminal article and subsequent book, Hibbs (1977, 1987) presents evidence of partisan effects on output, unemployment, and inflation in a dozen industrialized democracies. Left-wing administrations preside over periods of higher output growth, lower unemployment, and higher inflation compared to their right-wing counterparts. Further work has shown partisan effects on output growth rates and unemployment rates tend to be temporary, disappearing within a year or two of the election, while effects on inflation may be more permanent (Alesina 1988; Alesina and Sachs 1988; Alesina and Rosenthal 1995; Alesina et al. 1997; Grier 2008). There is also evidence that the magnitude of these partisan effects is related to the degree to which the election outcome was a surprise (Alesina et al. 1997; Drazen 2000; Grier 2008).

There exist two main theories to explain these facts: the traditional partisan theory (Hibbs 1977) and the rational partisan theory (Alesina 1987). The traditional partisan theory (PT) relies on adaptive inflation expectations to generate a relatively stable short-run Phillips curve. The rational partisan theory (RPT) is based on an expectations-augmented Phillips curve. Inflation expectations are assumed to be forward-looking and rational and thus only unexpected aggregate demand shocks affect output. An election serves as such a shock. On the eve of an election, it is uncertain which party will be in power next year. As a result, rational inflation expectations are an average of the preferred inflation policies of the parties weighted by the probabilities of each party being elected. When the election takes place, there is a sudden resolution of this uncertainty as a winner is produced and the party (and associated economic policy) in power next year is identified.

Because elected officials directly and indirectly control a wide range of fiscal, regulatory, and monetary policy instruments, policy is more than a choice of how strongly to stimulate aggregate demand. The purview of the fiscal policy authority includes spending and subsidies but also includes regulations concerning the source and magnitude of pension and health-care contributions, hiring and firing laws, replacement rates for the wages of the unemployed, the mix of personal and corporate taxes, the degree of progressivity in the income tax code, and a host of other policies with effects not captured by the classic tradeoff between unemployment and inflation. The issue is not simply the degree of fiscal or monetary stimulus but the distribution of net burdens of taxes and transfers, the direction of investment in public goods, and the design of market regulations.

For example, these models of monetary policy typically assume that output and employment move together and refer to them interchangeably. This may be broadly true in that output and employment may move together in response to aggregate demand stimulus. But

distributive and regulatory policy can decouple output and employment. To the extent that parties differ in these areas, politically induced macroeconomic volatility will not simply be an oscillation along the short-run Phillips curve.

There is also reason to believe that partisan differences, and hence politically induced macroeconomic volatility, vary by country. The left-right dichotomy originates from the seating arrangement in the first elected French Legislative Assembly of 1791 and was based on support for or opposition to the *ancient regime*, the dominant political issue of that particular time and place. From then on, the left-right terminology has been exported to a variety of polities and the salient issues in these polities have evolved. As a result, left vs. right has come to cover a host of different ideological differences: big government vs. limited government, intervention vs. *laissez faire*, redistribution vs. free market, equality vs. liberty, fair outcomes vs. fair processes, secularism vs. religion, collectivism vs. individualism. Broad application of the labels left and right has created the illusion of standardized political debate where, in fact, political cleavages are a function of the environment and thus vibrantly heterogeneous.

In short, the space of fiscal and regulatory policies is multi-dimensional and the primary political axis differentiating the dominant parties is endogenous. Citizens' preferences are distributed in a multi-dimensional policy-space and parties choose positions to represent citizens and win votes. Ideologically, one would expect the major axis of partisan differentiation to be the principal component of the underlying variation of preferences in the polity. As these preferences shift, so does the partisan axis. Thus salient partisan differences must be a function of the local environment.

As a matter of practical policy, challenges for economic policy and the array of possible solutions are naturally a function of the institutional environment within which alternate policies must operate. For example, the persistence of high unemployment in European countries is often attributed to inflexible labor laws governing the hiring and firing of workers suggesting that relaxing restrictions might lead to lower unemployment. As changing these laws has consequences for distribution and social justice, citizens, and thus political parties, differ in their support for the proposal. Generally, right parties have supported this solution to the unemployment problem while left parties have proposed a reduction in the work-week to spread the "lump of labor" across a greater number of people. In this case, the existing labor-market institution frames the issue and suggests a policy prescription with a particular set of distributional consequences. As a result, the existing labor market institution helps to define the partisan cleavage in these countries. Notice that in the United States, no such debate exists. The United States possesses a different set of labor market institutions and thus faces a different set of policy challenges admitting a different set of possible solutions leading to a different political cleavage. To the extent that policies must be incremental changes to the organization of the economy, existing institutions define the comprehensible policy alternatives and feasible macroeconomic tradeoffs and thus define the partisan differences over economic policy. In sum, we should expect that countries with different institutional environments have different left-right divides and thus different partisan political business cycles.

How then do we characterize the institutional environment of a country? Williamson (1985) and North (1990) have championed the idea that institutions—both firms and those governing intra-firm relationships—are a path-dependent social technology that evolves to reduce transaction costs of complex economic interactions. As a result of the path-dependent nature of institutional evolution, multiple systems may coexist. The literature on 'varieties of capitalism' emphasizes the systematic and persistent differences in the institutional organization of national economies. In particular, Hall and Soskice (2001) argue that complemen-

tarities between institutions of corporate governance, corporate finance, education and vocational training, wage bargaining, employment protection, technological standard setting, and technological transfer lead nations to cluster into three varieties of capitalism. In liberal market economies (LMEs), economic activity tends to be organized around arms-length transactions in competitive, decentralized markets with formal contracting. In coordinated market economies (CMEs), economic activity tends to take place via repeated collaborative relationships with informal contracting and network monitoring based on the exchange of private information. Mediterranean market economies (MMEs) represent a third cluster with “capacities for non-market coordination in the sphere of corporate finance and more liberal arrangements in the sphere of labor relations.” (Hall and Soskice 2001, p. 21)¹ Given the importance of institutions in framing policy issues and suggesting solutions with particular distributional and ideological consequences, one would expect economies with such broadly different institutional bases to be characterized by different political cleavages and, as a result, by different partisan political business cycles. Given the broad nature of these institutional families, containing many types of institutions, it is unfortunately not possible to develop predictions concerning which institutional family will be associated with what type of partisan political business cycle. This study is intended as an atheoretical first pass designed simply to demonstrate the existence of country-level heterogeneity in political cycles and tie it to the institutional setting via these broad families.

As I will show, there are three broad patterns in the response of forecasters to partisan surprise corresponding to the three varieties of capitalism defined by Hall and Soskice. In one set of countries (France, Italy, Spain), the left is expected to deliver higher output growth and lower unemployment than the right with higher short-term interest rates, but at the cost of higher inflation. This is the traditional partisan political business cycle. In another set of countries (United States, United Kingdom, Canada), the left is expected to deliver higher output growth than the right with higher short and long-term interest rates but no effect on unemployment and inflation—a relative boom for free. And in the third set of countries (Germany, Japan, the Netherlands, Norway, Sweden), the left is associated with lower output growth and lower interest rates—a relative bust—and with higher unemployment and higher inflation. The left is simply seen as less competent in general or as trading macroeconomic performance for other goals. These results indicate that politically induced macro-economic volatility is neither homogeneous across countries nor simply the result of oscillations along a SR Phillips curve.

3 Estimating partisan cycles

Alesina et al. (1997) attempt to differentiate between the traditional and the rational partisan mechanisms by looking at the duration of the partisan effects. They claim that while RPT predicts temporary effects, the traditional theory results in permanent partisan differences over the entire term of office. In a panel of 18 OECD democracies covering 1960–1993, they find that partisan differences in growth and unemployment are confined to the first year or two after the election while partisan differences in inflation persist throughout the entire

¹Hall and Soskice arrive at their classification by looking for countries sharing similar institutions along many dimensions. First, they quantify several of a country’s institutions, e.g., stock market capitalization, patterns of patent applications, employment protection, union density and the level at which bargaining takes place, etc. Then they look for clustering in the multi-dimensional “institution-space”. They find three such clusters and label these LMEs, CME, and MMEs.

term of government. While this work is valuable in documenting the behavior of the partisan political business cycle, it does not satisfactorily establish the transmission mechanism.

Essentially, *both* theories predict transitory effects in output growth and unemployment, albeit with different mechanisms governing the duration of the partisan effect. Alesina, Roubini, and Cohen (ARC) indicate that partisan effects decay within two years. They contend that this is consistent with wage-contracts with an average length of one to two years. However, it is also consistent with adaptive expectations that adjust sufficiently rapidly. The parameter values are plausible either way so it is difficult to reject either theory convincingly with a measurement of the cycle's length.

Furthermore, it seems likely that both mechanisms are simultaneously at work. Large firms with independent forecasting departments to inform investment, R&D, pricing, and production decisions might accurately be characterized by forward-looking expectations. The mom-and-pop grocery store on the corner is probably more adaptive, setting prices based on observation of recent changes in input prices rather than sophisticated projections of supply and demand from recent events. Other wage and price setters probably distribute on a continuum from those who incorporate relevant public information and project future states nearly as quickly and completely as professional forecasters to those who react mainly to recent, proximate changes in prices. This is essentially the spirit of recent models built on time-contingent price-adjustment (Mankiw and Reis 2002; Woodford 2002).

Since PT and RPT are both formulated to explain observed partisan political business cycles, any test based on the cycle itself is likely to suffer from low power. But the mechanisms require very different behavior of *expectations*. Under RPT, the resolution of electoral uncertainty acts as an unexpected shock to aggregate demand: it is the information revealed on Election Day that drives the partisan political business cycle. In PT, by contrast, information revealed on Election Day goes unremarked: expectations adjust only gradually as the policy of the new administration is revealed through the performance of the economy. These two sources of expectations changes operate on very different time scales. Thus to distinguish between them, we ought to look at the behavior of expectations in the neighborhood of the election. Under RPT, an election should trigger a sharp change in expectations. Under PT it will not.

However, characterizing the relative importance of Hibbs's and Alesina's mechanisms would require a broad set of macroeconomic expectations capturing the range of sophistication from Ford Motor Company to the corner market. Unfortunately, expectations data are hard to come by. Systematic recordings of professional forecasters' expectations for a broad range of countries and variables are no more than 20 years old. Data for the rest of the spectrum of wage and price setters remain almost non-existent. A systematic exploration of the entire spectrum must therefore be the subject of future work.² In this paper I look solely at the expectations of professional forecasters: presumably the rational, forward-looking end of the spectrum. Focusing solely on the expectations of professional forecasters isolates the rational partisan mechanism. This enables an investigation of how sharply the most attentive forecasters perceive the effective economic differences between the policies of the left and right to be and how and why this perceived difference varies across countries. In Sect. 8, I establish the efficiency with which these forecasters process political information to demonstrate that these forecasts actually represent forward-looking agents in the sense captured by RPT.

²See Mankiw et al. (2003) for a comparison of popular and professional inflation forecasts in the United States. They use the Michigan Survey of Consumer Attitudes and Behaviors which includes a direct question about inflation forecasts in a survey of ordinary citizens. To my knowledge, these kinds of data are relatively unique. It is unfortunately not available for a broader set of indicators and countries.

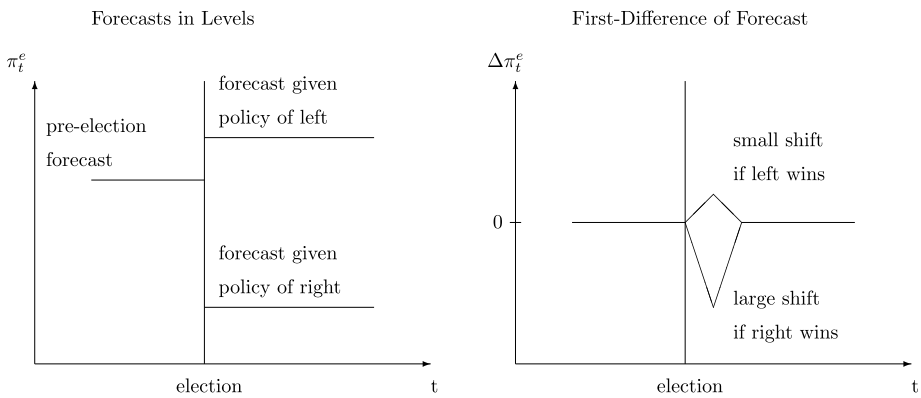


Fig. 2 The behavior of forward-looking forecasts near an election. (*At Left*) The pre-election forecast is a weighted average of the forecasts under alternative election results. The weights are given by the ex-ante probabilities of each party winning. In this example, a left-party victory is more likely. (*At Right*) Taking the first difference of the forecasts, we expect to see a discrete change at election time when the information is revealed, followed by a return to stability. The direction of the change depends on the identity of the victorious party. The magnitude of the change depends on the degree to which that victory was unexpected

Figure 2 shows how forward-looking expectations respond to the information revealed in an election and motivates the central measurement of this paper. Imagine an election in year t which will determine which party is in power (and thus which fiscal policy is implemented) in the following year, $t + 1$. Assume that the left and right-wing political alternatives espouse different fiscal policies with different consequences for a macroeconomic variable of interest such as inflation. To the extent that parties can credibly commit to policies before the election and forecasters can predict the implications of these policies, this implies that the forecast of year $t + 1$ inflation depends on which party wins the election. Of course, only one of these forecasts is actually observed; the other is counter-factual.

Meanwhile, pre-election inflation expectations are an average of the two inflation levels that would prevail under the alternate policies of the left and right weighted by the probability of each party actually winning the election. Once Election Day results are revealed, expectations change discontinuously to match the new policy. The more surprising the election result, the further the forecasts must shift. Similarly, the greater the distance between the policies of the left and right, the larger the shift. By measuring the degree to which the election was a surprise and the response of forecasts to that surprise, we can infer the perceived distance between the levels of inflation under the fiscal policies of the left and right. And by doing this for a variety of macroeconomic indicators we get a broad picture of the perceived difference between the fiscal policies of the left and right.

The central measurement of this paper is a set of simple regressions of the change in economic forecasts on the change in the political forecast during a narrow window around the election. What relationship between these two does existing theory suggest? The RPT model predicts that following political surprise in period t , unemployment temporarily diverges from the natural rate for period t . Inflation adjusts in period $t + 1$, bringing unemployment back to the natural rate in $t + 1$. Output is assumed to move counter to unemployment. Liberal (conservative) surprise delivers a one period drop (rise) in unemployment followed by an increase (decrease) in inflation.

Hence the prediction under this model depends on the horizon of the forecasts. A short forecast horizon should pick up an effect on unemployment and output but little effect on

inflation while a longer horizon would pick up little effect on unemployment and output but an effect on inflation. To the extent that liberal governments care more about output than inflation, we expect lower unemployment, higher output growth, and higher inflation following liberal surprise, with the magnitude depending on the length of the forecast horizon compared to the speed with which prices adjust.

But RPT is built on a modified Barro–Gordon (1983) framework in which the elected policy-maker directly controls inflation. Partisan political effects are more likely driven by fiscal and regulatory policy as they are under the direct purview of the elected officials. Under the assumption that liberal governments are more inclined to stimulus, the hypothesis remains that liberal government will result in higher output, lower unemployment, and higher inflation. However, many other aspects of the partisan cleavage may depend on the exact nature of the fiscal stimulus and the interplay between the fiscal and monetary policy authorities. Fiscal stimulus may be designed either to encourage household consumption or business investment. Regulatory policies may similarly affect the level of business investment. The pattern of partisan effects on components of GDP will depend on the partisan policy differences in a particular electorate. Equilibrium interest rates will clearly be a result of the degree to which monetary policy accommodates fiscal policy and hence there is no clear prediction from the partisan cleavage alone. Presumably, expert forecasters have some expectation about the likely reaction of the relevant central bank which is incorporated into their model of the economic consequences of partisan fiscal policy.

4 Two new data sets

I have obtained a panel of data recording forecasts of seven macroeconomic variables for 11 OECD countries from 1993–2004. I have also collected opinion poll data for the 30 elections in the countries and years covered by the forecast data. Comparing pre-election opinion polls—on which forecasters are assumed to base their political projections—to election results quantifies the degree to which the election result is a surprise. To my knowledge, this is the first study of election-induced macroeconomic volatility using expectations data. It is also one of a small number of studies which controls for the magnitude of election surprise.³ Matching these unique data sets, I regress the post-electoral change in economic forecasts on the election results to measure the anticipated difference between the policies of the left and the right. By looking at forecasts of several macroeconomic variables, I find that (a) forecasters respond strongly and swiftly to the information content of elections, (b) the perceived difference is more complex than an unemployment-inflation tradeoff and (c) the difference varies by country, with coordinated market economies exhibiting a markedly different policy cleavage than liberal and Mediterranean market economies.

4.1 A panel of economic forecasts

To measure economic forecasts, I employ forecast data from Consensus Economics. The data consist of monthly surveys covering 11 countries—Canada, France, Germany, Italy,

³Earlier studies using polling data include Chappell and Keech (1988), Alesina et al. (1997), Heckelman (2002), and Berleman and Markwardt (2003). In their Chap. 5, ARC use option-pricing techniques to convert pre-election polling data for the United States into election probabilities giving a more sophisticated measure of surprise. But in the rest of their book, they concentrate only on elections resulting in a change in government partisanship (left to right or vice versa). This is akin to assuming that the incumbent party is expected to win with certainty and thus elections where the incumbent wins generate zero surprise while those where the incumbent loses generate full surprise.

Table 1 Number of panelists per survey, by country

Country	Obs	Months	Mean	Std.Dev.	Min	Max
Canada	2006	139	14.4	1.6	11	19
France	2515	139	18.1	2.0	11	24
Germany	3772	139	27.1	2.2	20	32
Italy	1871	139	13.5	2.6	6	21
Japan	2701	139	19.4	2.3	12	26
Netherlands	1060	115	9.2	1.5	7	14
Norway	731	74	9.9	1.5	6	12
Spain	1417	115	12.3	2.0	7	19
Sweden	1417	115	12.3	2.2	7	17
UK	4037	138	29.3	5.1	19	39
USA	3632	139	26.1	3.0	19	33

Japan, Netherlands, Norway, Spain, Sweden, the United Kingdom, and the United States—from January 1993 to July 2004.⁴ For each country-month, the survey records individual forecasts on a number of macroeconomic variables by a number of panelists, typically commercial and investment banks, large firms, and think tanks. The exact number of panelists and the depth of the survey vary by country and month. Individual panelists are identified and can be tracked through the data. The duration of their stay in the sample varies by panelist: some panelists are in the sample consistently for years, others regularly miss a few surveys a year, and some simply make a brief one or two month appearance. Table 1 summarizes the number of panelists per country. The variables for which forecasts are formed for all countries include output growth, the growth of private consumption, the growth of business investment, consumer price inflation, 90-day interest rates, and 10-year bond yields. For the G7, forecasts are also made on the level of unemployment. Producer price inflation, the level of exports and imports, and the trade balance are also widely available while housing starts and auto sales are included for a few countries. Individual panelists do not necessarily deliver forecasts for every variable included in a country survey. Nonetheless, for the seven variables mentioned in the first two lists, which will form the core of this study, the response rate is extremely high.

Two aspects of these data are particularly exciting. First, the high frequency not only increases the number of observations but enables one to measure the high frequency responses to events that are likely to characterize expectations. A look at the summary statistics in Table 2 shows that expectations are often stable but are also capable of large changes in a single month. The $\Delta\xi$ variables are the first differences of a panelist's forecasts of variable ξ . The sample average for these first-differences is close to zero and the inter-quartile range is usually quite small. However, the standard deviation is large. This suggests that individual panelists adjust infrequently—perhaps every few months on average—but that such adjustments can be large and swift. Although many observations exhibit no adjustment, this does not necessarily imply infrequent updating of information or rule out rapid responses to

⁴The G7 are in the sample from the beginning; other countries have been added later. Norway was added to the sample in June 1998 while Spain, Netherlands, and Sweden were added in January 1995. Switzerland is part of the survey since June 1998 but has been dropped from this study due to difficulty finding Swiss polling data.

Table 2 Summary statistics of forecasts

Variable	Obs	Mean	Std. Dev.	Min	Max	Iqr
Output growth (<i>y</i>)	23817	2.47	0.84	−2.1	6.2	1.0
Consumption growth (<i>c</i>)	23706	2.31	0.91	−2.5	5.3	1.0
Investment growth (<i>i</i>)	23435	4.73	3.07	−14.7	20.0	3.4
Unemployment (<i>u</i>)	18879	7.66	2.76	0.6	13.5	6.2
Consumer Price Inflation (<i>cp</i>)	23463	2.06	1.05	−4.7	9.1	1.1
3-month rates (<i>sr</i>)	22166	4.37	1.96	0.0	13.5	2.3
10-year yield (<i>lr</i>)	22043	5.77	1.75	0.3	13.8	1.8
Δy	20996	−0.02	0.25	−3.2	3.0	0.0
Δc	20877	−0.01	0.25	−3.1	3.0	0.0
Δi	20535	−0.06	1.10	−17.3	16.1	0.1
Δu	16606	0.00	0.22	−2.0	2.0	0.0
Δcp	20521	−0.02	0.26	−4.0	4.0	0.0
Δsr	18809	−0.02	0.35	−3.2	2.7	0.2
Δlr	18930	−0.02	0.31	−2.7	3.0	0.2
SURPRISE > 0	6	0.24	0.38	0.02	1.00	
SURPRISE = 0	13	0	0	0	0	
SURPRISE < 0	9	−0.25	0.28	−0.96	−0.04	

$\Delta\xi$ refers to the one month change in forecast of variable ξ

IQR refers to the inter-quartile range: the difference between the 25th and 75th percentiles

events. High frequency data allows a more precise examination of the timing of panelists' responses to events and information.

The other exciting aspect of the data is the existence of multiple panelists per country. If each panelist is endowed with an information set and a forecasting model, then having multiple panelists is like a repeated experiment with multiple draws of information set and forecasting model. This increases the variation in the dataset and increases confidence in the generality of the results so long as the forecasts display independent variation. For each variable, I calculate the average pairwise correlation between the forecasts of two panelists located in the same country. This gives a measure of how much independent variation exists among different panelists' forecasts of the same variable. I then repeat the exercise for the first difference of each variable, which speaks to the cohesion in the sample over the short-run. The average pairwise correlation in output growth forecasts ranges between 0.41 and 0.61 which seems quite low given that the forecasters are observing the same country. This suggests that panelists are endowed with significantly different information sets and/or employ a variety of forecasting models. The average pairwise correlation of the month-to-month change in the forecast of output growth is much lower, ranging from 0.12 to 0.23. This indicates that the direction in which one panelist's forecast moves tells relatively little about the direction in which the concurrent forecast of another panelist (in the same country) will move. This is further information that panelists either focus on different sets of information or carry different interpretations of what that information means for future output growth. Thus the data indicate that panelists bring significantly different information to the sample, tend to move in different directions in the short-run, and tend to adjust every few periods rather than every month. Not all horses in the stable look alike. What I will show

is that despite this evidence of forecaster heterogeneity, elections spook the herd to move strongly in a certain direction.

For each variable, there are two forecast horizons: the current calendar year (current-year) and the following calendar year (year-ahead). This means that the forecast horizon is not actually a fixed distance from the date of the survey, but comes closer as the end of the year approaches and then leaps ahead again as the new calendar year is reached. For example, in the March 1996 survey, panelists record their forecasts for output growth for the 1996 calendar year and output growth for the 1997 calendar year. In the December 1996 survey, they record forecasts of output growth for the 1996 calendar year (which is all but over) and the 1997 calendar year. Then, in the January 1997 survey, panelists record forecasts for output growth over the 1997 calendar year and the 1998 calendar year. Clearly, as the year passes, the forecast horizons move closer to the survey date. As a result, forecasts for the current calendar year tend to become more tightly clustered toward the end of the year as the bulk of the uncertainty is resolved.⁵ This effect is also seen in the year-ahead data series, but is much weaker since the entire period with which the forecast is concerned remains in the future.

The central regression in this paper, (6), uses the first difference of a panelist's forecasts, $\Delta \xi_t^e$. However, if one were blindly to take the first difference of either the current year or the year-ahead data series, one would end up with data displaying spurious seasonality. Imagine that one is interested in using the year-ahead data to calculate a two-month measurement window for the variable ξ . Thus $\Delta \xi_t^e = \xi_{t+2}^e - \xi_t^e$. This generates a consistent data series for the months of January through October: in each case the forecast horizon is two months nearer on the far side of the measurement window than it is on the near side. For example, in April of 1992, we are taking the difference between $\xi_{June 1992}^e$, when the forecast horizon is 18 months distant, and $\xi_{April 1992}^e$, when the forecast horizon is 20 months distant. However, because the forecast horizon is tied to the calendar year, this causes a problem for December 1992. Now we are differencing a forecast for calendar year 1994 from a forecast for calendar year 1993: this is comparing apples and oranges. So we must use some of the data from the shorter forecast horizon, the current year data, to ensure a smooth data series. In the end, the forecast horizon on the far edge of the measurement window varies from 10 to 21 months depending on the time of year, but the difference in horizon between the leading and trailing edges of the measurement window is always two months.

4.2 Pre-electoral polling data

To measure the extent of the surprise contained in the election result requires a measure of the ex-ante expected election outcome. The natural place to look is polling data. Hibbs suggested using electoral preference polls over 15 years ago in his review article (Hibbs 1992) and several years prior, Chappell and Keech did so in their study of US unemployment (Chappell and Keech 1988). In recent years, interest in the effect of endogenous elections on the partisan theory has led a few authors to pursue time series polling data of party preference for other developed countries. Unfortunately, reliable time series polling data have proven difficult to acquire. The number of countries in those few studies that utilize polling data is quite small compared to typical cross-country studies of political business cycles which usually feature at least a dozen countries and sometimes as many as 100. Heckelman (2002)

⁵For the 90-day interest rate and the 10-year bond yield the forecast horizons are fixed at three months and 12 months ahead rather than being tied to the calendar year and thus do not exhibit this effect. For these variables I use the 12-month horizon.

studies Canada, Germany, and the United Kingdom, chosen “because of the availability of continuous polling data”. Berleman and Markwardt (2003) manage to assemble data for six countries (Australia, France, Germany, Sweden, the United Kingdom, and the United States) from a variety of sources. Because there is no central source of polling data, finding reliable time-series data for a lengthy period covering multiple elections is a challenge which must be revisited for each country in the study with no guarantee of compatibility.

Luckily, as I study information revealed in a short period around the election date, I do not need time series polling data. For my purposes, it is sufficient to measure the probability of either side achieving victory at the time when the pre-electoral economic forecasts are recorded: one to two months prior to the election. Comparing this probability to the actual election result gives a measure of the electoral surprise realized over the measurement window. Thus electoral surprise and the change in expectations are measured across the same window.

While pre-electoral polling data exist for many countries, it has not, to my knowledge, been compiled before. Table 8 documents the sources from which I have compiled my pre-electoral polling data. The data are based on in-person or telephone interviews featuring the question: “If a {general, parliamentary, presidential} election were held {tomorrow, Sunday}, which {party, candidate} would you vote for?” The raw data yield the frequency with which each candidate or party garnered support as well as frequencies of respondents who are uncertain. For each party i , I have defined v^i as the percentage of survey respondents who expressed an intention to vote for party i . Many polls repeat the question for those who express uncertainty, asking them “Are you leaning toward any particular party?” I have not counted these separate respondents, focusing solely on voters who profess to have made their decision. To form the final v^r (v^l) I have summed the v^i of all right (left)-wing parties which garner at least 5% of the votes in the election. I use the Mannheim Eurobarometer definitions to identify parties belonging to the left and right.

The data have been collected from three different types of sources: general election studies, public opinion polls by public opinion agencies, and opinion polls by major newspapers. My preferred source was general election studies. The goal is to gain a measure of pre-electoral political support at the front edge of the measurement window. Thus I selected the general election study if available and conducted fewer than three months prior to the election. If this was unavailable, I simply chose the most prominent opinion poll available during that period, either from a public opinion agency or a major newspaper (whose polls are usually conducted by public opinion agencies). All of the polls are scientific and feature at least 1000 respondents selected to represent the national electorate (no regional polls, no respondents ineligible to vote, no internet polls). In two cases (Italy 1994, Spain 1996), I could find no polling data so these have been coded missing. The sources for each election and the dates over which the field work for the polls were done are noted in Table 8. Macroeconomic data for the 11 countries have been gathered from the IMF and the EIU. All data are monthly with the exception of data on GDP growth which are quarterly.

5 The construction of key variables

5.1 A measure of political surprise

Data on election dates and outcomes have been assembled using Banks’ Political Handbook of the World (1988–2002) plus various National Election Institutes and include the date of the election, the vote shares of and seats allocated to each party, and the ideology of the

post-electoral government. The polling data are used in conjunction with data on the election outcome to produce SURPRISE, the variable measuring the partisan political surprise from an election. Its magnitude measures the extent of the electoral surprise; its sign indicates the partisan affiliation of the victorious government.

The first step in the construction of SURPRISE is to classify, using Banks’ Political Handbook of the World, the final government produced by each election on a five-point scale from left to right

$$\text{GOV} = \begin{cases} 1 : & \text{left} \\ 0.75 : & \text{center-left} \\ 0.5 : & \text{center} \\ 0.25 : & \text{center-right} \\ 0 : & \text{right} \end{cases}$$

Single party governments earn a pure left or right classification while coalitions encompassing parties with differing economic ideologies earn a diluted center-left, center-right, or even dead-center classification.

The next step is to determine which alternative governments the election is to decide between. For the United States, the alternatives are a Democratic president (GOV = 1) or a Republican president (GOV = 0). Similarly, for Britain the alternatives are a Labor government (GOV = 1) or a Conservative government (GOV = 0). But for Japan, the relevant electoral question during this era is not whether the LDP heads the government, but whether the LDP must form a coalition government which dilutes its policies (GOV = 0.25), or whether it can poll a majority (GOV = 0). Similarly, in the Canadian general elections of 1993, 1997, and 2000, support for the Liberal party was so overwhelming that the important electoral margin was not between pure left and pure right, but between pure left (GOV = 1) and center-left (GOV = 0.75). Finally, in multiple Dutch elections (1998, 2002, and 2003), three parties of roughly equal size implied that neither left nor right would obtain a majority. Thus in this case the alternatives were center-left (GOV = 0.75) and center-right (GOV = 0.25). Excepting these three countries, the alternatives are taken to be left-wing majority (GOV = 1) and right-wing majority (GOV = 0).

Having established the likely alternatives, the task is to assign ex-ante probabilities to these alternatives, or rather, to the probability that the more conservative alternative is realized on Election Day. This probability is generated using pre-electoral opinion poll data.⁶ Think of an opinion poll as a repeated draw from a trinomial distribution (voter prefers party R, party L, or some other party) with unknown parameters q_r and q_l which indicate the probability that a given respondent prefers party R or party L. If v^i denotes the vote-share party i receives in pre-electoral opinion polls and N is the size of the poll, then for large N , the difference $v_i - v_j$ is normally distributed with standard error

$$\sigma_{ij} = \sqrt{\left[\frac{v_i(1 - v_i)}{N} + \frac{v_j(1 - v_j)}{N} + 2 \frac{v_i v_j}{N} \right]} \tag{1}$$

⁶It is frequently argued that pre-election polls are unreliable predictors of election results. Recall, however, that for this study, we are not looking for an accurate predictor, but rather for a measure of forecasters pre-election expectations of the political results. So pre-election opinion polls need not be perfect predictors of the final outcome, they simply need to be the political predictors off which forecasters base their economic forecasts. This is a weaker constraint and one which, given the centrality of opinion polls in public debate and among the punditry, I find acceptable.

Then the probability that $q_r > q_l$ and thus that the right wing party will win Election Day given the polling results v_r, v_l is given by

$$P = \Phi\left(\frac{v_r - v_l}{\sigma_{rl}}\right) \quad (2)$$

where Φ is the cumulative standard normal distribution.⁷

This method works for the cases where the important margin is competition between two parties. For those elections (Japan 1993, 1996, 2000, 2003 and Canada 1993, 1997, 2000) where the electoral uncertainty concerns whether a single party i can win a majority of the seats, the standard error is much simpler:

$$\sigma_i = \sqrt{\frac{v'_i(1 - v'_i)}{N}}, \quad (3)$$

$$P = \Phi\left(\frac{v'_i - 0.5}{\sigma_i}\right) \quad (4)$$

The sole complication arises from the disconnect between vote share and seat share. While polling data measure expected vote share, parliamentary governments rely on support from a majority of seats. The well-known phenomenon is that small parties tend to garner a non-negligible fraction of the national vote, but are unable to muster a majority in a corresponding fraction of the districts and so fail to win seats commensurate with their vote share. Simply omitting parties which fail to poll 5% of the vote and renormalizing the remaining parties brings vote share and seat share broadly in line. Thus vote shares v'_i in (3) and (4) have been adjusted by omitting parties supported by less than 5% of the respondents.

The final value of SURPRISE is given by

$$\text{SURPRISE}_{ct} = G_{ct} - L_{ct} - (R_{ct} - L_{ct})P_{ct} \quad (5)$$

where G_{ct} is the value of the final government on the five-point GOV scale following the election in country c at time t , L_{ct} and R_{ct} are the values of the left and right alternatives on that same scale, and P_{ct} is the probability from (2) or (4). Each of these steps can be viewed in light of Fig. 2. Classifying the alternatives at stake in the election (R_{ct} and L_{ct}) on the five-point GOV scale is akin to measuring the distance between the liberal and conservative alternatives.⁸ Converting polling data to the expected probability of a conservative victory (P_{ct}) and then scaling by the distance between the two alternatives ($R_{ct} - L_{ct}$) gives a measure of the pre-election expectation. Subtracting this from the resulting government (G_{ct}) yields a measure of the degree to which the ex-post election result differed from the ex-ante expectation.

⁷The opinion polls I use range in size from 1000 to 9000 respondents.

⁸One commonly raised objection is that this method imposes a uniform left-right scale on all countries when clearly the right in England is rather different from the right in the United States and quite possibly less conservative on many issues than the American left. While this is true, it is not relevant to the measurement at hand. SURPRISE is based on the difference between the left and right in a particular country. Heterogeneity in the median of the national political spectrum washes out when this first difference is taken. On the other hand, the extent to which countries differ in the degree of polarization—the distance between the left and right—is suppressed by the common scale. There are, however, two mitigating factors. First, Center-Right and Center-Left classifications address the coarsest differences across countries. Second, the grouping of countries by variety of capitalism presumably mitigates the problem further as political systems are likely to be more similar within a group than across them.

The sample consists of 30 elections.⁹ Changes of power not associated with elections have been omitted due to the difficulty in obtaining measures of political surprise which are consistent with the surprise contained in the elections.¹⁰ The sample is well-balanced between the left and right: 14 left-wing governments, 14 right-wing governments, and two centrist coalitions.

Summary statistics for SURPRISE are included in Table 2. There are 13 elections for which the results generated negligible SURPRISE, six results generating positive (liberal) SURPRISE and nine results generating negative (conservative) SURPRISE. The average absolute value of SURPRISE during the non-negligible elections is roughly 0.25. The largest absolute values come from United States 2000 and Spain 2004 which accords well with qualitative accounts of these elections. For windows that do not contain an election, SURPRISE takes a value of zero.

This method for constructing ex-ante expectations makes two implicit assumptions. First, that polls reflect the fundamental preferences of the electorate modulo sampling error. Second, that the fundamental preferences of the electorate are relatively stable between the poll and the election and thus the extent of the ex-ante uncertainty in the outcome is given by the sample error of the poll. Wlezien and Eriksson (2002) have conducted a careful study of the time-series behavior of polls in US presidential elections from 1948 to 2000. They find that the party conventions seem to mark an important dividing line between early and late campaigns; each of which is characterized by distinctive behavior. During the early campaign (100–200 days preceding the election), campaign shocks are large but temporary. News can have a significant effect on voters' preferences but is forgotten by Election Day. Polls are highly volatile but most of this is due to transient shocks and survey error. Because voters tend eventually to forgive and forget, fundamental preferences are essentially a stationary time-series. During the late campaign (the final 100 days preceding election), campaign shocks are smaller but do not completely dissipate before Election Day. Preferences are an integrated series. However, the shocks, though permanent, are much smaller in magnitude. Hence any given poll conducted during the late campaign is a much better predictor of the final results than one taken during the early campaign. A simple linear regression of vote-division on poll results delivers an adjusted- R^2 of 0.7 for polls 60 days prior to the election, rising to 0.85 on the eve of the election.

Their results confirm that polling data are tightly tied to fundamental preferences, especially during the late campaign, and thus constitute an appropriate measure of ex-ante political expectations. They further suggest that some portion of the measurement error in a poll is due to the fact that fundamental preferences will move between the date of the poll and Election Day. The rest of the measurement error, and they emphasize that this remains the larger part, especially during the later period, is due to sampling. This additional source of uncertainty over the poll results—changes in fundamental preferences as well as survey error—means that my method overestimates the likelihood of the favored party winning. As a result, when the favored party does win, SURPRISE is an under-estimate of the changes in expectations a rational poll-watching agent would feel, but when the favored party loses,

⁹In most countries I have chosen legislative elections for the lower house. In the United States, the only presidential country in the data set, I look at presidential elections. This is consistent with Alesina, Roubini, and Cohen and other authors.

¹⁰For an attempt to combine both latent and electoral surprise, see Berlemann and Markwardt (2003).

SURPRISE is an over-estimate. Wlezien and Eriksson do emphasize that during the late campaign period these movements, and thus the source of error, are small. However, it is not clear in what direction this measurement error influences my coefficient estimates. Addressing this source of (unknown) bias would require more detailed time-series polling data and has thus been left for future work.

As Wlezien and Eriksson study only US presidential elections, generalizing their insights to a study of multiple countries with a variety of electoral systems and traditions requires caution. In particular, the 100 day cutoff between the early campaign, when voters are open-minded but less attentive, and the late campaign, when the reverse is true, is probably a result of the particular timing and structure of US presidential elections. Wlezien and Eriksson postulate that party conventions may instigate the shift from the first regime to the second by prompting voters to pay consistent attention. Nonetheless, it is not hard to believe—though evidence does not yet exist—that electorates in general exhibit some kind of transition from a period in which fundamental preferences are volatile but stationary to one in which preferences are integrated but display much lower volatility.

Wlezien and Eriksson's results highlight the importance of keeping my event window narrow so as to minimize this source of measurement error. Their categorization of different periods of campaign behavior further highlights the importance of ensuring that the ex-ante poll is taken during the period of "late campaign" behavior when polls more closely align with fundamental preferences.

5.2 The event window

Because the dependent variable is the change in forecasts, it is necessary to choose the number of months over which to compute that change. Uncertainty over the future fiscal policy is not resolved solely on Election Day: it may be resolved during the campaign as one party dominates the polls, or persist beyond the vote due to the post-election process of forming a governing coalition. The pattern of resolution is unique to an election. My purpose is to focus solely on elections as large and rapid changes in political forecast which are (relatively) easily measured. I focus on the election and formation of the governing coalition and leave aside the ups and downs of the campaign. I have chosen a two-month window in order to capture both the election and the formation of the governing coalition. This means that any dependent variable in a given month is the change in the forecast between the current survey and the survey two months earlier.¹¹

Using a two-month window generates autocorrelation in the LHS variable. If the LHS variable is the difference in the level of ξ_t across the past two months, but is recorded with monthly frequency, then these windows overlap. This means any particular month-to-month change in ξ_t is recorded in two consecutive values of $\Delta\xi_t$, guaranteeing that $\Delta\xi_t$ exhibits high autocorrelation. Thus I allow for autocorrelation in my error term.

¹¹ Rerunning the regressions with one- and three-month windows show the importance of post-election coalition formation. A three-month window gives similar point estimates but larger standard errors indicating that the larger window has simply picked up more noise with little extra signal. On the other hand, a one-month window results in smaller point estimates indicating that a large part of the movement in forecasts takes place the month following the election. This suggests that a significant part of election uncertainty is resolved by the post-electoral formation of the governing coalition rather than on Election Day.

6 Estimation: equation and technique

The basic specification for this study takes the form

$$\begin{aligned} \Delta \xi_{f,t}^e &= \alpha + \sum_{j=1}^3 \beta_j [X_{c,t-3(j-1)} - X_{c,t-3j}] \\ &+ \delta_1 \text{PoliticalSurprise}_{c,t} \\ &+ \delta_2 \text{PoliticalSurprise}_{c,t} * \text{MME}_c \\ &+ \delta_3 \text{PoliticalSurprise}_{c,t} * \text{CME}_c + \epsilon_{f,t}, \end{aligned} \quad (6)$$

$$\Delta \xi_{f,t}^e \equiv \xi_{f,t}^e - \xi_{f,t-2}^e \quad (7)$$

where ξ is one of the macroeconomic variables for which panelists record forecasts. I look at seven of these variables: GDP growth, CPI inflation, 90 day interest rates, 10 year bond yields, the level of unemployment, growth in household consumption, and growth in business investment. The first subscript, f , refers to the forecaster which is the cross-sectional unit of analysis. In cases where the variable is constant across all forecasters within the same country, this subscript has been changed to c to remind the reader (and the writer) of this fact. The second subscript refers to the time dimension of the panel and is indexed monthly. All forecast data is monthly. Data on realized variables are monthly except GDP growth data, which are quarterly. The dependent variable is the change in expectations over the measurement window. X is a vector of five macroeconomic variables: output growth, unemployment, inflation, and short and long-term rates. The summation on the RHS consists of the changes in the realized values of these variables over the three most recent quarters.¹² MME and CME are indicators of whether the country is a Mediterranean market economy or a coordinated market economy. Liberal Market economies comprise the omitted baseline. The classification is taken from Hall and Soskice (2001). The specification allows partisan political surprise to affect forecasts differently in different economies.

The next question is how to estimate (6). At each point in time for each country, we draw a number of different forecasting models (panelists) from the urn and record their forecasts. Having multiple panelists for each country is something like having repeated draws of the same experiment. The fact that several observations have been made on the same country on the same date raises the possibility of contemporaneous correlation: all the panelists from one country have something in common in their forecasting due to the fact that they are looking at the same country and change expectations at the same time in response to the same events. Country-specific shocks not accounted for by the set of explanatory variables are picked up by the error term, resulting in contemporaneous correlation of errors from forecasters in the same country.¹³ Positive correlation of errors among forecasters observing

¹²These are not quarters as defined by the calendar year but rather they are the three most recent three-month periods. Thus in May, the three most recent “quarters” cover February-May, November-February, and August-November. For GDP growth, this will actually correspond to the change over the past three calendar-year quarters because data are quarterly. But as the other series are monthly, this allows for the effects of the most recent past data.

¹³An excellent example of this is the terrorist attacks of September 11th, 2001. The data show a swift and sizable reaction by panelists predicting severe negative consequences for the US economy. This reaction in expectations was certainly not captured by any of the explanatory variables. And while panelists predicted dips in other countries as well, the movement of expectations was understandably much higher in the United

the same country (as is expected) would result in underestimation of the true standard errors. To correct for this, I cluster by country.

Another possible problem is serial correlation in the errors. Gallo et al. (2002) present evidence of herding behavior in Consensus data for the United States, United Kingdom, and Japan for 1993–1996 and suggest that such herding might generate first-order serial correlation in forecast errors as forecasters attempt to hit the group mean or follow a leader. The herding discussed by Gallo et al. would imply negative serial correlation in the errors of (6). Imagine that the economy is hit with a shock about which panelists have heterogeneous opinions, resulting in heterogeneous changes to their forecasts. Then, according to the story, before answering the second post-shock survey, the panelists look at their position relative to their peers in the first survey after the shock and adjust toward the mean forecast. This behavior would result in negative autocorrelation of my LHS variable as those who reacted most strongly to the event will then take a step back in the other direction once they have observed their peers. The story remains unchanged if the panelists all follow a particular leading analyst rather than reverting to the mean.¹⁴

However, visual inspection of the data suggests a different pattern. If one sorts forecasts from bullish to bearish, panelists seem to stake out relative positions within the array of forecasts and adhere to them for several consecutive surveys rather than herding toward the mean or toward a leading forecaster. This high persistence in the forecasts is probably due to the lengthy forecast horizon and thus the long feedback time. However, this implies autocorrelation in the level, but not necessarily in the first difference of the forecasts. Indeed, I find no evidence of the negative serial correlation in the first difference of the errors that would come from herding. Nonetheless, I do control for serial correlation because the two-month measurement window implies first-order serial correlation.

7 Results

Despite strong *a priori* reasons for believing that the underlying institutions of capitalism affect the nature of politically induced macroeconomic volatility, it is worth taking a broader look before estimating (6) to be sure that varieties of capitalism describe the underlying variation.

Figure 3 shows that partisan effects in output growth forecasts vary by institutional group. First the sample is sorted into three groups by variety of capitalism. Second, within each group of countries, elections are identified and split according to whether the liberal or the conservative alternative won. Then, for all the elections in an institutional group which generate the same direction of political surprise, the average monthly change in forecasts is plotted for nine months before and after the election. Such a series plots, for example, the average path of expectations in a liberal market economy before and after an election in which the liberal alternative was elected. Plotting the series for liberal and conservative outcomes is essentially a graph of the right-hand plot from Fig. 2. The extent that we see separation

States than for any other country. Clearly then, the error terms for the months after September 2001 were large and in the same direction for all US panelists. While 9/11/2001 is exceptional in magnitude, other country-specific events undoubtedly produce similar, if smaller, effects.

¹⁴This analysis assumes that a panelist's reaction to the event is uncorrelated with his previous position relative to the panel mean. If, in reaction to positive news, those who were previously pessimists react more strongly than those who were previously optimists, then we might support positive serial correlation of the errors as a result of herding.




Fig. 3 The evolution of output growth forecasts near an election

Table 3 Average pairwise correlation of forecasts within the same country

Country	y	Δy	CPI	ΔCPI
Canada	0.52	0.17	0.37	0.12
France	0.69	0.23	0.42	0.11
Germany	0.63	0.23	0.53	0.13
Italy	0.64	0.18	0.39	0.13
Japan	0.41	0.19	0.56	0.09
Netherlands	0.66	0.14	0.43	0.16
Norway	0.46	0.14	0.64	0.12
Spain	0.50	0.14	0.50	0.09
Sweden	0.45	0.13	0.49	0.10
UK	0.42	0.12	0.42	0.07
USA	0.53	0.22	0.46	0.10

$\Delta\xi$ is the one month change in forecast ξ

between the two lines in a single panel around the election is the extent to which forecasters respond differently to liberal and conservative surprise. And the extent to which the different panels display different patterns is the extent to which countries with different varieties of capitalism experience politically induced macroeconomic volatility of a different character. The three panels of Fig. 3 show three rather different patterns. In liberal market economies, conservative electoral surprise is greeted by a drop in output forecasts; in Mediterranean market economies there is a similar bias against conservative governments though the reaction is much less pronounced and less obviously tied to the election; and in coordinated market economies, the election of conservative governments is greeted with an increase in output forecasts and the election of liberal government by a decrease in output forecasts.

Figure 3 gives further evidence that the three groups display different responses to election surprise. But more importantly, looking at the time series shows that there is a real response to elections. The series for liberal and conservative victory separate in response to the election but move roughly together before and after. Now we can run (6) with confidence in the specification.

I've run (6) for seven different variables: growth rate of output, level of unemployment, change in household consumption, change in business investment, change in consumer prices, level of 90 day interest rates, and level of 10 year bond yield.¹⁵

Table 4 shows the raw coefficients from (6). The total effect of SURPRISE on in LMEs is captured by δ_1 , for MMEs it is $\delta_1 + \delta_2$, and for CMEs it is $\delta_1 + \delta_3$. Table 5 sums the proper terms to deliver the total partial effect by capitalist system for each of the seven macroeconomic variables over which forecasts are recorded.

The regression coefficient is the partial effect of a unit change in SURPRISE. SURPRISE = 1 means that the right was expected with certainty but that the left was actually victorious. Therefore, the regression coefficient, the partial change in forecasts due to a

¹⁵As mentioned in the Data section above, the original sample in 1989 consists of the G7; the four other countries are added in two waves in January 1995 (Netherlands, Spain, Sweden) and June 1998 (Norway). Since these countries have been added in descending order of GDP, and larger countries tend to have less volatile GDP growth, there is some small danger of bias in the GDP growth regressions due to sample bias. To be sure, I reran the regressions using only data since June 1998, the date at which the sample was full. The results are not markedly different from the full sample.

Table 4 Basic regression results

ξ	Dependent variable						
	GDP Growth	HH Cons.	Bus. Invest.	Unemp.	CPI	90-Day Rates	10 Year Yield
SURPRISE (δ_1)	0.699 [0.030]***	0.200 [0.045]***	3.033 [0.086]***	-0.026 [0.017]	-0.006 [0.015]	0.460 [0.043]***	0.184 [0.036]***
SURPRISE*MME (δ_2)	-0.097 [0.094]	0.236 [0.209]	-2.202 [0.221]***	-0.323 [0.122]**	0.346 [0.195]	0.421 [0.843]	-0.171 [0.714]
SURPRISE*CME (δ_3)	-1.869 [0.401]***	-1.439 [0.452]***	-5.575 [1.504]***	0.353 [0.112]**	0.218 [0.137]	-1.166 [0.330]***	-0.359 [0.070]***
Observations	21052	20935	20672	16558	20965	19259	18836

Standard errors in brackets

* significant at 10%

** significant at 5%

*** significant at 1%

GLS with dependent variable $\Delta \xi_t^c$

Controls: change in the realized values of output growth, unemployment, consumer price inflation, 3-month interest rates, and 10-year bond yields over the previous three quarters plus a constant

Table 5 Total partial effect by capitalist system

ξ	Dependent variable						
	GDP Growth	HH Cons.	Bus. Invest.	Unemp.	CPI	90-Day Rates	10 Year Yield
LMEs (δ_1)	0.699 [0.030]***	0.200 [0.045]***	3.033 [0.086]***	-0.026 [0.017]	-0.006 [0.015]	0.460 [0.043]***	0.184 [0.036]***
MMEs ($\delta_1 + \delta_2$)	0.602 [0.091]***	0.436 [0.209]*	0.831 [0.165]***	-0.35 [0.113]***	0.339 [0.189]	0.881 [0.856]	0.014 [0.717]
CMEs ($\delta_1 + \delta_3$)	-1.170 [0.404]**	-1.239 [0.441]**	-2.542 [1.536]	0.326 [0.116]**	0.212 [0.130]	-0.706 [0.326]*	-0.174 [0.075]**
Observations	21052	20935	20672	16558	20965	19259	18836

Standard errors in brackets

* significant at 10%

** significant at 5%

*** significant at 1%

GLS with dependent variable $\Delta\xi_t^e$

Controls: change in the realized values of output growth, unemployment, consumer price inflation, 3-month interest rates, and 10-year bond yields over the previous three quarters plus a constant

unit change in SURPRISE, represents the full perceived difference between the macroeconomic performance under the policies of the left and performance under the policies of the right.¹⁶ The data are measured in percentage points; thus the coefficient on output growth corresponds to a partisan difference in expected output growth of 0.7% for the three LMEs, 0.6% for the three MMEs, and 1.2% for the five CMEs in my sample during the period 1993–2004.¹⁷

Because the study contains relatively few events, there is always the possibility that the results are driven by an outlier. For each of the 30 elections, I have rerun the regressions with the election in question dropped. Point estimates respond in only a limited manner and the broad story being told is not affected implying that these results are not being driven by any single election.

Looking first at the baseline, δ_1 , from Table 4 shows that in LMEs, conservative governments are expected to produce a relative recession—decrease in both consumption and investment growth and consequent drop in interest rates—but with little effect on unemployment or inflation. Now look at the difference between LMEs and MMEs, captured by δ_2 in Table 4, or compare the total partial effect of SURPRISE in MMEs with that for LMEs by looking at Table 5. In MMEs, like LMEs, liberal governments are expected to produce a boost to consumption and investment growth and thus output growth but the effect differs between these groups of countries in three respects. First, forecasters expect the output growth increase in MMEs to be concentrated more on consumption and less on investment than that in LMEs. At the same time, forecasters in MMEs expect significantly lower unemployment under liberal governments. This indicates that the partisan divide over fiscal stimulus is different in LMEs and MMEs.

¹⁶In a race between pure left and right which is dead even according to the polls, we would expect an Election Day movement in forecasts of roughly half the magnitude of the reported coefficient.

¹⁷For comparison, Alesina et al. (1997) report a partisan difference in realized output growth of 1.4% for the United States during the period 1949–1994. See ARC Table 4.2

Table 6 Forecast errors in election and non-election years

ξ	Dependent variable						
	GDP Growth	HH Cons.	Bus. Invest.	Unemp.	CPI	90-Day Rates	10 Year Yield
Election Year Dummy (β)	0.187 [0.180]	-0.034 [0.192]	-0.242 [0.733]	0.129 [0.132]	-0.020 [0.102]	0.057 [0.054]	0.045 [0.045]
Constant ($\alpha\beta$)	0.183 [0.143]	-0.061 [0.152]	1.460 [0.516] ^{***}	0.307 [0.088] ^{***}	0.220 [0.091] ^{**}	0.022 [0.116]	0.035 [0.085]
Observations	28186	27202	27824	29110	24103	27601	26841

Standard errors in brackets

* significant at 10%

** significant at 5%

*** significant at 1%

GLS with dependent variable $\xi_t - E_{t-h}\xi_t$

Table 7 Forecaster's use of political information

ξ	Dependent variable						
	GDP Growth	HH Cons.	Bus. Invest.	Unemp.	CPI	90-Day Rates	10 Year Yield
Δ ActualGov (γ_1)	0.020 [0.057]	-0.004 [0.020]	0.755 [0.107]	0.049 [0.035]	-0.241 [0.027]***	-0.005 [0.070]	-0.324 [0.037]**
Δ ExpectedGov (γ_2)	0.224 [0.039]	-0.070 [0.014]	1.391 [0.594]	0.336 [0.117]	0.214 [0.126]*	0.035 [0.091]	0.042 [0.097]
Constant (α_γ)	-0.057 [0.139]	0.016 [0.138]	0.070 [0.518]***	0.057 [0.076]***	0.002 [0.084]**	-0.016 [0.118]	0.074 [0.087]
Observations	28163	27179	27801	29087	24096	27582	26823

Standard errors in brackets

* significant at 10%

** significant at 5%

*** significant at 1%

GLS with dependent variable $\xi_t - E_{t-h}\xi_t$

Table 8 Opinion poll data sources. Question: “If an election were held tomorrow, for which candidate/party would you vote?”

Country	Election date	Source	Time to election
Canada	10.25.1993	Canadian National Election Study	5 weeks – 1 day
Canada	6.2.1997	Canadian National Election Study	5 weeks – 1 day
Canada	11.27.2000	Canadian National Election Study	5 weeks – 1 day
Canada	6.28.2004	Ipsos-Reid	5 weeks – 1 week
France	3.28.1993	Eurobarometer	1 month
France	6.1.1997	Eurobarometer	3 months
France	6.16.2002	Le Figaro (Ipsos)	1 month
Germany	10.16.1994	German Election Study (Politbarometer)	7 weeks – 1 week
Germany	9.27.1998	German Election Study (Politbarometer)	6 weeks – 1 week
Germany	9.22.2002	Politscope	1 month
Italy	3.27.1994	.	.
Italy	4.21.1996	Corriere della Sera (Ispo-Cra/Nielsen)	1 month
Italy	5.13.2001	Corriere della Sera (Ispo-Cra/Nielsen)	2 months
Japan	7.18.1993	NHK	2 weeks
Japan	10.20.1996	Asahi Shimbun	2 weeks
Japan	6.25.2000	Asahi Shimbun	2 weeks
Japan	11.9.2003	Asahi Shimbun	2 weeks
Netherlands	5.6.1998	Dutch Parliamentary Election Study	5 weeks – 1 day
Netherlands	5.15.2002	Dutch Parliamentary Election Study	5 weeks – 1 day
Netherlands	1.23.2003	NRC Handelsblad	2 weeks
Norway	9.10.2001	Verdes Gang (Norsk Gallup)	1 month
Spain	6.6.1993	Eurobarometer	3 months
Spain	3.3.1996	.	.
Spain	3.12.2000	El Pais (Demoscopia)	1 week
Spain	3.14.2004	Instituto Opina para la Cadena SER (Pulsometro)	2 weeks

Table 8 (Continued)

Country	Election date	Source	Time to election
Sweden	10.20.1998	Swedish Election Study	2 months – 1 week
Sweden	9.15.2002	SIFO	3 months
UK	5.1.1997	Eurobarometer	2 months
UK	6.7.2001	British Election Study	1 month – 1 day
USA	11.5.1996	National Election Study	2 months – 1 day
USA	12.12.2000	National Election Study	2 months – 1 day

Second, unlike their counterparts in LMEs, forecasters in MMEs expect a significant increase in inflation along with the increased output growth. One explanation is that the forecast horizon is typically longer in MMEs than LMEs because their elections occur, on average, earlier in the year. In my sample, the average MME election takes place on the 113rd day of the year whereas the average LME and CME elections take place on the 230th and 232nd days of the year. This means that the forecast horizon for output, inflation, and unemployment (but not interest rates which have a fixed twelve month horizon) is, on average, four months longer in MMEs than LMEs and CMEs. Perhaps the effect of fiscal policy on unemployment and inflation is substantially greater in months 8 to 20 than months 4 to 16. I cannot rule this out. More likely, however, is that there exists either a difference in the partisan cleavage between LMEs and MMEs, or some structural difference between the economies such that a given partisan difference has different effects.

A likely candidate for the different reactions of inflation forecasts is the difference in central bank independence between the two groups. A more conservative central bank in the sense of Rogoff (1985) places lower weight on output fluctuations. As a result, a given partisan shift in fiscal policy has a lesser effect on inflation. The literature on central bank independence documents the relatively high independence of the central banks in the United States and Canada compared to much lower independence in pre-ECB Spain and Italy (Cukierman 1992). The three MMEs average just 1.75 on Alesina and Summers' (1993) index of CBI while the three LMEs average 2.7 (and the five CMEs average 2.6).

However, while CBI might explain patterns in inflation forecasts, it cannot shed light on the behavior of forecasts of unemployment, investment, and consumption. The fact that liberal governments in MMEs are expected to lower unemployment more than their counterparts in LMEs and to boost consumption more than investment is probably due to the different nature of the partisan divide on fiscal policy. In other words, liberal governments in LMEs produce expansions via one set of fiscal policies while liberal governments in MMEs produce expansion via a different set of policies. However, a thorough investigation of this question will have to wait for a model of strategic interaction between the fiscal and monetary authorities.¹⁸

The last interesting aspect concerns the different effects on the yield curve and constitutes further evidence that the interaction between the fiscal and monetary authorities is an important determinant of the character of partisan cycles. In LMEs, it is expected that the boom from liberal governments will result in a moderate increase in next year's three-month rates and next year's 10-year rates. In MMEs, the forecast is for a very strong effect on next year's three-month rates but virtually no effect on next year's 10-year rates. In other words, in LMEs, the difference between the left and right is a translation of the forecasted yield curve whereas in MMEs a the difference is a *rotation* of the forecasted yield curve. Evans and Marshall (2001) argue that fiscal policy, to the extent that it affects households' marginal propensity to consume, ought to result in a translation of the yield curve; the only policy shock that results in a rotation of the yield curve is a monetary policy shock. These results, then, suggest that forecasters in the 1990s believed not only that fiscal, but also monetary policy in MMEs varied with the elected policy-makers. Apparently economic forecasters agreed with Alesina and Summers that central banks in MMEs were less autonomous than those in LMEs.

¹⁸Dixit and Lambertini (2003) do model the strategic interaction of independent fiscal and monetary policy authorities but their model is based on a new-Keynesian framework and thus a completely forward-looking Phillips curve. It is thus unsuited to the current purposes.

Shifting attention to CMEs, I find that the partisan cleavage in these countries is radically different than that in MMEs or LMEs. In CMEs, it is *conservative* governments which are expected to produce an output boom relative to their liberal counterparts. The increased output growth associated with conservative governments in CMEs is, relative to that associated with liberal governments in LMEs, more strongly visible in consumption than in investment and is accompanied by a decrease in unemployment, again suggesting a significantly different partisan cleavage over fiscal policy. But though they produce a boom, conservative governments are also associated with a *decrease* in inflation. As far as interest rates, conservative governments are associated with higher rates in both the short and the long term: a translation of the yield curve. This suggests that in CMEs, partisan differences in policy are confined to fiscal policy, which is consistent with the relatively high average index of CBI among these countries. But in CMEs, conservative fiscal policy is perceived to be strictly superior to liberal fiscal policy, delivering higher output growth, lower unemployment, and lower inflation.

On the whole, the three groups are characterized by significantly different patterns. In the LMEs and MMEs, it is liberal parties who are associated with expansion while in CMEs it is the conservative parties. The character of that expansion also differs across the groups. In LMEs, forecasters expect the left to achieve extra growth, without an effect on unemployment, at virtually no cost to inflation. In MMEs, forecasters expect the left to deliver both higher output growth and lower unemployment, but at the cost of higher inflation. And in CMEs, forecasters expect the right to deliver higher growth, lower inflation, and lower unemployment. The clear implication is that partisan differences over fiscal policy vary considerably by country.

8 Forecast quality

This paper measures the change in economic forecasts induced by a change in expected political outcomes and, from this response, infers the perceived difference between the two parties' platforms. The degree to which this perceived difference corresponds to the actual difference between the parties' platforms depends on the ability of the forecasters to properly process election news. I look at this in two parts. First, are forecast errors larger in election years? Second, are forecast errors correlated with the surprising and/or unsurprising components of political surprise which were revealed during the previous month (and thus ought to be incorporated if agents properly process the information)?

$$\xi_t - E_{t-h}\xi_t = \alpha_\beta + \beta \text{ElectionYearDummy}_{t-h} + \mu_t \quad (8)$$

$$\xi_t - E_{t-h}\xi_t = \alpha_\gamma + \gamma_1 \Delta(\text{ActualGov})_{t-h} + \gamma_2 \Delta(\text{ExpectedGov})_{t-h} + \nu_t \quad (9)$$

Estimates of α_β reveal that forecasters systematically overestimated investment, unemployment, and inflation during this period. But this mistake does not seem to be indicative of an inability to process political information. The failure to reject $\beta = 0$ for any of the seven variables suggests that forecasts are not systematically worse in election years. Estimates of γ_1 suggest that forecasters systematically underestimate the negative (positive) effect of a liberal (conservative) government on ten-year bond yields. However, in general, the estimates are insignificant, suggesting that the political information, including both the surprising component revealed in the election (γ_2) and the component revealed prior to the election (γ_1), is well incorporated in post-election forecasts.

Gallo et al. have suggested that Consensus forecasts often suffer from persistent errors because forecast horizons are long and thus feedback is slow. Moreover, if the underlying economy about which panelists are forecasting evolves, slow feedback may ensure that the time-scale on which errors are perceived and eliminated is not much faster than that over which new sources of error are generated. This is perhaps true concerning their ability to properly process the array of economic events which, in the aggregate, push the macroeconomy this way or that but can be difficult to properly weigh and interpret. However, these results suggest that Consensus forecasters process the information content of (highly visible) national elections in a manner approximating the canonical rational, forward-looking agent with a properly specified model.

9 Conclusion

This paper has developed a new method for the measurement of politically induced macroeconomic fluctuations. Measuring the change in economic forecasts as a function of the change in expected policy delivers a measure of the perceived difference between the policies of the left and right. The results demonstrate that forecasters do respond swiftly and significantly to the information content of elections. However, the effective difference between the policies of the left and the right is more complex than the prevailing view. Policy is not a uni-dimensional choice of the degree of aggregate demand stimulus, but also involves both regulatory policy and the choice of the method by which to deliver fiscal stimulus. Because the policy-space is multi-dimensional, the axis of variation between left and right varies by country. The institutions by which an economy is organized affect the framing of and feasible solutions to the policy issues which define left and right. As a result, the effective difference between fiscal policies of the left and right depend on the underlying political economy. Among the 11 advanced economies that comprise my sample, there exist three distinct patterns in the response of forecasts to electoral surprise corresponding to the three distinct varieties of capitalism. In liberal market economies, forecasters expect the left to achieve higher output growth, without affecting unemployment, at virtually no cost to inflation. In Mediterranean market economies, forecasters expect the left to deliver both higher output growth and lower unemployment, but at the cost of higher inflation. And in coordinated market economies, forecasters expect the right to deliver higher growth, lower inflation, and lower unemployment. Policy cleavages between the left and the right, and hence politically induced macroeconomic volatility, vary a great deal across different systems of capitalism.

These results are not only a reminder that policy is multi-dimensional, they are important evidence that economic agents perceive significant differences in partisan policies. Moreover, they are evidence that politically induced macroeconomic volatility is not uniform, but is a product of the institutional environment. To the extent that this paper establishes the importance of the institutions of wage bargaining, vocational training, shareholder protection, and others in both influencing policy and channeling its effects, the next step is building an understanding of how they do so.

This suggests at least three directions for future research. First, examining the reaction of different types of forecasters would give us an idea of the degree to which various different pricing agents in the economy react to political information. Second, opinion polls are gradually giving way to political markets which generate continuous measures of political forecasts. Use of such data would enable this kind of study to be extended to include campaigns as well, thus capturing a wider span of political information which may cause

macro-economic volatility. Finally, as forecast series become long enough to cover multiple elections, this approach can be extended to countries from other regions such as Latin America and Africa. The attendant increase in institutional diversity may help shed light on the links between specific institutions and policy.

Acknowledgements The author would like to thank Jon Bendor, Olivier Blanchard, Bob Hall, Douglas Hibbs, Francisco Rodriguez, William Shughart, Justin Wolfers, Romain Wacziarg, and a pair of anonymous referees for helpful comments and to seminar participants at the Stanford Graduate School of Business, the Stanford Department of Economics, Georgetown School of Foreign Service, Wesleyan University Department of Economics, Claremont McKenna Department of Economics, Reed College Department of Economics, the 8th INFER Annual Conference, and the First World Meeting of the Public Choice Societies for useful feedback. Further thanks to Peter Latusek and the Stanford GSB Jackson Library for help acquiring the data.

References

- Alesina, A. (1987). Macroeconomic policy in a two-party system as a repeated game. *Quarterly Journal of Economics*, 102, 651–678.
- Alesina, A. (1988). Macroeconomics and politics. *NBER Macroeconomics Annual*, 11–55.
- Alesina, A. (1989). Inflation, unemployment, and politics in industrial democracies. *Economic Policy*, 8, 55–98.
- Alesina, A., & Rosenthal, H. (1995). *Partisan politics, divided government, and the economy*. Cambridge: Cambridge University Press.
- Alesina, A., & Summers, L. (1993). Central bank independence and macroeconomic performance. *Journal of Money, Credit, and Banking*, 25, 157–162.
- Alesina, A., & Sachs, J. (1988). Political parties and the business cycle in the United States, 1948–1984. *Journal of Money, Credit, and Banking*, 20, 63–82.
- Alesina, A., Roubini, N., & Cohen, G. (1997). *Political cycles and the macroeconomy*. Cambridge: MIT Press.
- Banks, A., Overstreet, W., & Muller, T. (1988–2002). *Political handbook of the world*. Washington: CQ Press.
- Barro, R., & Gordon, D. (1983). Rules, discretion, and reputation in a model of monetary policy. *Journal of Monetary Economics*, 12, 101–121.
- Berlemane, M., & Markwardt, G. (2003) *Partisan cycles and pre-electoral uncertainty*. Unpublished manuscript.
- Chappell, H., & Keech, W. (1988). The unemployment consequences of partisan monetary regimes. *Southern Economic Journal*, 55, 107–122.
- Cukierman, A. (1992). *Central bank strategy, credibility, and independence*. Cambridge: MIT Press.
- Dixit, A., & Lambertini, L. (2003). Interactions of commitment and discretion in monetary and fiscal policies. *American Economic Review*, 93, 1522–42.
- Drazen, A. (2000). The political business cycle after 25 years. *NBER Macroeconomics Annual*, 15, 75–117.
- Evans, C., & Marshall, D. (2001). *Economic determinants of the nominal treasury yield curve*. WP-01-16, Federal Reserve Bank of Chicago.
- Gallo, G., Granger, C., & Jeon, Y. (2002). Copycats and common swings: the impact of the use of forecasts in information sets. *IMF Staff Papers*, 49, 4–21.
- Grier, K. (2008). US presidential elections and real GDP growth, 1961–2004. *Public Choice*, 135, 337–352.
- Hall, P., & Soskice, D. (2001). An introduction to varieties of capitalism. In Hall, P. & Soskice, D. (Eds.), *Varieties of capitalism* pp. 1–68. Oxford: Oxford University Press.
- Heckelman, J. (2002). Variable rational partisan business cycles: theory and some evidence. *Canadian Journal of Economics*, 35, 568–585.
- Hibbs, D. (1977). Political parties and macroeconomic policy. *American Political Science Review*, 71, 1467–1487.
- Hibbs, D. (1987). *The political economy of industrial democracies*. Cambridge: Harvard University Press.
- Hibbs, D. (1992). Partisan theory after fifteen years. *European Journal of Political Economy*, 8, 361–373.
- Mankiw, N. G., & Reis, R. (2002). Sticky information: a model of monetary non-neutrality and structural slumps. *NBER Working Paper 8614*.
- Mankiw, N. G., Reis, R., & Wolfers, J. (2003). Disagreement about inflation expectations. *NBER Working Paper No. 9796*.

- North, D. (1990). *Institutions, institutional change, and economic performance*. Cambridge: Cambridge University Press.
- Rogoff, K. (1985). The optimal commitment to an intermediate monetary target. *Quarterly Journal of Economics*, *100*, 1169–90.
- Williamson, O. (1985). *The economic institutions of capitalism*. New York: Free Press.
- Wlezien, C., & Erikson, R. (2002). The timeline of presidential election campaigns. *Journal of Politics*, *64*, 969–993.
- Woodford, M. (2002). Inflation stabilization and welfare. *Contributions to Macroeconomics*, *2*(1), article 1.