FLOATING VOTERS IN U.S. PRESIDENTIAL ELECTIONS, 1948-2000*

John Zaller
UCLA

*I thank Willem Saris, Barbara Geddes, Mark Hunt, Larry Bartels, and Matt Baum for many helpful comments on various drafts of this paper. Advice and raised eyebrows from Jeff Lewis were also helpful. I especially thank the editors of this volume, Paul Sniderman and Willem Saris, for their infinite patience with me in the preparation of this paper. It made a huge difference.

Poorly informed voters are more likely to shift back and forth between the parties in presidential campaigns. They are also less likely to cast votes that can be explained in terms of their self-described policy positions. These findings, along with side evidence of "non-attitudes" and "ideological innocence" among the less informed, feed the impression that they cast their ballots haphazardly. They tend to be, as early voting studies put it, "floaters" (Daudt, 1961). Less kindly, they are sometimes called know-nothing voters.

This paper focuses on these low information voters in U.S. presidential elections. Contrary to the widespread view, it finds that they are often more responsive to the political content of elections than their better-informed counterparts. Low information voters are more apt to reward incumbents who preside over strong national economies and punish those who do not. Poorly informed voters are also more responsive to the ideological positions of the candidates. And finally, low information voters are at least as likely as other voters, and perhaps more so, to respond to presidential success or failure in foreign affairs.

High information voters tend, in contrast, to vote loyally for their customary party. Candidates, issues, and national conditions may vary greatly, but high information voters do not. The stable partisanship of the highly informed is the foil against which we observe the greater sensitivity of less informed voters to the political content of elections.

Reaching this conclusion requires patient attention to problems of measurement and model specification. One big problem is that analysis must rely on survey measures of presidential vote choice that are, as we shall see, biased and hence in need of adjustment. Another is that there are multiple ways of measuring key variables in the analysis, some of which produce different results than others. Measurement of candidate ideology relative to the ideological position of the median voter is particularly difficult. Neither existing theory nor research convention provides firm guidance on how to deal with these problems. On top of all this, we have only 14 presidential elections for which good survey data are available, which makes it hard to tell which model specifications are capturing something real and which are merely capitalizing on chance.

Until recently, the standard approach to these problems has been to find a robust specification and report it as a "best model." It is increasingly recognized, however, that the best model tradition makes for intellectual sandcastles, unable to withstand the next tide of new evidence (Geddes, 2002; Bartels, 1996). The alternative, as I try to practice it in this paper, is to work through results from a variety of models, hoping for stable patterns across specification but willing to report such instability as turns up.

The paper has six primary sections. The first identifies a set of national-level forces likely to create swings in the presidential vote. The second section identifies plausible measures of these forces, including alternative measures. The third section evaluates the alternative measures. The fourth section presents the main empirical findings and tests their robustness. A brief fifth section evaluates the paper's findings in light of continuing disagreements in the political behavior literature about voter rationality. A sixth section discusses why past research has so long failed to detect the responsiveness of poorly informed voters to important dynamic elements in presidential contests -- and what methodological lessons should be drawn from this failure.

I. THEORETICAL BACKGROUND

Information and political behavior

Understanding the voting behavior of low information voters has both a scientific and a normative motivation.

The normative motivation is obvious. Countless studies attest that the overall level of political information in the electorate is low. When, on top of this, election outcomes turn disproportionately on swings among the least informed voters, one is naturally concerned to know what principles animate the uninformed and what their behavior bodes for the political system as a whole. Obtaining an accurate assessment of how low information voters behave takes up the bulk of this paper, but I also consider the broader question of what animates low information voters in section V.

As regards scientific motivation, many political behavior studies continue to ignore informational differences among citizens and their effects on political action. When these differences are noticed, they are often accommodated by means of a simple direct effect term that cannot capture interaction effects -- even though the interactive effects of information are often more important than direct effects (Stimson, 1975; Sniderman, 1993, Sniderman, Brody, and Tetlock, 1991; Delli Carpini and Keeter, 1995; Zaller, 1996). The consequence of neglecting information effects can be incorrect models and misleading accounts of the dynamics of mass political behavior.

Political information is not the only factor that might be examined in a study like this one. Education, political interest, participation in politics, and media exposure belong to the same family of "political involvement" variables and produce some of the same effects. But research has shown that information often generates the biggest effects on political behavior (Luskin, 1987; Zaller, 1990; Price and Zaller, 1992). And, from a normative perspective, we are usually more interested in what a citizen actually knows about politics than in his education, emotional involvement, or level of media consumption.

<u>Determinants of presidential voting</u>

The flow of votes between presidential elections depends greatly on the tendency of voters to engage in retrospective voting – the tendency, that is, to look back over the incumbent party's performance in office and vote accordingly (Downs, 1957; Key, 1966; Kramer, 1971; Fiorina, 1981; Rosenstone, 1983). There is wide agreement that stewardship of the economy is an extremely important dimension of incumbent party performance, at least in the United States (Lewis-Beck, 1990; Bartels and Zaller, 201; Wlezien, 2001).

But what general considerations, besides economic performance, exert an important influence on presidential elections? This is a daunting problem because there are numerous possible influences and few elections in which to evaluate their effects. Consider the following list of factors that might plausibly affect election outcomes: Party performance on race, foreign policy, or social welfare; differences in campaign spending; factional splits within either party;

how long the incumbent party has been in power; balancing party control of Congress; the presence of third parties; the ideological location of the party nominees; political or personal scandal; the degree to which media coverage favors one candidate or the other; whether the incumbent president has had a challenge for re-nomination.

It seems likely that, with enough elections to study, most or perhaps all of these factors could be shown to have some real effect in at least some elections. But with only a handful of presidential elections to use as evidence, it is clearly impossible to test the effects of every factor that may plausibly -- and probably actually does -- have some degree of influence.

Of the many possible factors, I propose to focus on the four for which, as it seems to me, the theoretical and empirical warrants are especially strong. Besides economic performance, they are handling of foreign policy, the general ideological location of the candidates, and the length of time a party has controlled the White House.

An incumbent administration's management of foreign policy is the natural complement of its management of the economy – the other major domain in which voters may wish to play the role of "rational god of vengeance and reward," As V. O. Key Jr. put it. In this regard, Doug Hibbs (1987, 2000) has provided strong evidence that unsuccessful war affects presidential voting. There is also evidence that foreign policy actions short of war, such as President Jimmy Carter's handling of the Iran hostage crisis, can affect voting behavior (Iyengar and Kinder, 1987; see more generally, Aldrich, Sullivan, Borgida, 1989).

Voters may not only look backwards at how the incumbent administration has managed the country, but forward at what the candidates promise to do if elected. In a cross-election study of this kind, these promises can be summarized in terms of a candidate's general ideological location. For at least five decades, theorists have known that, in a two-party system, the optimal location for candidates on political issues is at the center of the ideological spectrum (Downs, 1957). But because party activists pressure candidates to take non-centrist positions (Aldrich, 1995, chapter 6), candidates sometimes stray from the middle-of-the-road. When they do, they tend to punished by voters (Rosenstone, 1983; MacKuen et al, 2002; see also Wright 1994 and Moon, 2002.[1]

The fourth variable, which might be called “incumbent party fatigue,” attempts to capture a variety of processes that occur when a party has been in the White House one or more terms. Low values often indicate that a political innovator – e.g., someone like Franklin D. Roosevelt or Ronald Reagan – may have swept into office. High values suggest that the incumbent party is getting down to its second string – Bush to succeed Reagan, Gore to succeed Clinton. High values may also indicate an increasingly querulous in-party and an increasingly frustrated out-party, one ready to shed old ways and “do what needs to be done” to win. The nomination of Bill Clinton in 1992 as a “New Democrat”, and of apolitical Dwight Eisenhower in 1952, are examples of out-parties willing to break with old ways in order to win.[2] Ralph Nader’s showing in 2000, compared to that of Pat Buchanan in the same year, suggests the intra-party dynamics that arise when incumbent party fatigue is high.

I propose that this set of four variables captures a broad fraction of what presidential elections are about – retrospective voting on the economy and foreign policy, responsiveness to the center of public opinion, and political innovation versus staleness. Each concept is, moreover, highly variable across cases rather than fixed for most of them (as, for example, “major political scandal” or “civil rights innovations” would be). Such variability is obviously helpful in a small dataset. Other scholars might propose other concepts, but I believe it would be hard to argue that my list of four is a bad one.

Two other factors that I might have examined are the incumbent president’s approval ratings and “macro-partisanship,” or mean level of partisan attachment in the electorate. The difficulty

---

[1] Rosenstone's model of the presidential vote is specified with state-level measures of key variables, including economic performance and ideological location of voters. Hence, although he examines nine elections in the period 1948 to 1980, he has about 9 X 50 observations, which enables him to obtain strong results for his ideology measure. Indeed, Rosenstone obtains results for two separate measures of candidate ideology, racial liberalism and New Deal welfare liberalism.

[2] Many scholars believe that the wide open presidential nomination process precludes parties from acting as rationally as this description indicates. See, however, “Beating Reform: The Resurgence of Parties in Presidential Nominations, 1980 to 2000” by Marty Cohen *et al*.

with both of these factors is that they would require a two-stage analysis – first to show how events affect presidential approval and partisanship among voters at different levels of information, and then to show the effect of approval and partisanship. This would require more effort and space than are feasible here.

II. Measurement of Key Variables

This voting study, then, focuses on handling of foreign policy, economic prosperity, ideological centrism, and Party Fatigue as determinants of the two-party presidential vote, and how, if at all, voter information mediates their impact on the vote. This requires measurement of a total of six variables, all but one (party fatigue) of which presents problems. As we shall now see, some of these problems are extremely thorny.

Measuring the vote: The accuracy of CPS/NES surveys

In this study, as in almost all other studies of the presidential vote over multi-year periods, I confine my analysis to vote shares for the two major parties, the Democrats and Republicans. Third party votes, when they occur, are omitted. This is not because third parties are unimportant, but because responsiveness of low information voters to third party candidates would require major study in its own right .

Since the aim of this paper is to examine the effect of information on voting, it must rely on survey measures of the vote rather than the official vote itself. The survey measures it uses are from the series of presidential election studies that began in 1948 at the Center for Political Studies (CPS) at the University of Michigan and has continued since 1980 as the National Election Studies (NES), which has also been located at the University of Michigan. Each of these election surveys contains a pre-election and a post-election wave and a host of independent variables, including political information.

Like any data, the CPS/NES surveys are prey to both random and systematic error. Hence I open my analysis with the question: How well do the Michigan surveys reflect the actual vote? To answer the question, Figure 1 shows a scatterplot of the incumbent party's actual share of the two-party vote against the CPS/NES estimate. As can be seen, the relationship is strong but by no means perfect. For a regression of survey vote share on actual vote share, one would hope that all points would fall close to a line with an intercept of 0 and a slope of 1.0 – that is, the line y = x + chance error. The actual data depart considerably from this expectation:

$$\text{Survey Vote} = \underset{(5.8)}{-17.7} + \underset{(.11)}{1.34}*\text{Actual Vote} + \text{Chance Error} \qquad \text{adj. } r2 = .92$$

INSERT FIGURE 1 ABOUT HERE

The problem, as an examination of Figure 1 reveals, is a general tendency toward over-report for the winner. Winning candidates tend to fall above the y = x line, which indicates that their share of the survey vote is higher than it should be. And candidates who get less than 50 percent of the actual vote fall below the y = x line. On average, the winning candidate gets about two percentage points more support in the CPS/NES surveys than at the ballot box.

This over-report for the winner is not large as survey error goes. But it is too big to be due to chance, and it is big enough that, if left uncorrected, it would create a bias in the analysis, for this reason: Suppose that economic performance causes the actual vote, and that the actual vote causes, in turn, a small amount of over-report for the winner in surveys. Given this, an incumbent who won because the economy is strong will win by an even greater margin in the reported vote, thereby making the economy seem like it affects the vote more than it does. The same tendency toward overestimates could occur for any variable that affected the actual vote.

Preliminary analysis bore out this fear. In a variety of tests, I found that the estimated effect of economic performance on the incumbent vote was always too large when the CPS/NES measure of the vote was substituted for the actual vote. Estimates of the effect of foreign policy performance and candidate moderation were also consistently too big, sometimes double what they should be.

Can a mere 2 percent over-report for the winner really lead to such distortion? Yes it can. Given that the SD of the incumbent vote is 5.8 percent, a 2 percent over-report is a large fraction of what the aggregate variables are trying to explain. Correction is therefore imperative.

In much research, it is hard to correct for, or sometimes even to detect, survey error. In this case, however, we know from the official voting returns what the overall survey vote ought to be, and knowing this, we know exactly what the error in the survey vote is.

The next step, then, is to correct for the error. To do so, I first calculate the fraction of individual voters who must have over-reported in order to produce the aggregate amount of over-report that exists. Then, working with the 0-1 individual-level data, I subtract that fraction from the value of each reported vote for the winner. Here is an example of the method. Suppose that in a sample of 100 voters, 60 percent report voting for the winner, but that the winner got only 50.01 percent of the actual vote. Given these numbers and ignoring sampling fluctuation, 10 of the 60 who claimed to have voted for the winner must have over-reported. I do not know which particular voters they are, so I assume that each of the winner's supporters has a one-sixth chance of having over-reported. I therefore count each of their votes for the winner as only 5/6ths of a vote. Thus, instead of a vote variable with values of 0 and 1, I have a vote variable with values of 0 and .833. If I adjust each year's individual-level vote variable in this manner, the result is that the mean of the NES vote exactly matches the mean of the actual vote.

When I used the corrected vote variable in analyses of the effects of the economy and other aggregate determinants of the vote, I obtained coefficient estimates that were essentially identical to those obtained from using the actual vote. This was true whether I used the survey vote in the form of 14 aggregate means or individual-level regressions having some 16,000 cases. In other words, the correction worked extremely well.

But, although producing a perfectly calibrated vote measure, this correction procedure has an important shortcoming. It is the assumption that all kinds of voters – including voters who differ by political information -- have been equally likely to over-report. As we shall see below, a test of this assumption makes it appear plausible. Still, in a study focusing on differences between low and high information voters, I do not want the findings to be vulnerable to the suspicion that results may have been driven by differences in over-report rather than differences in voting.

For this reason, I undertake a parallel analysis of vote intention from the pre-election wave of the NES surveys. If the analysis of vote intention in pre-election surveys leads to the same conclusions as analysis of reported vote in post-election surveys, confidence in the conclusions will be greater.

A survey question about vote intention has been carried on the pre-election wave of all 14 of CPS/NES surveys and typically reads as follows:

Who do you think you will vote for in the election for President?

We have seen that the CPS/NES post-election measure of vote contains a significant bias. But what biases, if any, afflict the CPS/NES measure of vote intention? The main concern would be that CPS/NES samples, which are collected over the post-Labor Day campaign, might fail to accurately capture vote intentions at any particular point. As an initial test of accuracy, Figure 2 shows a scatterplot of the mean CPS/NES measure of vote intention against the Gallup organization's early October poll in elections from 1948 to 2000. I shall have more to say about the correspondence of these two sets of data, but for the moment their association is reassuring. Indeed, the scatter of points falls almost exactly on the Y=X line.

INSERT FIGURE 2 ABOUT HERE

But how is intended vote related to actual vote? The statistical relationship between the two measures over the period 1948 to 2000 is:

$$\text{Vote Intention (CPS/NES)} = \underset{(9.2)}{-14.2} + \underset{(.20)}{1.29}\ \text{Actual Vote} \qquad \text{adj. r2} = .77$$

Thus, over-report for the winner seems to be present in the pre-election survey and to about the same degree as in the post-election report of actual vote. Because of this, we might suspect that the over-report derives from a similar psychological tendency in both cases, namely, an impulse to associate oneself with the winner. Yet if this mechanism were at work, it would lead to more over-report for the winner as election day grew closer, since it would be increasingly clear who the winner actually was. Yet this is not what happens. Indeed, it is the opposite. Surveys taken close to the election show essentially no over-report for the winner. Thus, a regression of vote intention *as measured in the last 14 days of the campaign* on actual vote yields the following:

$$\text{Vote Intention (CPS/NES)} = \underset{(11)}{-3.03} + \underset{(.22)}{1.08}\ \text{Actual Vote} \qquad \text{adj. r2} = .66$$

These results, which are close to the y = x line, strongly suggest that a psychological drive to over-report support for the winning candidate is not the source of the problem in the vote intention variable. What, then, is the problem? In other analysis, it turns out that many presidential races tighten at the end. Thus, whatever the winning candidate's margin on election day, his margin was probably greater at an earlier point in the campaign. Because of this, vote intentions averaged over the entire campaign overstate the support that exists on election day. The overstatement is not a bias, but an accurate reflection of vote intentions over the whole the campaign.

In light of this, I would like to be able to focus on vote intention data at very end of the campaign. However, this yields too few cases for analysis of individual-level differences. I must therefore examine vote intentions as they evolve over the course of the campaign. This approach has the side-benefit benefit of enabling us to see how the preferences of low and high information voters are shaped by the campaign.

<u>Level of Political Information</u>

I turn now to a brief discussion of the measurement of political information. The advantage of the CPS/NES election series, compared to other survey data, is that they typically carry items designed as tests of political information. The main exception is the 1952 study, in which I had to build a scale that more nearly resembles an omnibus measure of political involvement than political information. But even the 1952 scale has some information items and, in other years, it is possible to build scales that are mostly or entirely based on tests of political information.[3] In recoding the items and combining them into scales, I followed the general procedures described

[3] The 1956 and 1960 surveys were also scarce on information items, but they were part of a panel study that, in its 1958 wave, carried numerous good items. Hence I merged responses to the 1958 information tests with the 1956 and 1960 studies, building instruments for respondents who did not take part in the 1958 survey.

in Zaller (1992). The SPSS code for constructing these scales in the various election studies and merging them is available on my webpage.[4]

In constructing the information scales, I sought to avoid missing data. Hence, if a respondent failed to answer one or several of the items in an information scale, I assigned that respondent an average score based on the items to which he did respond. If respondents missed many or most items, I imputed information scores based on age, education, other demographics, and other indicators of political involvement, such as interest in politics. The weights in the imputation were determined by OLS regression within the 14 separate election studies, with r-squares in the imputation equations averaging about .50. No attempt was made to adjust standard errors in the main analysis to reflect the uncertainty involved in data imputation over the 14 individual CPS/NES studies. After data imputation, missing data on the information scales averaged only about two tenths of one percent of the respondents who reported voting for one of the major party candidates in a post-election survey.

Several studies indicate that, despite increases in educational attainment, Americans are about as well-informed about politics today as they were in the 1950s (Smith, 1989; Deli Carpini and Keeter, 199s). In light of these studies, I have constructed my political information scales so as to produce a constant mean and variance across time. The usual way of achieving this purpose – conversion to z-scores – results in scales having significant inter-election differences in skews and endpoints. Hence, much of the analysis that follows relies on three-category scales or five-category scales in which cutpoints have been calculated within each separate election study.

Measuring economic performance

The difficulty in measuring economic performance is a true embarrassment of riches – more data and combinations of data than can be meaningfully evaluated on the basis of a mere 14

---

4 The code runs on the 1948_1997 National Election Studies CD-ROM, plus the version of the 2000 study that was on the NES webpage in December, 2001. This code is organized so that Tables 3 and 4 below, which are the heart of the paper, can be replicated with a single SPSS command. My webpage is www.sscnet.ucla.edu/polisci/faculty/zaller/.

data points. Among the primary data series made available by the U.S. Government are inflation, unemployment, real disposable income (RDI), and gross domestic product (GDP). All of these measures, and more, have been tried, but RDI generally works best in the U.S.

But even a decision to focus on RDI alone leaves an open question. RDI can be measured over varying time periods — the incumbent's whole term of office, the year of the election, or just one or two quarters before the election (Lewis-Beck, 1992.

The most elegant approach to the problem of time periods is that of Hibbs (2000), who develops a model that uses economic data from every part of a president's term except the first quarter, when the new leader cannot reasonably be held accountable for performance. The Hibbs model explicitly estimates whether voters place an equal weight on all quarters, place more weight on more recent quarters, or possibly even place all weight on economic performance in the quarter of the election.

In applying his model to data from 1952 to 1996, Hibbs (2000) found that voters place more weight on economic performance in more recent quarters, but that they place some weight on economic performance over the whole term. More specifically, Hibbs found that economic performance in the four quarters immediately prior to the election gets a weight of about 60 percent, compared to a 40 percent weight for performance in the other 11 quarters of the presidential term. In re-estimating Hibbs model on data from 1948 to 2000, I found that the four quarters immediately prior to the election get weights that vary between 50 and 80 percent, depending on which other variables are included in the model.

According to any of these results, it is a mistake to use economic performance data from only a few quarters. But although the Hibbs model avoids the need to make an arbitrary assumption about how voters weight past economic performance, it does so at a cost. The cost is an extra parameter. That is, there is one parameter to capture the <u>impact</u> of economic performance on the vote, and one extra parameter to estimate how voters <u>weight</u> recent compared to temporally distant economic performance. In a dataset having only 14 elections and several independent variables, inclusion of an extra parameter is not trivial.[5] A further problem is that statistical estimates of the impact parameters and the weight parameters are correlated, which complicates comparisons of both direct economic impacts and the impacts of other variables across groups.

I will deal with the difficulties posed by the extra parameter in the Hibbs model by estimating the weight parameter from the data on actual election outcomes and then using that estimate as a fixed parameter in the estimation of group differences in the survey data. In this way, I reap the advantage of Hibbs' weighting routine – a scheme based on a president's entire term rather than an arbitrary and rigid assumption about which quarters matter -- without incurring the cost of an extra parameter in the estimates that are central to my analysis.

Hibbs' approach, although theoretically elegant, has no guaranteed empirical advantage. Hence I test it against two other approaches. The first is economic performance in the 12 months prior to the election (quarters 12 to 15 of the presidential term).[6] My other measure of economic performance comes from a recent paper by Nadeau and Lewis-Beck (2001). They maintain that data on actual economic performance may be less consequential than public *perceptions* of economic performance. And they point out that public perceptions may embody the effects of a great many factors, as they seem relevant to the public at the time, rather than only Real Disposable Income. On this theory, they use the Michigan Survey on Consumer Attitudes and Behavior to develop a measure of these perceptions. This measure is simply the percent who think "business conditions" have gotten better over the last year minus the percent who think they

---

[5] I also worry that taking two parameters to measure the impact of the economy, along with the strong likelihood that economic performance is better measured than ideological location, may have a distorting impact in a small dataset. At the same time, there is some indication in the data that different groups do discount the past to differing extents. In particular, high information voters appear to take a longer view of economic performance than low information voters, who seem to respond most strongly to the immediate past. This is consistent with Krause's (1997) analysis of education differences in perceptions of economic data. Nonetheless, this difference, while probably real as regards presidential voting, did not seem large or stable enough to justify the effort of complicating my analysis to capture it. On balance, then, I am comfortable in fixing the weight parameter at its average population value.

[6] In previous work, I have also used RDI in the 12 months of the calendar year of the election – that is, quarters 13 to 16 of the presidential term. However, that measure is correlated at .95 with RDI in the previous 12 months – that is, quarters 12 to 15. Also, I use the calendar year measure below as an instrument. In these circumstances, carrying calendar year RDI would add more clutter than information to the analysis, so I omit it.

are worse. They call this the National Business Index, or NBI. Nadeau and Lewis-Beck show that, for elections from 1956 to 1996, NBI outperforms election-year RDI in accounting for presidential election outcomes.

The NBI measure, however, has may be endogenous to the election outcome. As the authors point out, it is possible that perceptions of economic performance – which are measured close in time to each election -- are colored by people's vote intentions. For example, people intending to vote for a Republican incumbent may be thereby induced to report that the economy is strong, unconsciously telling themselves, "If I'm voting for this Republican incumbent, the economy must be good." As Duch, Palmer and Anderson (2000) have shown, and as we shall see later in this paper, there is strong reason to believe that this kind of psychological process operates in the context of NES surveys. It would, I assume, be less likely to occur in surveys devoted to consumer affairs, but is still possible. Another problem with the NBI measure is that it is available only for elections from 1956 to 2000.[7] But both difficulties can be solved by building an instrument for economic perceptions and using it to estimate perceptions for the missing years. I have done so[8] and propose to use it along with the other variables I have described.

The NBI measure is scaled in rather awkward units: percent who think the economy is doing better minus the percent who think it is worse. To facilitate interpretation of results and comparison with the other measures, I have converted this measure to the same scale as the simple RDI measure. That is, I have given it the same mean and SD as the simple RDI measure.

This gives us a total of three measures of economic performance that are, as would be expected, strongly correlated. Yet the measures are by no means identical, as shown in the following matrix.[9] The availability of these theoretically equivalent but empirically different indicators affords us leverage in testing the stability of the effects of different versions of other variables (and vice versa).

---

7 I thank Richard Nadeau for making the 2000 value on his measure available to me.
7 I thank Richard Nadeau for making the 2000 value on his measure available to me.

8 As instruments, I use unemployment, inflation, GDP, and RDI in the year of the election. These variables yield a model with an r-square of .88 and an adjusted r-square of .80.

9 The Hibbs measure is based on a weight parameter of .822, as described below.

	Previous 12 Months RDI	Hibbs weighted measure
National Business Conditions (instrument)	0.66	0.80
Hibbs weighted measure	0.85	

One final wrinkle. In the course of gathering the data for this analysis, I came across a complication, namely, the considerable variability that exists in government estimates of economic performance. The Bureau of Economic Analysis (BEA), a subsidiary unit of the Department of Commerce, releases estimates of quarterly performance within a month of so of the end of each quarter, but describes them as preliminary and subject to change. In fact, the estimates sometimes change a fair amount, and they continue to change for years and even decades as better information and analytic techniques become available. For example, preliminary BEA estimates of RDI in the first four quarters prior to each election– i.e., the RDI estimates available on election day – correlate with revised estimates of this same quantity at the level of .87 in the period 1948 to 1996. Squaring this number, one gets an r-square of .76, which implies that 24 percent of the variance in the preliminary measure, as aggregated over four quarters, is error variance. Reliability in more recent years is greater, however.

The discrepancy between the preliminary and revised figures may be politically consequential. For example, it is widely believed that George Bush lost the 1992 election because of the economy's poor performance and that Bill Clinton won re-election in 1996 because the economy was stronger. Statistics released at the time of each election support this view: According to the BEA's preliminary estimates, average change in RDI in the first three quarters of 1992 was 2.37 percent, whereas in Clinton's reelection year of 1996, this figure was somewhat higher at 2.95. Yet in later BEA estimates, the two election years came to be indistinguishable in terms of economic performance. The most recent (but never final) estimates are that RDI was 1.80 percent in the first three quarters of 1992 and 1.79 percent in the first three quarters of 1996.

It is not clear which set of estimates should be used to explain presidential election outcomes. But, as Bartels (1996) has pointed out, this uncertainty is actually an opportunity to discover the mechanism that underlies the well-known effect of the economy on presidential elections. If the preliminary government estimates, which are well publicized prior to the election, have more effect on the vote, it would suggest that the political effects of the economy are mediated by official government reports of it. Thus, it is the “reported economy” that swings elections. But if, on the other hand, later and presumably more accurate estimates of economic performance provide a better statistical explanation for the swings in the presidential vote, it would indicate that voters respond more to the "real economy" than the "reported economy."

In my preliminary analyses, neither the “real” nor the “reported” economy had a consistent advantage over the other in predicting the presidential vote. However, the “real” measure tended to be somewhat stronger. It is also far easier to collect in a reliable manner.[10] Hence, my analysis will rely on the BEA’s adjusted reports of economic performance.

Measuring candidate positions on issues

Measuring the role of issues in presidential elections is the most difficult problem in this paper, as different approaches give quite different results. There are actually two parts to the problem: First, measuring the positions of the candidates, and second, measuring the aggregate position of voters.

The first scholar to use candidate ideology in an aggregate model of presidential elections is Rosenstone (1983). He obtained his measure by asking some 40 scholars of presidential politics

[10] Tedious issues arise in calculation of the “reported” measure. One is that what the BEA reports changes over the period 1948 to 1996. Early in the period, it reports only disposable income, rather than real disposable income. It is possible for me to make the inflation adjustment – indeed, I have done so -- but the question arises: Is the resulting measure still a “reported” measure if I have post-adjusted it? On the other hand, if I fail to post- adjust it, then I would be using different measures in different periods. Another difficulty is that news reports emphasize different measures in different elections, depending on the problems the country faces. It is arguable that a true measure of the “reported” economy ought to take into account differences in what is actually reported from year-to-year. However, this would be a larger task than I can presently undertake.

to rate each major party nominee from 1948 to 1980 on two ideological dimensions, New Deal social welfare liberalism and racial liberalism. As he explains,

> [Raters] were instructed not to judge how the public perceived the candidates, or to recall the results of public opinion polls. Rather, I asked the scholars to score the candidates' actual positions on these dimensions "the way an insightful political observer of the day would have evaluated the actions and positions of the candidate prior to the election." (p. 174)

Rosenstone shows that his ratings are a strong predictor of the presidential vote. The ratings are, however, open to the possibility of bias. Any person knowledgeable enough to rate the ideological location of candidates over a stretch of several decades years also knows the Downsian theory of the median voter.[11] If so, there might be a correlation between moderation and vote share because the expert raters, knowing that the winner "should" be more moderate, rated him so. In light of this concern, I will take care to show that, insofar as I rely on the Rosenstone ratings, this kind of bias is not driving the results.

A plausible source of evaluations for more recent elections is the CPS/NES election series, which asked its respondents from 1972 to 2000 to place both major party candidates on a seven-point scale running from extremely liberal to extremely conservative. These ratings are unlikely to be influenced by the Downsian theory of the median voter. Of course, they might not be much influenced by the candidate's ideological location either. The problem, as Converse pointed out nearly four decades ago, is that the bulk of the electorate lacks the political sophistication to use ideological language. Yet, one can simply restrict attention to survey respondents who score sufficiently high on political information to know what ideology is. Hence, in the analysis that follows, I use candidate location estimates only from NES voters who score one or more SDs above average on information and who rate the Democratic candidate at least one point to the

---

[11] Although Rosenstone uses his measures to forecast elections, all of the data on which I rely involve elections whose outcomes were known at the time Rosenstone collected his ratings.

left of the Republican candidate.[12] This procedure uses roughly the top 15 percent of the electorate as if it were a panel of experts judging the ideological location of candidates.[13]

We have, then, two partially overlapping measures of candidate location – the Rosenstone expert ratings, which span elections from 1948 to 1980, and the NES ratings, which run from 1972 to 2000. For my analysis, I recoded the 1 to 7 NES measure so that -3 was the most conservative position, +3 was the most liberal position, and 0 was the midpoint. I then spliced the two measures into a single scale, using the three elections (hence six candidates) on which they overlap as a means of scaling the Rosenstone ratings to the NES metric. In the final measure, I use the re-scaled Rosenstone's ratings for the elections from 1948 to 1968, and the CPS/NES ratings from 1972 to 2000. Rosenstone's ratings for the 1972, 1976, and 1980 elections were used only to put the two sets of ratings on a common metric.[14] Figure 3A displays the resulting measure over the period 1948 to 2000.

INSERT FIGURE 3 ABOUT HERE

---

[12] The number of highly informed voters who failed to see the Democratic candidate as one or more scale-points to the left of the Republican was small – about 3 percent of the eligible pool. Their ratings have not been used in this analysis.

[13] I made an extensive effort to apply the Brady-Sniderman (1985) model of group location to data on candidate location. Although this model is elegant, it did not seem to work well on candidate location data. A central assumption in the model is that survey respondents want, all else equal, to make accurate estimates of ideological location. From this it is deduced that political information ought to be associated with convergence on a common rating reflecting the actual location of the object being rated. This did not occur. More informed raters were more divergent in their ratings than less informed ones; moreover, this unexpected tendency persisted even when the model was run only on respondents scoring +1 SD or higher on political information and after additional control variables were added, including an instrument for candidate affect. For a few of the 16 presidential candidates from 1972 to 2000, the expected positive coefficient for political information was obtained, but in most cases it was not, and in an overall test the sign on information was statistically significant but "incorrect."

However, the ratings produced by this method and by the simple method described in the text are quite similar; After transforming the candidate location estimates into estimates of relative extremism, the correlation between results from the method described in the text and the Brady-Sniderman results was .99. In the end, I decided to forgo systematic comparison of the performance of the two measures because they are so similar and to rely on the method described in the text because it is simpler to calculate and therefore to replicate.

[14] To accomplish the splicing, I collapsed left-right locations into extremity locations by folding the scales at their midpoint and taking absolute values. The two sets of ratings – but only six common points! – were then regressed on one another. The resulting coefficients were used to scale the two series to a common metric and, as the final step, the Republican scores were multiplied by -1 to restore the ratings to left-right space.

Independent assessment of the validity of these ratings is available. Of the 28 major party presidential candidates from 1948 to 2000, 13 served in the House or Senate before running for high office. Keith Poole has kindly supplied me with a measure of the ideological coloration of the voting records of these 13 candidates in the form of scores on his Common Space scale of ideological location (Poole, 1998). The scores are on a metric from +1 to -1, where +1 is most conservative. For purposes of comparison, I converted the Common Space scores to measures of extremity by taking their absolute value and correlating them with the absolute value of the Rosenstone/NES rankings of these same candidates. The correlation was 0.76.

Poole also supplied measures of the ideological coloration of eight Presidents who ran for a second term, where the measure was based on the ideological location of CQ Presidential Support Roll Calls – that is, the roll calls most important to the president's legislative program. Again, I converted these left-right scores into absolute values and correlated them with the absolute values of the Rosenstone/NES rankings of the same candidates at the time they ran for re-election. The correlation was .38. This low correlation is entirely due to Jimmy Carter, whose presidential program was very liberal. If Carter's 1980 score is eliminated, the correlation between the NES/Rosenstone ratings and those of Poole rises to .89.

For another validity check, Figure 3A also gives data on the ideological thrust of the platforms of the two parties in this period, as measured by a group called the "Comparative Manifesto Project" and reported in *Mapping Policy Preferences: Estimates for Parties, Electors, and Governments, 1945-1998* (Budge *et al*.).[16] The measure is created by subtracting the percent liberal sentences in a party's platform from the percent conservative. Separate calculations are made for the Democratic and Republican platforms.

Notably, the platform measure does not attempt to code the substance of individual platform statements. Thus, a platform sentence calling for the abolition of private property would count the same as a sentence calling for more spending on public education. This is surely a weak point in

[16] I received the 2000 data courtesy of Robert Erikson.

the measure. However, it is likely that a platform calling for abolition of private property would call for many other leftist programs and few if any conservative ones, so that, in the end, rough justice could be done to the party's overall position.

A couple of points in the platform data should be noted. One is that, as in Poole's data, Carter in 1980 stands out as a staunch liberal. In fact, if we evaluate Carter in terms of the number of liberal and conservative sentences in his party platform, he appears as the second-most liberal Democrat to win his party's nomination in the whole post-war period, with a more liberal ranking even than George McGovern

Both parties also appear, from the platform data but not the candidate rating data, to make major excursions into enemy territory. In 1956 and 1960, the Republican party presented itself as liberal, and in the 1990s, the Democrats wore conservative clothing. Indeed, in 1992, the Democrats were more conservative than Republicans had been in the 1970s and almost as conservative as the Goldwater Republicans of 1964. This seems to me an exaggeration of Clinton's "New Democrat" thrust to the center.

In general, one can see from a glance at Figure 3 that the data on candidate positions and on party platform positions do not fit well together. The ideological positions of the Democratic candidates are only slightly correlated with the positions of their party platforms, r = .01. On the Republican side, however, the correlation is positive and moderately large, +.69. Thus, Republican candidates seem to have roughly the same ideological positions as their party platforms, but Democratic candidates do not.

It seems to me that, faced as we are with two quite different measures of an important concept, the appropriate methodological choice is to use both of them and see how they work, and this is what I shall do. I will refer to the measures as Candidate Moderation and Party Moderation.

I turn now to the second part of the problem of assessing the role of issues in presidential elections: How to measure the aggregate position of voters. Recall that, in the theory of the median voter, the candidate or party who is closer to the median of public opinion is expected to get the most votes. So it is the midpoint of public opinion that we must measure. This is a vexing problem to which I devote much space below. For introductory purposes, however, let us assume that the center of public opinion corresponds to the natural midpoints of the Party Moderation and Candidate Moderation scales, which are 0 in each case. Given this assumption, we can calculate which candidate is closer to the center by simply summing their scores, as shown in Figure 3B. Seeing how this is works is essential to much that follows, so the reader should attend carefully to the next few sentences. Look first at the candidate locations of the two candidates in 1964. These are in Figure 3A. The Republican (Goldwater) is rated as about -2.1 and the Democrat (Johnson) is about +1.5. Hence the sum is .60, meaning that Johnson is about .60 points closer to the center than Goldwater. Accordingly, the number plotted in Figure 3B is +.60. Given the scoring in Figure 3B, all such positive numbers indicate that the Democrat is closer to the zero-point and negative numbers indicate that the Republican is closer to the zero-point. Thus, to take another example, the negative value for 1972 indicates that Republicans are rated as closer to the zero-point of the scale and, by provisional assumption, closer to the center of public opinion as well. The theoretical expectation, to be tested below, is that the candidate or party that is closer to the midpoint of the electorate will have an electorate advantage over an opponent who is further away. Positive scores on moderation are expected to increase the Democratic vote share and negative ones are expected to increase the Republican vote share. Or, when I re-scale all variables to reflect vote for the incumbent, high scores on the moderation variables are expected to increase the incumbent party's vote share.

With this in mind, let us look carefully at Figure 3B. Even granting that we are not yet sure that the center of public opinion is the 0-point on these scales, some of the party ratings continue to look odd. In 1992, Bill Clinton has a huge moderation advantage over George Bush, fully as large as Johnson's advantage over Goldwater in 1964 and much bigger than Nixon's advantage over McGovern in 1972. Overall, the net Candidate Moderation scores in Figure 3B are correlated with the net Party Moderation scores at +.44, which is almost statistically significant on a one-tailed test but not impressively large.

Yet both scores may be capturing something that is real. Consider, for example, the Democratic party's extreme liberalism score in 1980. Carter may have projected a personal image of moderation, but he struggled with Senator Edward Kennedy for the nomination in 1980 and used the party platform as a means of assuaging Kennedy's liberal supporters. Hence the Democratic platform in that year really was quite liberal. The 1956 and 1960 elections were the high tide of "me too" Republicanism even though the GOP chose the coldest of Cold Warriors to head its ticket in 1960, Richard Nixon. And Clinton was, by some accounts, a dyed-in-the-wool liberal who nonetheless knew the value of appealing to the ideological center. Perhaps the platform data are registering these facts. It then becomes an empirical question to see whether voters put more weight on the party measure or the candidate measure.

All of this assumes, however, that public opinion is somewhere around the zero point of these scales. Figure 4 presents data bearing on this question. The shorter data series is the mean of voters' self-placements from 1972 to 2000 on the same ideology scale that has been used to locate the positions of the candidates in these years. The difference is that the ratings in Figure 4 are averaged over all self-reported voters who responded to the question, whether or not they score high on political sophistication. The second series shows James Stimson's "Policy Mood" measure, which is scored in the direction of liberalism. Stimson's measure summarizes the overall movement of public opinion on domestic political issues – abortion, civil rights, death penalty, welfare policy, and every other issue that pollsters ask about – as reflected in questions that have been asked often enough to permit discernment of trends.

Again, the two data series do not fit together well. Since the NES ideology measure presumably gives citizens' summaries of their views on the same issues that Stimson's measure summarizes, changes in the two ought to be correlated. They are, but not strongly. The correlation of data in this figure is .38. Moreover, the NES measure places the center of public opinion somewhat right of center, whereas the Stimson measure puts it left of center.

INSERT FIGURE 4 ABOUT HERE

And here is another disconnect: If the opinion data in Figure 4 capture electorally significant trends in public opinion, we might expect to find reasonably large positive correlations between changes in public opinion and changes in the positions of ambitious candidates and parties. Yet we do not. The largest correlation in absolute terms is “wrongly” signed – candidates become more liberal as the public becomes more conservative – and the rest are small. None is statistically significant.

No doubt some readers are a bit frustrated by these contradictory data. The effort to relate candidate and party moderation to the center of public opinion is not proving easy. Yet we are closer to a solution than it may seem. Let us reflect on what we are seeing.

As regards the disconnected trends in public opinion in Figure 4, what seems to happen is that as the public becomes, for example, more conservative on issues like welfare or defense spending, citizens do not all begin calling themselves conservatives. Rather, they move with the times on particular issues without changing their expressed ideologies. Although I had not anticipated this disconnect, I probably should have. Consider the following thought experiment: Suppose we had a continuous measure of the public’s self-location on an abstract measure of racial liberalism over the last 100 years. As we know, the public has become dramatically more liberal in the concrete positions it takes on race issues (Schuman, Bobo, Steehl, 1997). Yet, we would almost certainly not find citizens rating themselves as increasingly racially liberal over the whole period and “extremely liberal” on race today. This is because they would not, as an analyst might wish them to do, take the beginning of the data series in 1900 as the baseline for their evaluations. Rather, they would evaluate themselves in relation to the current policy status quo, which has become more liberal on race. Thus, if the public’s self-locations on general racial liberalism changed at all, they would probably shift to the right over time, as some citizens became satisfied with the new status quo and others opposed further change.

Wlezien (1995) has found that exactly this sort of adjustment process operates strongly in the domain of public preferences for government spending. Across five policy domains, aggregate changes in public preferences for spending are negatively related to spending decisions, such that the public adjusts its preferences for spending downward (upward) when appropriations increase (decrease).[17]

The upshot of this analysis is it would be difficult to specify a priori how temporal changes in the public's self-rating on ideology ought to vary in relation to temporal changes on concrete issues. The two kinds of attitudes could move in tandem, but it is easy to imagine that changes would be temporally uncorrelated or even negatively correlated as people adjust their positions in light of changes in the real world.

It is, moreover, likely that the same adjustment process operates on candidates as they take positions on issues. Suppose that public opinion on social welfare issues moves to the right over some period of time. We would then expect nationally ambitious politicians to move to the right as well. The liberal politician who initially favored more welfare spending would now perhaps favor a "hard look" at existing programs and perhaps a few cuts. The conservative who initially favored cuts would now want deeper cuts and an end to federal involvement. Thus, the liberal, although forced to move right, would remain somewhat to the left of general sentiment; the conservative, who has had the good fortune of seeing the public move in his direction, would likewise retain a stance somewhat to the right of the general consensus. Each moves in an absolute sense while maintaining his position in a relative sense.

[17] A further complication is that changes in the public's ideological positions are likely to be extremely prone to measurement error – much more so than, for example, the attitudes on specific spending questions that Wlezein examined. Thus, Box-Steffensmeirer and De Boef (2001) find in a recent paper on time trends in macropartisanship and macroideology find that, for the large majority of citizens, these trends are almost unrelated. As they comment: Citizens whose ideology and partisanship are tightly linked are a relatively small and generally stable segment of the electorate… The [much larger] inattentive public shows much wider gaps between ideology and partisanship series and no relationship between the two (p. 245).

Along with many other scholars, Box-Steffensmeirer and De Boef also note that Americans tend to embrace conservative values even as they regularly support liberal policies and vote for Democratic politicians. Americans are, as Free and Cantril (1967) put it, philosophical conservatives and operational liberals. Hence it would be dubious to accept the mean level of mass sentiment on the ideology scale, which can be seen in Figure 4 to be somewhat right of center, as an indication of its true ideological location or how it votes in elections.

Party platforms might be equally responsive to a change in the public's welfare attitudes and yet equally stable in their relative positions. The Democratic party platform might embrace the need to "end welfare as we know it," but nonetheless propose a flurry of liberal amendments to the new conservative status quo. The Republican platform might likewise embrace the new reality but seek more change. Again, relative ideological divergence might remain stable even as actual policy positions change.

To summarize: One cannot make an estimate of the effect of candidate moderation on vote share without saying something about where the center of opinion is. This is why, in discussing Figure 3, I made the provisional assumption that the zero-point on the ideology scales represented the true center of mass opinion. I now suggest, though still on a provisional basis and subject to later testing, that this assumption may be valid. I suggest, that is, that politicians ambitious for national office stay more or less centered on public opinion most of the time, changing on particular issues but not on the degree they are left or right of center. To be sure, an extremist candidate may happen to capture a party nomination one year and a moderate the next. Or party platform writers may experiment with different techniques of appealing to the public. Or an important group of politicians may try to break out of the pack with the claim that it understands better than other politicians what the public really wants.[18] But such deviations from the overall center of public opinion may occur, and may even turn out to be electorally consequential, without invalidating the idea that, in the main, national politicians continuously recalibrate so as to stay close to public opinion.

If all this were true, the zero-point of the left-right scale in the NES survey, or in the scale used by Rosenstone's expert raters, could be the <u>moving center</u> of a measure of <u>relative</u> left-rightness. As I show in the next section, the data on presidential voting are most consistent with this view. Yet I do not force it on the data. Rather, I permit it to emerge from the most plausible empirical tests I can devise.

---

[18] The Republican party's "Contract with America" in 1994 may be an example of this.

Measuring foreign policy performance

The leading model of U.S. presidential elections is what Doug Hibbs calls a "Bread and Peace" model. Peace – or, actually, war – in the Hibbs model is measured in terms of battlefield deaths. In practice, however, the only important wars in the period of his study are Korea and Vietnam, and they produced nearly equivalent numbers of deaths by the time of the 1952 and 1968 elections. Hence, the war variable is not much more than a dummy variable with two cases scored 1 and the rest 0.

But are the Vietnam and Korea the only foreign policy events that affected presidential elections? It seems unlikely. The notion that foreign policy affects public opinion and vice versa is pervasive in the modern study of international relations (Aldrich et al., 1989; Fearon, 1994; Holsti, 1996; Schultz, 1998; Sobel, 2001).

In an effort to measure a broader range of foreign policy influences, I asked a group of foreign policy scholars to rate the election-year success of presidents in advancing American interests in the international sphere. More specifically, I asked the scholars to:

> evaluate success or failure of the administration's conduct of foreign policy in defending or advancing American interests, as seen from the perspective of a person who follows events in the newspapers and judges them in light of mainstream values of his own day.

Raters were to make their evaluations on a numbered seven-point scale, ranging from "highly unsuccessful" (-3) to "highly successful" (+3). The zero point was not labeled. My instructions focused on the year of the election because the Hibbs measure of economic performance indicates that voters give most weight to events close to the election. The exactly language is included as Appendix A.

A total of nine sets of rankings were obtained.[19] Their rankings are displayed graphically in Figure 5, using the original scale units of +3 to -3. The solid line shows the mean ranking and the dotted lines show +/- 2 SEs from the mean, as calculated separately at each point.

---

[19] Nine of the ten scholars I contacted provided ratings, with the tenth demurred from lack of knowledge of some cases. I thank Richard Eichenberg of Tufts University, Ole Hosti of Duke

INSERT FIGURE 5 ABOUT HERE

There are two points to notice in Figure 5. The most important is that, in the minds of foreign policy scholars, post-war presidential performance was notable for much more than just the Korean and Vietnam wars. Moreover, they achieved a great deal of consensus on the following cases that do no involve war:

- Truman, with a mean rating of +2.1 in 1948, was an extremely effective election-year foreign policy president. The raters did not reveal the reasons for their judgments, but they surely reflected passage of the Marshall Plan by Congress in April of the election year and the Berlin Airlift, which occurred during the fall campaign.
- Richard Nixon in 1972 was also an effective foreign policy leader in the year of his re-election . No doubt this judgment reflected his famous trip to China in February of that year and peace negotiations with North Vietnam, which came to a head during the campaign and seemed on verge of success.
- Jimmy Carter had a very bad foreign policy year in 1980. Surely this judgment reflected the failure of his military attempt to free the hostages seized by Iran and the collapse of negotiations to free them on the weekend prior to the election. It may also have reflected Carter's unpopular decision to enforce an American boycott of the Olympics, protested by American athletes, following the Soviet invasion of Afghanistan.
- George Bush the elder, who prided himself on being an effective foreign policy leader, was actually quite successful. No doubt this judgment reflected the collapse of the Soviet Union and its reconstitution, with American encouragement, as a nominal democracy at the end of his administration. It may have also reflected American success in the Gulf War, the aftermath of which was often in the headlines in 1992.

The other point to notice in Figure 5 is the surprisingly high mean rating assigned to 1952. Seven of the raters gave Truman low scores, but one gave him a +1 and the other a +2. When I asked the latter why, he emailed back: "by 1952 most of the stupid decisions vis a vis Korea were past, and that year the U.S. signed a fair number of formal alliance treaties, including a treaty with

---

University, Rick Sobel of Harvard University, and from UCLA: Matt Baum, Russ Burgos, Debbie Larsen, Marc Trachtenberg, Ken Schultz and Art Stein.

Japan." This comment points up what I now feel is a mistake in my instructions to the raters. In them, I set the time frame as follows: "Evaluations should be based on events in the 12 months or so preceding each presidential election … and matters occurring in closer proximity to election day should be given somewhat more weight." I now feel I should have used language that set a somewhat broader time frame.

Despite some cases of significant rater disagreement, a scale based on a simple average of the eight sets of ratings has an alpha reliability of .92, which is a relatively high value. The first factor of a principal component analysis of the items, which could also be used as a scale, explains 64 percent of the variance. In the analysis that follows, I will use the additive scale. This choice, in allowing possible errors and idiosyncrasy their full weight, creates a conservative bias in the scale.

There is, in fact, reason to be conservative in working with this scale. As with the expert rankings of candidate location, we must be concerned that the experts may know which presidents were popular and shade their ratings accordingly. In this regard, I note that the zero-order correlation between these rankings and the election results is not particularly high at $r = .45$ ($p = .05$, one-tailed). By way of comparison, the correlation between change in Real Disposable Income and election results is .73; between Candidate Moderation and the vote the correlation is .23.

To protect against the possibility of bias, I created a "Consensus Scale." It assigns values only to the six cases which the coders all or nearly all agreed were important. On the plus side, these were 1948, 1972, and 1992; on the negative side, they were 1952, 1968 and 1980. To avoid unconscious shading, the positive cases were all given scores of +1, the negative cases were all given scores of -1, and the rest were given 0.

The Consensus scale is crude in the sense that it rates the Vietnam War, in which 35,000 Americans died by the time of the 1968 election, as equal in gravity to several events that involved little or no bloodshed. But it is a good variable in the sense that it has a strong basis in expert judgments and yet is largely free of concern that it has been shaded in the direction of presidents who simply happened to be electorally popular. The Consensus scale has a zero-order correlation with the first Expert scale of .91 and a correlation with the incumbent vote of .43.

III. Validity Checks of Theory and Measures

This section undertakes preliminary analyses to validate the general theoretical ideas and resolve the measurement issues described above, beginning with an assessment of the effects of voter differences in information on swings in presidential voting.

Are low information voters really different?

Since Converse (1962) and Dryer (1971), it has been known that preference change within election campaigns tends to be concentrated among the less politically interested and informed strata of the electorate. Immediately below is an update of those results from elections from 1948 to 2000.

	Low Information		Middle Information		High Information
Percent unstable from pre-election to post-election	26.5%	20.6%	16.6%	14.5%	11.8%

The measure of instability in these data is a straightforward one: Whether respondents voting in the election report the same preference in both the pre- and post-election waves of the CPS/NES surveys. Movements from undecided to a vote decision, or from one candidate to the other, count equally as instances of instability. As it happens, most of this instability is due to movement from undecided to a vote choice. If we limit the analysis to voters who state an initial preference and subsequently change, the change rates are:

	Low Information		Middle Information		High Information
Percent change from pre-election to post-election	8.8%	6.5%	5.8%	3.6%	5.6%

Thus, two thirds or more of the change that occurs over the course of the campaign consists of forming an initial preference, and about one third consists of changing one's initial preference.

But these data do not address an arguably more important type of change, namely, party change between elections. It is consistent with these data, though perhaps unlikely, that within campaign change is due to voters returning home to the party they supported in the last election. Berelson et al. (1954) emphasized this sort of change in their study of the 1948 election.

Figure 6 offers two ways of assessing between election change. The graph on the left shows percent support for the Democratic candidate by voters according to levels of political information, with a separate line for each of the 14 elections from 1948 to 2000. The lines themselves were generated by means of bivariate regressions of Democratic vote on political information within each election study. The numbers below the graph show the standard deviation of the inter-election vote swing at each level of a five-point measure of political information. As can be seen both visually and by the statistical descriptions, there is far more inter-election variability among low information voters than among high information ones.

A difficulty with this approach, however, is that it fails to distinguish secular change from wild swings. If, for example, low information voters were migrating steadily from one party to the other over a series of elections, we would be mistaken to describe the movement in terms of "instability" or "floating." Only if voters swing back and forth would we describe the change in such terms.

INSERT FIGURE 6 ABOUT HERE

In light of this ambiguity, the graph at the right shows swings between pairs of adjacent elections, with each trend line formed by subtracting Democratic support in one election from Democratic support in the last. Although all 14 elections are involved in these calculations, there are only 13 lines, since there was no baseline from which to calculate swings in the earliest election study in 1948. Again, the standard deviations of the inter-election swings are shown at the bottom of the graph. The SDs show, as previously, that low information voters are far more volatile than high information ones.

These data establish that low information voters contribute disproportionately to inter-election swings, as specified by Daudt's (1961) floating voter hypothesis. What remains is to find out whether they swing erratically, or whether there is a political basis for their movements.

It is appropriate to use one-tailed tests of statistical significance when one has clearly defined expectations based on theory or prior research. Such expectations exist as regards the direct effects of economic performance, foreign policy performance, candidate moderation, and party fatigue. However, I have no basis for expectations with respect to how information might interact with these factors. Hence, all tests of statistical significance involving interactions with political information will be reported as two-tailed tests.

Testing the independent variables

Prior research has identified four theoretically important causes of swings in the presidential vote – economic performance, foreign policy performance, policy responsiveness, and party fatigue. I take the importance of these causal forces as given and ask in this section which measures of them work best. I also attempt to settle two other preliminary issues: What value should we use as the center-point of the electorate for the purpose of assessing the effects of Candidate Moderation? Are expert ratings of candidate location safe from spurious attribution of moderation to the winner?

INSERT TABLE 1 ABOUT HERE

Table 1 starts us off with nine models of the presidential vote to test the various combinations of measures. Through the blizzard of coefficients, we can see that all of the measures work at least some of the time, and that one cluster of measures – involving economic performance – works all of the time. Nonetheless, the analyst looking for "the" right specification is bound to be disappointed. Party Moderation is the stronger measure of policy responsiveness when the Hibbs measure of economic performance is used, but Candidate Moderation dominates when the NBI measure of economic performance is used. Party Fatigue is a very strong performer in several models, but a non-starter in others, with even the direction of the effect ambiguous.

One bit of helpful news is the functioning of the two new measures of foreign policy performance: Both achieve conventional statistical significance in all models, and they always do better than the “wars only” version of foreign policy performance. But stay tuned, there are surprises ahead even for this variable.

As these results show, p-values are imperfect indicators of the reliability of an effect. A variable can be highly statistically significant in several specifications and suffer a sign reversal in the next. The reason for the fallibility is well understood in statistical theory but often forgotten in practice: P-values give the uncertainty of an estimate, conditional on the assumption that the model is true and all measures are the correct ones. Yet in the present problem, and in the large majority of practical problems, the truth about these matters is not known. To rely on any particular set of p-values is therefore misleading.

Bayesian model averaging is one approach to this problem. The researcher tests as many models as he considers plausible and reports the results in the form of weighted averages. The underlying philosophy is that a portfolio of models is better than results from any one model, given that the true model is not known. The weights in the weighted averages are called “Bayes factors” and are intended to reflect the overall goodness of fit of the model to the data.[20] Results from models with high Bayes factors get more weight in the weighted average than results from weaker models. (For a very non-technical introduction to model averaging, see Bartels and Zaller, 2001; for a moderately technical introduction, see Bartels, 1996.)

Bayes weighting is a useful way to summarize the ungainly data in Table 1. Note, for example, that the Bayes factors associated with the NBI measure of economic performance are notably higher than those for the other models. That is, the sum of the Bayes factors associated with three NBI models comes to about 35,000, compared to 7,000 for the simple RDI models and 32,000 for the Hibbs models. This indicates we should put a good deal more weight on results from the NBI model. If we transform these numbers to probabilities – that is, numbers that sum to 1 – the Bayes weights for the NBI, simple RDI, and Hibbs models are .47, .10, and .43.

[20] The Bayes factors resemble the adjusted r-square statistic, except that they impose a much heavier penalty for extra parameters.

These results suggest the superiority for the NBI instrument, but they are not the final word. For one thing, the substantive impacts of the three variables are about the same. The coefficients for the NBI variable are 2.59, 2.05, and 2.14; when weighted by their Bayes factors, these values average 2.07. As mentioned, NBI has been scaled in units of Real Disposable Income, so the estimate of 2.07 means that if RDI goes up by one point, the vote for the incumbent presidential party goes about by about 2.07 points. To judge whether this is a little or a lot, consider that three of the 14 post-war election winners would not have been winners if they lost the benefit of a mere one-half point change in RDI; 7 of the 14 winners would not have been winners if they lost the benefit of a 2 point shift in economic performance. Effect sizes for the simple RDI and Hibbs measures are, despite their lower overall Bayes factors, the same or perhaps larger.[21]

Information not reflected in Bayes factors is also relevant to deciding the importance of a variable. I am, for example, impressed that the adjusted zero-order correlation between the Hibbs measure and incumbent vote is .68, whereas for the NBI measure this correlation is only .45. Thus, the NBI measure attains superior performance only when control variables are present, which suggests weakness.[22] The Hibbs measure also has the theoretically attractive feature of estimating rather than arbitrarily stipulating how voters discount economic performance from early in the president's term. Even the simple 12-month RDI measure has something to recommend it: Its adjusted zero-order correlation with the vote is .70, larger than that for either of the other measures. In light of all these considerations, it seems appropriate to give all three economic variables equal play in the analysis that follows.

The Bayes factors are also useful in assessing the three foreign policy measures. The models containing the "wars only," Expert, and Consensus measures of foreign policy performance yield Bayes probability weights of .02, .89, and .10, respectively. These weights do

---

21 The SDs of the NBI, RDI, and Hibbs measures are 1.4, 1.8, and 1.8.

22 Recall that I am using the NBI measure in the form of an instrument. The adjusted zero-order correlation between the original measure and the vote in the period 1956 to 2000 is .73. I used the adjusted r's to take account of the fact that the Hibbs measure uses two parameters.

not impose a penalty for the possibility of endogeneity bias in the Expert measure, but they do suggest the empirical limitation of the “wars only” measure. The “wars only” measure is also theoretically narrow in that, as we have seen, it fails to take account of several cases of presidential foreign policy performance that are intuitively strong candidates for inclusion in a vote model. For this combination of theoretical and empirical reasons, I shall demote the “wars only” variable to the backburner in this analysis, mentioning results but not attaching much importance to them.[23]

The impacts of the other two foreign policy variables are moderately large. The Bayes-weighted coefficient for the Expert ratings of foreign policy performance is 2.03. This means that every one-point shift in average ratings on the +/-3 scale is associated with a swing of about two percent of the two-party vote. Thus, the difference between a masterful performance like Truman’s in 1948 (average rating = +2.1) and an ineffective one like Carter’s in 1980 (average rating = -1.8) can be calculated as 2.04 X 3.9 = 8. Since Carter got 44.7 percent of the two-party vote in 1980,[24] an additional 8 percent would have been enough to reverse the election outcome. Obviously, no one should be literal-minded about such counterfactual estimates. But it is realistic to believe that the close elections of 1948, 1960, 1968, 1976, and 2000 were within range of foreign policy performance to have made a decisive difference.[25]

---

[23] Because of the scholarly importance of Hibbs’ “Bread and Peace model” – that is, a model containing RDI performance and the “wars only” foreign policy variable – I will provide some statistics on it. For the standard version of the model over the period 1948 to 2000, my estimate of the Bayes factor is about 1320. The Bayes factor for the standard Hibbs model plus Party Fatigue is 888. Thus, adding Party Fatigue worsens Hibbs’ standard model. But if Party Fatigue is added and the Consensus Foreign Policy variable is used instead of “wars only,” the Bayes factor rises to 6292. This odd combination of results appears due to the fact that the Korean War occurs when the term variable has a value of 5 and that the Vietnam War occurs when its value is 2. Given this, the Fatigue variable adds little to the Bread and Peace model when “wars only” is used to measure foreign policy, but a substantial amount when a broader measure of foreign policy is used. The Consensus scale and the Fatigue variable are virtually uncorrelated, but the correlation between “wars only” and Fatigue is .51.

[24] Recall that all calculations in this paper are based on the Democrat-Republican vote. In 1980, Carter got 41.1 percent of the total vote, but after setting aside John Anderson’s third-party vote, Carter’s share of the two-party vote is as 44.7. Winning the two-party vote is normally enough to win the electoral college, as in 1948, 1968, and1992 – but not, of course, 2000.

[25] The mean rating for 1968 on the Expert scale was -2.3, which indicates that Vietnam cost the Democrats about 4.6 points in that election. In 1960 and 1976, the means (scored in the direction of the incumbent party) were -.2 and -.3, which indicate that effects on the two-party vote of the incumbent of -.4 and -.6. Nixon would have won in 1960 with an additional .4 percent

As noted earlier, there is a possibility of bias in the Expert ratings scale. But for the more conservative Consensus scale, the Bayes-weighted effect is 3.12 points per event. This scale assigned scores of -1 to foreign policy disasters (1952, 1968 and 1980), scores of +1 to policy successes (1948, 1972, and 1992), and 0 to other cases. By this measure, then, the difference between Truman's performance in 1948 and Carter's in 1980 is 2 X 3.12 = 6.2, which is in the same ballpark as the estimate from the first measure and also big enough to have reversed the actual election outcome.[26] If anyone wonders whether presidents face electoral incentives in their management of foreign affairs, these estimates should end the puzzlement.

The overall pattern of results also somewhat assuages concern about bias in the Expert measure. The measure would be expected to provide better evidence – and perhaps therefore larger estimates – because it includes much more information of the expert judges than the blunt Consensus measure. We now see that it does provide larger estimates, but not greatly larger. This makes it seem a plausible measure. Still, despite its lower Bayes factors, the Consensus scale is the one to be most relied upon because it is free of the taint of bias.

The "party fatigue" variable also has a large effect, with a Bayes-weighted coefficient of -2.3. This implies that when, for example, a vice-president succeeds a two-term president and runs on his own for the first time – at which time the party fatigue score is 2 – he will do 2 X 2.3 = 4.6 points less well, all else equal. If he nonetheless wins and runs a second time, as did George Bush senior, the expected decrement to his vote share when he runs the second time would be

---

of the vote; Ford need a bit more than the .6 percent that he apparently lost from poor handling of foreign policy. In 2000, the incumbent party's average rating was +.45, which indicates a gain of about 1 percent to Gore. Of course, all of these estimates are highly uncertain in statistical terms, but it is still fair to say that these elections were within the range that foreign policy performance could easily have affected, given the magnitude of the apparent effects.

[26] To test whether the Korean and Vietnam wars had effects above that of other kinds of foreign policy performance, I included both the traditional War variable (1's for 52 and 68, 0's otherwise) and the new foreign policy scales in the standard models. The War variable could then pick up the added effect of war, if any. As it turned out, there was no added effect. The coefficients for the War variable often had the wrong sign (positive) and were never close to statistical significance, while the newer scales retained their effect and were always at or close to statistical significance, despite the multicollinear presence of the War variable in the regression. Thus, the new variables appear to do a good job of capturing the effect of war as well as the effect of other kinds of performance.

6.9 points. The estimate further implies that Stevenson lost 11.5 points in 1952 by virtue of seeking a sixth Democratic term.

These estimates could be overly large.[27] But if even roughly accurate, they imply a bias in models that try to estimate a Korean war effect in 1952 without including a control for party fatigue. Moreover, independent reason exists to take the Party Fatigue effect seriously. As discussed earlier, one of the mechanisms implicated in the Party Fatigue effect is strenuous effort by the out-party to do whatever it takes win. Figure 7 provides insight into what hungry out-parties do. It shows, that is, the relationship between "time in the wilderness" – i.e., number of terms since the party last controlled the White House -- and the ideological distance of the out-party nominee from the scale center-point. The correlation between these two variables is r= -.65, p=.01, two-tailed. The correlation with respect to platform extremism is also negative, though only weakly so (r= -.25, p=.38, two-tailed). Meanwhile, long incumbency in the White House has little if any effect on the moderation of the in-party. For the in-party, the correlation between length of incumbency and distance from the center is +.09 (ns), and the correlation with platform statements is r= +.27 (p=.35, two-tailed).

INSERT FIGURE 7 ABOUT HERE

All in all, these results strongly corroborate the decision to include Party Fatigue in the analysis. To leave it out would be to invite bias.

The effect of moderation on vote choice is obscured by two factors. One is multicollinearity between relative Candidate Moderation and Party Fatigue, as just described. The two variables are correlated at r = -.53, p=.06, two-tailed. For small datasets and high r-squares, a correlation

---

[27] To assess the plausibility of the estimates, I estimated a vote model over all U.S. elections from 1900 to 2000, using election-year GNP, War (1918, 1944, 1952, and 1968), an interaction for GNP X year, year, and Party Fatigue. The adjusted r-square on the equation was .46 and the estimated value of Party Fatigue was -2.7, p=.01, one-tailed. If I omit the 1912 election, in which former Republican president Theodore Roosevelt ran as a Third Party candidate in the fifth year of Republican incumbency in the White House, the effect of Party Fatigue falls to -1.2, p=.15, one-tailed. Yet, intra-party splits and squabbles are exactly the sort of "fatigue effects" one expects to observe when a party has been in office a long time, so omitting the 1912 case would seem unfair.

of this magnitude makes stable estimates difficult. Another difficulty is that the two moderation variables, Party Moderation and Candidate Moderation, tend to work against each other – if one is bigger, the other is smaller. But the Bayes factors are again helpful. The Bayes-weighted coefficient for Candidate Moderation across all models is 1.2; for Party Moderation, it is 0.7. The simple unweighted averages are 1.7 and 1.4, but the Candidate Moderation variable does better in the stronger models, which is why it dominates in the comparison of the Bayes weighted coefficients.

The size of the Candidate Moderation effect may be interpreted as follows: If one of the candidates moved one point further from the center in the metric of the NES 7-point scale and the equivalent of one point on the party moderation scale, the expected cost would a vote loss of about 1.9 percentage points. The effect of even a half-point shift of these scales would be enough to reverse the popular vote outcome in 3 of 14 elections, including the most recent one.

Because a roughly one-point difference between the Democratic and Republican candidates is the largest difference that comes up in this data – between Nixon and McGovern in 1972 and between Carter and Reagan in 1980 (see Figure 3) -- this estimate of the moderation effect may seem rather small.[29] Yet, in one important sense, it is the biggest, because it is both big enough to swing many elections and also a factor politicians can largely control. Indeed, one reason the Candidate Moderation effect shows up only weakly is no doubt that politicians do, for the most part, control it – especially, as Figure 7 shows, when they are motivated to do so. Politicians are now so keenly aware of the importance of moderation in presidential elections that the chances for political scientists to observe its effects via a Goldwater- or McGovern-style candidacy will come up only infrequently.

In any case, moderation remains a highly important force in American elections, as we shall see in more detail below.[30] First, however, we must deal with some other specification issues.

---

[29] The SD of the candidate moderation variable is .35.

[30] I should mention that I also tested an alternative, smoothed measure of the Party Moderation variable. As proposed by Erikson, MacKuen, and Stimson (2001), this variable is a weighted moving average of party positions rather than positions in the year of the election. My reconstruction of this variable, however, failed to generate important effects in any of my

Fixing the weight parameter. As explained, the Hibbs variable uses an extra parameter to establish voters' time discount on past economic performance. The Bayes weighting scheme strongly discourages extra parameters – much more strongly so than the adjusted r-square statistic -- so the relatively high Bayes factors for the Hibbs model indicate that the weight parameter is justified. But even so, I do not want to carry this parameter through the rest of the analysis. My concern is that random fluctuation in the estimate of the weight in later models, which will be based on noisy survey data rather than the actual vote, will ripple through other parameter estimates. Since I have a highly reliable measure of the weight coefficient from the actual vote, I will use it as a fixed value in the analysis of survey data that follows. With weighting by the Bayes factors, the value of the weight parameter in Table 1 is .824.[31]

The center of the electorate. It is now time to settle the nagging question of how the midpoint of the electorate should be estimated. Analysis of this problem is best approached by means of a simple regression model in which a "Policy Moderation" variable is used to explain the Democratic party share of the presidential vote, as follows:

Democratic vote = b0 +b1[Democrat's Relative Policy Moderation]

To concretely measure the Democrat's relative policy moderation, we can write:

Democratic vote = b0 +b1 [{Center - Democrat's Ideology} + {Center - Republican's Ideology}]

Let us (unrealistically) suppose that the Democrat's position on the ideology scale is +3, the Republican's ideology is -2, and that the center of public opinion is +1. The Democrat, then, is 2 points from center of opinion and the Republican is 3 points from center, thus leaving the Democrat a net 1 point closer to the center. The Democrat, then, scores higher on Moderation. Given the numbers in the example and the model above, the expected sign on the coefficient on

---

specifications. These authors use the smoothed measure in a specification in which macropartisanship is an independent variable. I would like, for the sake of completeness, to test their full specification in relation mine, but constraints of publication space and especially time prevent me from doing so. As the reader can see, the present analysis is quite complicated as it is.

31 The Hibbs measure is built as a standard variable from quarterly economic data, using the weighting parameter in the formula supplied by Hibbs (2000); once built, it can be used as simply another variable in OLS regressions.

Democratic Moderation would be positive: Greater relative moderation for the Democrat yields more votes for the Democrat.

If the center of the electorate remains fixed, it would be absorbed in the constant. But it is possible that the center of the electorate moves over time.[32] Let us, then, take Stimson's Policy Mood variable as our operational measure of the center of opinion. The idea would be that temporal fluctuations in this measure, as shown in Figure 4, are the moving midpoint of public opinion against which Candidate Moderation must be measured. Since the Mood variable has a mean of about 60 and an SD of about 5, it needs rescaling to articulate with the Moderation variable, which has a theoretical range from -3 to +3. If we accomplish this transformation through a two-parameter linear function and use DI and RI to represent the ideology of the two candidates, our vote model becomes

Democratic vote = b0 +b1[{ b2 + b3*Mood -DI} + { b2 + b3*Mood -RI}]

Democratic vote = b0 +b1*b2 + b1*b3*Mood -b1*DI +b1*b2 + b1*b3*Mood - b1*RI

Democratic vote = (b0 + 2b1*b2) + 2*b1* b3*Mood +b1*(-DI-RI)

Democratic vote = B0 + B1*Mood + b1*(-DI -RI)

This exercise shows that adding the Mood variable to a vote model serves to control for the moving center of public opinion – if, of course, the Mood variable is the measure of the moving center of public opinion that determines which candidates are moderate and which are extreme. Yet when I re-score these variables in terms of incumbent party vote rather than Democratic party vote, and then add an appropriately scored Mood variable to the models in Table 1, its coefficient never approaches statistical significance and doesn't have a consistent sign.[33] When I use the NES ideology variable to center the Candidate Moderation variable, I get similarly mixed and anemic results.[34] Moreover, adjusted r-squares slightly decrease and the Bayes factor

---

[32] When, moreover, I switch to estimation of incumbent vote rather than Democratic vote, failure to include a term for the center of opinion would result in bias unless the midpoint were actually zero.

[33] Since incumbent vote is the dependent variable in Table 1, I multiply the Mood variable by an index that scores Democratic incumbents as +1 and Republican incumbents as -1.

[34] I set the mean of the ideology variable to 0 in from 1948 to 1968, when the Rosenstone scores constitute the Candidate Moderation variable, and to the values shown in Figure 4 (with a

scores, which are sensitive to overfitting, plummet in the presence of an extra variable to measure the center of public opinion.

On the basis of these results, I conclude that it is reasonable to use 0 as the fixed midpoint of the Party Moderation and Candidate Moderation scales. In doing so, I implicitly assume that individuals who rate the candidates on ideological location – whether NES respondents or "experts" – implicitly use the scale as if 0 were the center point of current public opinion. In terms of my earlier example of racial liberalism over the last 100 years, I am assuming that "history" gets quickly absorbed into current ratings – most likely because politicians are continuously alert to changes in opinion and stay roughly centered on it.

<u>Are the Expert Raters Biased?</u> I turn now to the other threat to the validity of the moderation variable. This is the possibility that Rosenstone's expert raters – upon whom I rely for candidate location scores in the period from 1948 to 1968 – are scoring the winning candidate as more centrist because they know who won the election and think that, by the theory of the median voter, the winner "should" be more moderate.

To test for this possibility, I created an interaction term between Candidate Moderation and a dummy variable for the period from 1948 to 1968. I then added these variables to the standard models. A positive sign on the interaction would indicate that Moderation had more effect in the elections from 1948 to 1968, which is when ratings were done by possibly biased experts. As it turned out, however, the coefficients were more often negative than positive and never approached statistical significance. Hence, I conclude that there is no evidence that raters were biased in their assignment of candidate location scores.

---

sign change for changes in party of incumbent) in other years. Rosenstone asked his experts to center their scores on the center of public opinion, so centering is not theoretically indicated for his ratings.

A positive coefficient would suggest that the true mean of the voting electorate is slightly right of center, as the NES ideology variable indicates. A negative sign would indicate that it is actually left of center. In six trials – 3 economic variables X 2 foreign policy trials – it was slightly positive three times and slightly negative three times.

Dynamics of Political Campaigns. We face one final preliminary. As discussed earlier, it was necessary to correct the CPS/NES measure of reported vote because some respondents over-reported for the winner. The correction works well overall, but it is insensitive to the possibility of individual differences in over-report. For this reason, I am conducting a parallel analysis of vote intentions as they evolve during the campaign.

Figure 8 provides a vivid view of this evolution, but it requires a few sentences of explanation. To create the figure, I began with measures of vote intention at three points in the campaign, as estimated from roughly comparable Gallup surveys 1) after the party conventions in each election year, 2) in early October, and 3) over the weekend before the election. I have run each of these "vote intention" measures against economic performance, foreign policy performance, the two moderation variables, and Party Fatigue. The resulting sets of coefficients, along with coefficients for the official vote, are presented graphically in the left panel of Figure 8. The economic performance variable is the instrumented version of NBI and the foreign policy variable is the Expert scale, as indicated by the analysis in Table 1.[35]

Look first at the point at the upper left of the left graph, where the coefficient value is about 12. This means that, at the time of the Gallup post-convention poll, each point of Candidate Moderation is worth 12 points of vote intention for the incumbent party. As indicated, however, the maximum value Candidate Moderation has ever attained is +.94, and its SD is only .47, So an appropriate way to interpret the effect of Candidate Moderation at the start of the fall campaign is to say that a change of 1 SD on this variable is associated with a change of six percentage points in vote intention (.47 X 12 = 5.6). The downward slope of the line indicates that the Moderation effect diminishes over the course of the campaign. By the October poll, the effect has declined to 10 points, and by the weekend before the election, it has fallen to 6 points. Finally, the effect on election day is about 3.7 points. What this shows is that moderate candidates have a big edge at the start of the campaign, but lose most of that advantage by

[35] All forms of these variables show much the same pattern, except for the wars only version of the foreign policy performance, which gives choppy results

election day. These effects, to reiterate, are based on three sets of polls plus the election results over the period 1948 to 2000.[36]

Economic performance, foreign policy performance, and party fatigue show no clear time trend. Party moderation shows a trend toward zero, but since its initial sign is "wrong" – voters appear to reward party extremity even while punishing candidate extremity – it is hard to believe Party Moderation had any real effect in the first place.[37]

The right-hand part of Figure 8 shows these same relationships in the CPS/NES data series. These estimates have been obtained from an individual-level regression containing the same variables as in the analysis of the Gallup data, plus an interaction between each variable and "days before the election" to capture changes during the campaign. That is, each aggregate variable is represented in the regression by a direct effect term plus an interaction with time of interview.[38] Because the 1948 CPS survey contains no day of interview variable, the CPS/NES estimates in Figure 8 span only the period from 1952 to 2000.

As can be seen in the figure, the general pattern of estimates from the individual-level data is quite similar to the pattern from the Gallup data.

INSERT FIGURE 8 ABOUT HERE

Four important conclusions follow from this analysis of vote intention. The first and most important is that the Gallup and CPS/NES data produce generally similar results. This tends to validate the use of CPS/NES data for an analysis of the effect of political information on vote intention.

---

36 For a view that campaigns do not affect campaign preferences, see Erikson and Wlezien (1999), which is based on analysis of the 1999 campaign.

37 I am trying to avoid flooding the reader with statistics, but let me note parenthetically that the correctly signed Candidate Moderation variable easily attains statistical significance in all three Gallup models, whereas wrongly signed Party Moderation does not. Since multicollinearity is most likely a problem here, I re-ran these models without Party Moderation. Candidate Moderation was not much affected, declining from an early peak of 16 to 7 on the weekend before the vote to 2.85 on election day.

38 A direct term for time and a measure of incumbent party identification, described in the main analysis below, are also included. This individual-level model did not contain fixed-effect terms – as may be obvious, since it does contain direct effect terms. The standard errors of these estimates are incorrect, but the estimates themselves are unbiased, and that is sufficient for present purposes. I estimate properly specified fixed-effect models below.

The second point is that, as hinted earlier, Candidate Moderation is a more important factor in U.S. elections than it appears to be in the analysis of the official vote in Table 1. In both the CPS/NES and Gallup data, it has quite large effects early in the campaign and is still important on the eve of the election. This feeds my suspicion that the moderation effect is badly estimated in Table 1 because of multicollinearity. But even if the Table 1 results are accurate, the "diminishing moderation" effect remains a substantively important campaign effect.

The third conclusion is that the Party Moderation variable may not be more important than it appears in Table 1. At no point in the Gallup data, or in the election itself, does Party Moderation have an important effect in the expected direction. Although I suspect party platforms do have some real effect on presidential elections, the effect does not show up with sufficient clarity in the preliminary analyses to justify carrying an additional variable into the analysis of information interactions in the next section. Thus, Party Moderation joins the "wars only" variant of foreign policy performance as a backburner measurement of an important concept. Results for it will be reported but not emphasized.[39]

The final point concerns the effect of presidential campaigns more generally: It is often argued that campaigns affect voters by "educating" them about the state of the country and the positions of the candidates (Gelman and King, 1993). However, no important educational effect is apparent in Figure 8. The fall campaign begins with presidential vote intentions strongly affected by economic performance, foreign policy performance, and the positions of the candidates. The campaign then fails to increase sensitivity to economic conditions, and appears to actually decrease sensitivity to moderation.

Recall, however, that most campaign change consists of change from undecided to support for one of the candidates. Thus it is likely that much of the diminishing moderation effect consists of voters gradually "coming home" to their customary party despite its candidate's lack of moderation. If so, the supposed educational effect of political campaigns could exist, but get

[39] The pattern in Figure 8 looks very similar if the Hibbs measure or the RDI measure is used instead of the NBI measure of economic performance: Candidate Moderation has a big effect early in the race and diminishes over time; Party Moderation never has a big effect.

swamped by the effect of partisan mobilization. As we shall see below, some modest evidence exists for this scenario.

MAIN FINDINGS: LOW INFORMATION VOTERS IN PRESIDENTIAL ELECTIONS

We are now ready for the main findings of the paper. Without further ado, I introduce them in Figure 9, which shows three scatterplot of economic performance (Hibbs variant) against the incumbent party. Each scatterplot is based on the aggregate vote among low, middle, and high information voters in the CPS/NES surveys. The clear pattern is that low information voters are more sensitive to economic performance than high information voters. To judge by the slopes in Figure 9, they are about 2.5 times more sensitive. The same substantive pattern shows up for the other two measures of economic performance.

INSERT FIGURE 9 ABOUT HERE

The points for 1952 and 1968, which are the years of the Korean and Vietnam wars, are also labeled. From an eyeball scan of these data, low information voters react at least as strongly to these painful events as other voters. But a close look at the war elections suggests the fragility of the data. High information voters seem highly sensitive to the Korean War, falling 11 percentage points below their expected vote in 1952. Yet in 1968, high information voters support the incumbent candidate at a rate slightly above the economic performance trendline. So what is really going on: Are high information voters more sensitive to war, less sensitive to war, or what?

Without much doubt, chance is playing tricks with us. Apart from 1948, when there are only about 130 voters per point, each cell in Figure 8 has an average of 400 voters per point, with an SD of 75. The SE of each of these points is then about +/- 3 percentage points. In view of this, we should obviously avoid paying too much attention to pairs of individual points. Yet, as regards the effect of war among high information voters, a pair of points is all we have!

The difficulty of discerning clear patterns in these data is reduced by variables, like the new foreign policy measures, that assign values to numerous points rather than just two. But the problem never entirely disappears, as we shall have occasion to observe.

Table 2 provides a broader but less detailed view of the effect of information on voting. It does so in the form of partial correlations between vote and causes of the vote, within each of five information groups. The basic data are 14 aggregate means per level of information, as in Figure 9. The reason for partial correlations is to remove the effect of party of the incumbent, as was also done in Figure 9.

INSERT TABLE 2 ABOUT HERE

Even allowing for distortions that bivariate correlations can introduce, the data in Table 2 continue to feed the impression that low information voters respond to politics in a coherent fashion. Low information voters are, by slight but consistent margins, more responsive to candidate moderation than any other group, and the same is true for foreign policy performance as variously measured. The low information vote is somewhat less correlated with economic performance than are the votes for some other groups, but given the high correlations that exist with other variables, this may be almost a mathematical necessity.[40] Correlations with Party Fatigue indicate no strong pattern.

The next step, then, is to test the patterns in Figure 9 in a multiple regression. Before doing so, I need to discuss a specification issue and a scaling issue, starting with the former. As noted earlier, any analysis of Candidate Moderation must make an assumption about the midpoint of the electorate. As regards the electorate as a whole, the most plausible assumption is that the midpoint corresponds to the center-point of scales measuring ideology. However, we are now about to analyze voting within subgroups defined by levels of political information, and within these groups the zero-point assumption is untenable. It is, in other words, untenable to assume that the ideological center of low information groups, which tend to vote Democratic, is the same as it is for high information groups, which tend to vote Republican. An easy and standard fix for this problem is to add a dummy variable for the party of the incumbent president. This dummy is not necessary for the electorate as a whole, where it has been assumed that the mean is zero,

40 A variable cannot have high bivariate correlations with more than a handful of variables unless those other variables are highly intercorrelated, which these are not.

but is needed within subgroups of the electorate. By the logic developed above, this dummy variable sets the center-point of subgroup opinion according to whether the incumbent party candidate is a Democrat or Republican. But there is a concern: This approach requires us to assume that the center-points of the several information subgroups are stable across the time of the study – that is, that the liberal/Democratic propensity of low information voters is roughly constant from 1948 to 2000, and that the conservative/Republican propensity of high information voters is likewise roughly constant. This seems a strong assumption, but Figure 10 shows that it is empirically plausible. The figure shows the amount by which voters in different information subgroups vote more or less Democratic than the electorate as a whole in different time periods. The results indicate a remarkable stability of partisan coloration within every group but one: High information voters became more Democratic in the 1990s, such that they are now roughly as Democratic as the rest of the electorate rather than less so. Whether the change is permanent or a response to particular events is hard to say at present.

INSERT FIGURE 10 ABOUT HERE

These results amount to a lucky break for my analysis. If groups as defined by information level were less stable in their partisan bias than they are, it would require complicated modeling to control for their changing center of political gravity. But in view of the results in Figure 10, I can with reasonable safety use a simple variable for party of the incumbent to capture the ideological center in subgroup regressions.

Now the scaling issue. For the sake of stability, I wish to present average results rather than results from individual coefficients. But coefficient averaging makes no sense unless the underlying variables are on a common metric. As noted above, I have already re-scaled the NBI instrument into the units of the simple RDI measure. And from here onward, I shall use a version of the Hibbs measure that has also been scaled to have the same mean and SD as the simple RDI measure.

Figure 11 and Table 3 now present the principal empirical findings of the paper. The figure shows multiple regression results within each of five information subgroups, based on 14 aggregate mean votes for the incumbent party from 1948 to 2000. The purpose of the figure is to show how the impact of each aggregate variable varies by level of information. (A sample of the regressions which the figure summarizes are reported in table form in Appendix B.) Table 3 estimates these same information interactions in an individual-level model involving some 16,000 individual observations rather than 14 aggregate means.

The aggregate and individual-level analysis are complementary rather than redundant: The aggregate analysis – because based on a separate regression within each subgroup – places no restriction on how information can interact with economic performance and other determinants of the vote. In particular, non-monotonicity is allowed. The downside is that purely chance variation is given wide scope to produce oddball results. The individual-level analysis, for its part, imposes a strong constraint on the data – that any interactions with information must be strictly linear. But in so doing, it protects from chance variation, tests the statistical significance of any linear interactions that turn up, and accommodates control variables. Any patterns that show up in both forms of analysis will thereby gain credibility. Patterns that show up in only one will be somewhat suspect.

Look first at the upper left panel of Figure 11, which is a graph of the coefficients for economic performance in six regressions. The reason for six regressions is that I have crossed 3 measures of economic performance with 2 measures of foreign policy performance. The solid black line is the Bayes-weighted average of these coefficients; the dotted lines are the coefficients from the separate regressions.

For respondents in the bottom quintile of information, the average effect of economic performance on the vote is about 2. This means that for every 1 point change in Real Disposable Income, incumbent party vote increases by about 2. points among low information voters. This effect falls to close to zero for high information voters. The data exhibit some non-monotonicity, but the pattern does not appear strong or general enough across other variables to be credible. (I return to this issue below.)

Next look at the effects of foreign policy performance. The three solid black lines show effects for the Consensus scale and the three dotted lines show effects for the Expert scale. Because the units of the underlying scales are different, I present no Bayes weighted average, but the results suggest that low-information voters may be the most responsive to foreign policy performance. However, this trend is heavily driven by voters at the bottom of the information scale, so we should withhold judgment until results from the individual-level model are examined.

Results for Candidate Moderation are shown at the lower left of Figure 11. As in the graph for economic performance, the figure shows trends for each of six individual regressions plus a Bayes-weighted average. These data show a very strong gradient by political information. Again, I do not recommend taking the hint of non-monotonicity seriously, but note that the average effect of Candidate Moderation among high information groups is about -11. Taken at face value, this is an assertion that high information voters rather severely penalize moderate candidates and reward extreme ones. Although this possibility may seem initially plausible – high information voters are, after all, the most ideologically oriented segment of the electorate – it seems most likely the sort of odd statistical bounce that one fears in a small dataset. As discussed earlier, the Party Fatigue variable correlates at -.53 with Candidate Moderation. A common effect of such multicollinearity is off-setting over-estimation of the effect of each variable, and this is what is most likely happening here: If the Party Fatigue variable is omitted from the model, the Candidate Moderation effect among high information voters falls to a much more plausible value – essentially zero.[41] See also Table 2.

Nonetheless, the finding is sufficiently robust to merit consideration. If real, how could it be explained? I can imagine that high information ideologues remain stubbornly or even happily loyal to their party when it nominates a candidate like Goldwater or McGovern. That high information centrists actually flock en masse to extremist candidates is more dubious, but Aldrich (1983, 1996, chapter 6) has argued that ideological activists might participate more in campaigns

[41] When, however, the fatigue variable is eliminated from the low information model, the positive effect of Candidate Moderation remains robust.

and perhaps even turn out to vote in greater numbers when offered, in Goldwater's words, "a choice, not an echo." Thus, the wrongly signed coefficient on moderation among high information voters could be evidence of sophisticated ideologues surging to the polls – or at least claiming in NES surveys to have surged to the polls – to support a favorite. This possibility still seems to me dubious, both because of the statistical problems mentioned and because the surge appears so large. Nonetheless, the data clearly lean strongly in that direction and so must be taken seriously. I return to this point below.

In the meantime, let us turn to the final variable in Figure 11, Party Fatigue. Here we find a non-monotonic pattern. Yet the non-monotonicity depends on a big impact among high information voters and, once again, I suspect an artifact: Multicollinearity exists with respect to Candidate Moderation, and if this variable is omitted, the impact of Party Fatigue among high information voters is about -2, its value for other information groups. Hence, my tentative conclusion is that we lack consistent evidence that the effect of the Party Fatigue variable varies by political information. (Later analyses continue to fail to clarify whether Party Fatigue interacts with information.)

Table 3 tests these same interactions at the individual level, but with several control variables: Party identification, Race, white southerner, and education. To capture party realignment among white southerners, this term for this group is interacted with a term for "post Civil Rights Act." Each of the individual-level variables is included as a main effect term and as an interacted term that is multiplied by -1 when the Republicans are the incumbent party.

The most important of the control variables is party identification, coded 1 for party of the incumbent, -1 for party of the challenger, and 0 otherwise. In view of Brody's (1991) finding that the direction of party attachment is stable but that strength varies with political context, I expected this scoring to minimize endogeneity, but the variable nonetheless responds to economic performance, candidate moderation, and even foreign policy performance. Hence I ran a first-stage equation to purge the party variable of these endogenous influences.[43] The 1948 study lacks a party identification variable, so rather than omit either the variable or the case, I gave all 1948 respondents a value of 0, which is the mean of the purged party variable.[44]

In order to obtain correct standard errors in a pooled cross-section analysis, one must add a 0-1 dummy variable for each election year except one. These dummy terms make it impossible to estimate the direct effects of aggregate variables, since the dummies explain all across-year variation in the data. However, interactions between political information and each of the aggregate variables can be included to explain within-year variation, and these interactions are our principal interest.

The model thus specified contains six aggregate-level variables that interact with political information – economic performance, foreign policy performance, moderation, party fatigue, incumbent party, and party identification -- plus a handful of controls. Given that we have 16,000 individual-level cases, this may seem like a modest number of variables. But crucial variance still comes from only 14 aggregate cases, so the instability and multicollinearity present in the previous regressions are by no means absent from these.

The results, in the first three columns of Table 3, are easy to summarize because they vary little across specification. Political information has highly statistically significant negative

---

[43] The first stage equation includes year dummies, interactions between information and each aggregate variable, and information. The interactions are necessary because party attachment among low information voters is more responsive to the aggregate variables than is party attachment among high information voters.

[44] One other compromise was necessary to retain the 1948 sample in the analysis: Since the 1948 study did not carry a variable for state or region of residence, I assigned the constant value of .15 to the South variable in this year, which was the proportion of respondents who lived in the South in the 1952 study. Race and a truncated education variable were included in the 1948 study, fortunately. Data imputation via BUGS or Amelia should perhaps have been used to assign party identification and South values in the 1948 sample, but I have not had time to do that.

Less than one percent of voters are also missing on the education variable in various years; I assigned these case the mean of the education variable in the year of their study. Approximately 0.2 percent of respondents are missing on political information and are omitted for this reason. Except for survey respondents who were not re-interviewed in the post-election study, these are the only cases missing from the analysis in Table 3.

interactions with economic performance, Candidate Moderation, and Party Fatigue. Information also interacts with foreign policy management, though more weakly.

Note that all these interactions are negative. This should be interpreted as follows: Economic performance, foreign policy performance, and candidate moderation all have direct positive effects on the vote. The negative interactions with information indicate that those positive effects diminish as information increases. Thus, low information voters are more responsive.

For Party Fatigue, however, the direct effect on the vote is negative. That is, vote share goes down as Party Fatigue increases. Thus, a negative interaction with Party Fatigue indicates that the negative effect increases as information increases. Thus, high information variables are more responsive to Party Fatigue. This, however, is a difficult effect to interpret. Party Fatigue is a grab-bag of effects that go with long party control of the White House – the determination of the challengers, the complacency of the in-party, a resort to second-string candidates, among others. Voters presumably respond to these individual elements rather than to the general phenomenon. In view of these difficulties, I have little to say about Party Fatigue beyond the bland statement that it seems to be an important factor in explaining voting behavior. Moreover, as we shall see, this interaction fares poorly in some later tests.

With respect to foreign policy performance, the interactions are also all negative, but none is statistically. Nonetheless, these modest estimates imply non-trivial information effects: That a voter who scores -1.2 SD on information places about 60 percent more weight on foreign policy concerns than a voter who scores +1.2 SDs on information.[45] (The reader can best judge the magnitudes of other variables from Figure 11.)

INSERT TABLE 3 ABOUT HERE

Here is a digest of results from alternative specifications:

---

[45] I calculate this as follows: The simple average of the three estimates in columns 4, 5, and 6 is -.75. Given that the SD of information is .94 and the Bayes weighted main effect of Consensus foreign policy in Table 1 is 3.12, the ratio of effects for +1.2 SD versus -1.2 SD is 3.83/2.42 = 1.58. The Bayes' weighted estimate of the information effect would be higher, but Bayes weights don't make sense to me in this context.

- If the Expert foreign policy scale is used instead of the Consensus scale, levels of statistical significance and magnitudes increase marginally. If the "wars only" form of the foreign policy is substituted into these models, the sign of the interaction indicates in all three models that low information voters punish incumbents more for war than high information respondents do, but the results do not approach statistical significance..

- If the party identification variables are omitted from the analysis, the interaction terms become somewhat statistically weaker but nothing of importance changes.

- As noted, party identification is affected by national conditions and candidate moderation, making it endogenous to the voting process. If, however, the exogenous causes of party identification are included in the analysis, as they are, it is not strictly necessary to purge party of their effects. One can simply use party as an ordinary control. When this is done, all the key interactions gain in statistical significance. Most notably, the information interaction with foreign policy becomes strongly negative and statistically significant at the .01 level in all three models, two-tailed. I do not focus on this unexpected result, but it should be reported.

- As discussed earlier, high information voters show an implausible boomerang response to Candidate Moderation, such that they appear to punish moderates rather than extremists. It is natural to wonder whether this odd result is perhaps driving the negative interaction between information and Candidate Moderation in Table 3. It is not. If voters in the top fifth of the information scale are filtered out of the regressions in Table 3, the interaction between Candidate Moderation and information becomes smaller but remains statistically significant at the .05 level, two-tailed, in all cases.[46]

- If a term for Party Moderation X Information is added to these models, it gets a relatively small positive coefficient that is statistically significant. Its presence undermines the effect of foreign policy managements and magnifies the larger negative interaction between information

[46] The interactions with economic performance are more affected, becoming non-significant and reversing in one case. The negative interaction with foreign policy performance turns slightly positive. Since there is no concern about an artifact with respect to these effects, this weakening seems unimportant, but I nonetheless report it.

and Candidate Moderation. Thus we get the odd result that high information voters reward Party Moderation but severely punish Candidate Moderation, especially so with Party Moderation controlled. This is most likely a pathology of multicollinearity, as noted earlier. If, to combat the problem, the two variables are simply added together to create a single scale, the new scale generates results generally stronger than Candidate Moderation alone, most notably for the foreign policy interaction, which statistically significant at the .05 level, two-tailed, in all three trials with the new scale.[47] Despite these strong results, I remain skeptical of this variable, as indicated in my initial discussion of it.

Taken altogether, the individual-level results reinforce those from the aggregate analysis: The two information interactions that looked strong in Figure 11 – with economic performance and candidate moderation -- turn out in individual-level models with ample controls to be statistically significant and reasonably robust. The interaction with foreign policy that looked uncertain in Figure 11 remains statistically marginal but a nonetheless substantively important effect. Only with respect to Party Fatigue do we get mixed signals – a questionable interaction in Figure 11 and a strong one in Table 3.

Given all that has been written about the low levels of information that characterize the lower strata of the electorate, these results are important. They suggest much more political responsiveness on the part of the less informed than existing scholarship suggests exists. However, the findings still face an important threat – the possibility that they may stem from survey bias.

As discussed earlier, CPS/NES respondents tend to overstate their votes for the winning candidate by about 2 percentage points, an amount that tends to inflate the impact of aggregate variables on the vote. I corrected for the over-report, but did so in a way that assumed all respondents were equally likely to over-report. But what if, as seems plausible, low information

[47] In columns 4 through 6 of Table 3 and in Table 4, this combined variable produces consistently strong results for the foreign policy interaction, though somewhat weaker effects for moderation.

respondents are more likely to over-report than high information ones? This could, in principle, explain much if not all of the evidence of information interactions examined so far.

To assess this threat, I created a new vote variable on the assumption that low information voters are three times more likely to over-report than high information voters and two times more likely to do so than middle information voters. I assumed, in other words, that over-report occurs in proportions of 3 -2 -1 among low, middle, and high information voters. Results for this "over-corrected" variable are shown in columns 4 through 6 of Table 3.

As can be seen, the results for the interactions with economic performance and moderation hold up well. The weakest interaction terms – with foreign policy – fall by an average of about 60 percent. At this magnitude, they indicate that a low information voter is only 20 percent more responsive to foreign policy management than a high information voter.

On the whole, these results are also reassuring. Even under the assumption that low information voters are much more likely to over-report than high information voters -- an assumption that, as implemented, guaranteed weaker interactions with information -- the former remain more responsive to the aggregate determinants of elections than high information voters.

Before reaching a final conclusion, however, we must examine vote intention. If the pre-election measure of vote intention shows the same pattern as the post-election measure of reported vote, it would significantly increase confidence in the results.

As it turns out, it does mostly reinforce them. The caveat is due to the fact that, as we saw in Figure 8, vote intentions change over the course of the campaign. Vote intentions at the start of the campaign do not exhibit the information interactions evident in Table 3, but they do at the end of the campaign.

Evidence on this point is in Figure 12, which shows results for aggregate vote intention at two points in the campaign and for self-reported vote. Table 4 presents these results in an individual-level model directly comparable to Table 3. In both Figure 12 and Table 4, the analysis is limited to persons who reported in a post-election interview that they had voted. This restriction entails the loss of some actual voters who were not interviewed after the election. But is also entails two gains: It excludes many non-voters from the pre-election analysis of vote intentions, as is appropriate in a study of actual voting; and it makes it meaningful to compare vote intentions with self-reported votes, because they are based on the same individuals.

INSERT FIGURE 12 AND TABLE 4 ABOUT HERE

The two campaign measurements were obtained by dichotomizing the CPS/NES pre-election survey into an early and a late measurement. The early measurement is based on interviews conducted 30 or more days before the election; the second covers interviews from the remaining period of the campaign. Although the CPS did a pre-election interview in 1948, it does not include a time of interview variable, so the CPS/NES campaign readings run only from 1952 to 2000.

The CPS/NES surveys were not intended to be broken into time periods in this way. Interviews were conducted on different schedules in different years, and in most studies neither half of a study, taken separately, is a proper random sample. Nonetheless, as we saw in Figure 8, these data produce overall results quite close to what the much stronger Gallup data produce. Given this, it is worth looking at these data. Even though imperfect, they provide what is in some ways the best window we have on the pattern of campaign effects over a half century and some 16,000 individual survey respondents.

Let us begin with the coefficients for economic performance, which are shown in the upper left of the figure. The broken line depicts the effects for low information voters; the solid line is for middle information voters; and the dotted line shows effects for high information voters. The coefficients are Bayes-weighted averages across six measures of economic performance; I do not present the individual regressions because this would entail the clutter of 18 more lines.

As can be seen, low information respondents are not more responsive to economic performance at the start of the campaign, but become so in the last month. Table 4 presents results from individual-level models of vote intention in the second period of the campaign that confirm the visual impression. For these regressions, the data from 1948, all of which were collected in October or November, are added in. (I omit any 1948 data from the figure in order to maintain visual comparability across the three time periods.)

The pattern for Candidate Moderation is shown at the upper right. Here low information voters are more responsive at the start of the campaign and remain so at the end. The individual-level models in Table 4 confirm this pattern. Thus, results for vote intention again match those for vote report.

The two panels in the middle of Figure 12 are best described as confused. Both at the end of the campaign period and in the reported vote, low information voters are more responsive to foreign policy management than high information voters, but middle information voters do not follow a clear middle path – or any clear path.

These are noisy aggregate data at their worst. But still we must ask: Why the change from earlier results in Table 3 and Figure 11? The answer, in a word, is 1948. This case has been omitted from Figure 12 because the 1948 study did not carry a time of interview variable. However, all of the 1948 interviews were done in October and November, and presumably most of them after the first few days in October. Hence all or nearly all belong in the "late campaign" period. With that in mind, 1948 has been included in the late campaign regressions in Table 4, and there we find a pattern in which low information vote intentions are more responsive to foreign policy management than high information vote intentions. This effect, though statistically weak, is on average almost exactly as big it is in the first three columns of Table 3, which are based on more than twice as many cases. Thus, when all the vote intention data are examined, there is support for the notion that low information voters are as responsive as high information voters, and probably somewhat more so, to foreign policy concerns.

The final results in Figure 12 are for the Party Fatigue variable. Here the pattern is again muddy, and in this case the individual-level data in Table 4 are unable to clarify it.

Before closing out the empirical analysis, I have four matters to wrap up:

- The tests in Table 4 involve, as indicated, voters interviewed in the second half of the campaign. This was because voters interviewed early in the campaign show a markedly different pattern of vote intentions in Figure 12. The most notable difference was that low information voters were not initially more responsive to economic performance than high information voters. But how about after the campaign? Had these same low information voters become more responsive to economic conditions than high information voters? The answer is yes. When formally tested in regressions of the form of Table 3, vote intentions in the early part of the campaign were not affected by an interaction between economic performance and information. However, when the vote reports of these same individuals were examined after the campaign, the interaction was strongly present.[49] The campaign, thus, functioned to prime the economic concerns of low information voters.
- We saw in Figure 11 that high information voters seem to reward rather than punish extremism. That implausible result is also present in the vote reports in Figure 12, which are based on a three-way rather than five-way division of the vote. Note, however, that this result is not present in the vote intention data for these same high information voters. Indeed, at the start of the campaign, high information voters consistently punish extremists: The average value of their Candidate Moderation coefficients is +6.3, with an average p-value of .15. Only in the post-election data do the moderation coefficients turn negative. Although it is possible that an initially positive response to moderation turns negative by the end of the campaign, it seems more likely that the negative estimate is an artifact of multicollinearity in a small dataset. Still, future research should keep an eye out for the possibility that the effect reflects actual behavior.

---

[49] More specifically, I re-ran the six regressions in Table 3 of reported vote and “over-corrected” reported for the half-sample of respondents who had been interviewed in the first half of the campaign. Despite the reduced sample, interactions were statistically significant at the .05 level, two-tailed, in five of six trials and at .11 in the sixth.

[52] In a set of six tests, the education interaction with foreign policy was always statistically significant, such that poorly educated were more responsive to foreign policy management. These education results are based on using a six-point education variable, or the nearest to that range possible in a given survey. The five-point education variable available in the NES Cumfile produces notably weaker results.

- Figure 11 also suggested the possibility of non-monotonic interactions with political information, such that voters near the bottom of the information scale, rather than at the very bottom, might be most responsive to the aggregate determinants of the vote. To my eye, the suggestion is a weak one, since the data are noisy and non-monotonicity is not always present. Most likely, the responsiveness curve slopes gently downward from low to middle information and turns more steeply downward at that point. But the idea of non-monotonicity is certainly plausible, so it is useful to formally evaluate it. This is not straightforward to do. Adding a series of interaction terms with the square of information would be to invite a multicollinear explosion. So I developed a more constrained model on the following logic. Suppose that we take the absolute value of an information variable that runs from -2 to +2. The result would be a variable, folded on the original 0-point, in which low scores on information were the same as high scores. If this variable had a larger effect than the original variable, we would take it as evidence of a non-monotonic effect, since high scores have the same effect as low ones. Now suppose that we started with the same information scale, added +1 to it, took absolute values, and tested its effect in a model. If the effect of this variable, folded at the value of -1, was bigger than other forms of the variable, we could take it as evidence of a non-monotonic effect with the inflection point at -1 rather than zero. This is the approach that I followed, except that I created a model to estimate the inflection point. The results disclosed only minimal hints of non-monotonicity in the overall effect of the information variable. However, further exploration of the question is warranted. It would not be surprising to find that middle to low information voters are most responsive to the political content of elections, and that only high information voters are so strongly tied to parties.
- Finally, it is natural to wonder what happens when variables other than political information are used to capture voters' engagement with the political process. I have briefly tried two in individual-level models, formal education and interest in the campaign. As would be expected from past research, both variables produce results that resemble those obtained from an information variable, but are not as consistently strong. Formal education generates no interaction with economic performance, but quite strong interactions with foreign policy

performance and with candidate moderation.[52] A simple three-point scale of "interest in the campaign" produces some statistically significant interactions with economic performance, moderation, and foreign policy performance. Even when strong, the results from the political interest variable are difficult to interpret, because political interest may wax and wane in ways highly endogenous to vote choice, whereas political information is an unusually stable individual-level characteristic (Delli Carpini and Keeter, 1996). Campaign interest may capture little more than the partisan exuberance of voters supporting their customary party. Still, the result is worth noting.

Summary. The empirical heart of this paper consists of columns one to three of Table 3 and Table 4. The results reported in these tables consistently support the idea of a robust interaction by which low information voters are more responsive to economic performance and candidate moderation. The results also consistently support the idea of a less robust but not substantively small interaction of the same sort with foreign policy management.

In an effort to communicate fragility as well as strength of these patterns, I have examined them in aggregate form, where considerable fragility is apparent. As I have argued, the overall sense one gets from the aggregate data is consistent with the individual-level findings: Patterns that are strong in the aggregate data are strong in the individual data; patterns weak in the aggregate are weak in the individual data..

VI. Theoretical Conclusion: Low Information Voters are Not Fools

It is conventional to begin scholarly monographs with a review of the literature that prepares readers for the results to follow. This paper contains no such preparation, for two reasons. The first is that it is an overwhelmingly empirical effort to answer a question that is inherently interesting and important, and that seemed enough. The second is that few readers would have attached much credence to an attempt to prepare the ground for the results I have presented.

Indeed, even after seeing my results, some scholars initially insisted – and perhaps would still insist – that I must simply have reverse-coded the information variable.

But now the question arises: How shall we understand these findings that few if any experts in the field, including myself, would confidently have predicted? How do we explain how low information voters – voters so poorly informed that, in many cases, they do not know which party controls Congress – manage to be more responsive to national conditions and even the ideological positions of the candidates than high information voters?

Most of the answer, I believe, is contained in the foundational classics of the Michigan school of political behavior – *The American Voter* (Campbell et al., 1960), "Nature of Belief Systems in Mass Publics" (Converse, 1964), and "Information Flow and the Stability of Partisan Attitudes" (Converse, 1962).

Near the end of Chapter 10, which is surely the best known chapter of *The American Voter,* Campbell *et al* observe that "It appears that quite generally in this period [from 1948 to 1956], people who paid little attention to politics were contributing very disproportionately to partisan change" (p. 264). The Michigan scholars offered three reasons why political involvement was associated with stability. One was that the ideological commitments of highly involved voters kept them firmly attached to the party that was more congenial to their beliefs. Another was that involvement itself strengthened partisan attachment:

> Once a person has acquired some embryonic party attachment, it is easy for him to discover that most events in the ambiguous world of politics redound to the credit of his chosen party. As his perception of his party's virtue gains momentum in this manner, so his loyalty to it strengthens, and this fact in turn increases the probability that future events will be interpreted in a fashion that supports his partisan inclination… typically, party identification is not only sustained but strengthened by the passing show of political acts and actors. (p. 165).

Two years later, Converse proposed another kind of learning mechanism in his paper on "Information Flow and the Stability of Partisan Attitudes." His central argument was that exposure to politics builds up informational ballast that engenders inertial resistance to change.

The American Voter offered complementary reasons for finding a lack of stability among the less involved. One was a preoccupation with "the nature of the times," including economic bad times and good times, among less involved voters. Another was that voters at the very bottom of the "Levels of Conceptualization" measure exhibited a marked drop-off in strength of party attachment. Voters in this lowest group, called the "No issue content" category, include some diehard party voters, such as one who said, "I'm a Democrat , that's all I know." Yet, the largest bloc of voters in the bottom level is classified "candidate oriented." Here are excerpts from the three cases of candidate-oriented voters that the *American Voter* gives us:

- I vote for who I think is the best man. The last two elections have been mudslinging, and the Democrats are responsible for this… They've [the Republicans] just had the best man in the last two elections, that's all I can say.
- I just like [Eisenhower], the way things have gone…. That's really all I know.
- I hate the darned back-biting… I don't like the cut-throat business – condemn another man and then shake hands with him five minutes later…As a man I like Eisenhower better … he is not so apt to cut your throat.

In many ways, the present paper has done little more than buttress these classic views. Politically involved voters, as measured in this paper by political information, are far more stable than the less involved. In good times and bad, these voters stick with their party. And strongly ideological candidates do not put them off. At the low end of the spectrum, we find voters who lack party loyalty and who are highly sensitive to the nature of the economic times.

The one finding in the paper that may not entirely fit with the general view of the *American Voter* is that low information voters are more sensitive to the ideological positions of candidates. But these same kinds of voters are, according to the Levels of Conceptualization framework, sensitive to candidate qualities and, as it seems from their open-ended remarks, to partisan negativity. It is hard to be an extremist without slinging mud, so if we imagine that low information voters see politicians like Goldwater and McGovern less as ideologues than as hard-edged, mudslinging *politicians*, then their sensitivity to extremism is perhaps intelligible.

I must add, however, that I do not think back-biting and mudslinging are exactly what low information voters dislike. The respondent in the *American Voter* who blamed Democrats for mudslinging in the 1952 and 1956 elections seemed somehow to overlook McCarthyism, which was not a Democratic theme. This person also liked Franklin Roosevelt, whose 1936 campaign was one of the most unabashedly negative in American history.[53] So we should probably not take views of low information voters at face value on this point. They may like those ideologues and extremists whose negative views resonate with a populist or authoritarian view of the world. In 1968, for example, George Wallace put on a world-class mudslinging campaign, yet we find a robust negative correlation between political information and rating Wallace high on the NES feeling thermometer.[54] What low information respondents probably really do not like are candidates who take an aggressively principled or high-minded approach to politics. In the staid context of American politics, this could often mean disliking candidates who are ideologically extreme.

But this argument does not account for quite all of the evidence. Ronald Reagan in 1980 was the second most extreme Republican candidate of the post-war period, only a bit less so than Goldwater (see Figure 3). Yet Reagan was not hard-edged, and if anyone slung mud in the 1980 election, it was the moderate Jimmy Carter, who blasted away at Reagan's extreme policies until the press criticism for "meanness" put him on the defensive.[55] So this might be considered a critical test: Did genial but extreme Reagan, in opposing the moderate but hard-edged Carter,

[53] In the 1936 campaign Roosevelt described big business people as "economic royalists" who are "unanimous in their hate for me -- and I welcome their hatred." He added: "I should like to have it said of my first administration that in it the forces of selfishness and of lust for power met their match; I would like to have it said of my second administration that in it these forces met their master."

[54] In a multiple regression that controls for race, party, age, income, gender, and being a southerner, respondents scoring near the bottom of the information scale were, all else equal, 20 points warmer toward Wallace than persons scoring near the top of the information scale. Information is third in substantive impact on affect for Wallace, just behind race and southerner.

[55] In a famous incident, Carter was asked by Barbara Walters in a televised interview about the "meanness" of his remarks and he admitted that he may have gone too far.

nonetheless suffer an extremism penalty among low information voters? The best evidence on this point is from a fragile point in the left-hand panel of Figure 9 – that for the 1980 election. This point indicates that Carter finished 4.4 points above the economic trend line among low information voters in that year. In terms of the accounting scheme I have been using, this is evidence that low information voters penalized Reagan for his relative extremism.[56]

The evidence thus continues to stubbornly assert that low information voters are responsive to the relative moderation of the candidates on the issues. Yet any argument that low information voters are actually more sensitive to issues than high information voters must be carefully made. The way I would put it would be that low and high information voters respond to different elements of the political landscape than high information ones. High information voters are more responsive to partisanship and to ideological principle, prominently including the general philosophies of the two parties. Since parties and party philosophies are mostly fixed from one election to the next, responsiveness to them results in stability. Low information voters, by contrast, are responsive to candidates relatively independent of parties, which results in instability.

But what, concretely, are low information voters responding to? The Candidate Moderation variable takes its highest absolute values in these elections:

- 1964. Goldwater, insisting that "extremism in the defense of liberty is no vice," has a record of opposition to social security and efforts to end racial segregation. He has also hinted at using nuclear weapons in Vietnam and joked about lobbing a nuclear weapon "into the men's room at the Kremlin."

[56] For a comparable analysis using the RDI measure, Carter is 3.7 points above trend and for the NBI measure he is 4.7 points above trend. These results are without counting 1980 as a negative foreign policy event, which makes Carter's above-trend performance even more notable. The number of cases in the cell for low information voters in that year is 269. Recall that, unlike most figures in this paper, Figure 9 is organized in terms of the Democratic percentage of the two-party vote, with informational differences in mean Democratic vote removed so that all three panels in Figure 9 center on 50 percent. The unadjusted Democratic vote for Carter in 1980 among low information respondents was 55.9 percent.

- 1968. Humphrey is remembered by liberals as the timid supporter of Johnson's Vietnam policy, but Rosenstone's coders saw him as a civil rights firebrand (in a year of racial disturbances) and a supporter of federal job guarantees. This made him highly liberal on both the racial and New Deal dimension of politics.
- 1980. In this third run for the presidency, Reagan sought to present himself as a moderate. But as late as 1976, he had spoken against social security and the Tennessee Valley Authority. He was also accused of being a "cowboy" on foreign policy matters.
- 1972. McGovern favored a "reverse income tax" as a way of redistributing wealth to the poor, an immediate pull-out of U.S. troops from a war that seemed to be ending on U.S. terms, busing to end racial segregation of public schools in the North, and major cuts in defense spending.

Thus each of these "extreme" candidates challenged important elements of the post-war status quo – a status quo that was presumably where it was because mainstream opinion wanted it there. One would not need exquisite sensitivity to politics to recognize such challenges; everyone who showed up to vote must have known. If high information voters – including the moderates among them -- nonetheless tended more often to vote loyally for their customary party, it must be because they were more forgiving than low information voters.

Throughout this discussion, I have described low information voters as "more responsive" and "more sensitive" to Candidate Moderation, but this may be the wrong way to frame it, for there is a real sense in which low information voters actually *lag behind* high information voters rather than race. Look again at Figure 8, which shows how the penalty paid by extremist candidates diminishes as the campaign progresses. Presumably this means that extremist candidates gradually gain the support of voters who initially shunned them. Earlier data suggests that this gain comes more from mobilization of the undecided than from conversion of the committed. In Figure 12 we further see that extremist candidates gain support within all voter groups. High information voters, who are heavy consumers of political communication, begin campaigns relatively supportive of extremists, and become still more so as they receive more communication in the course of the campaign. Low information voters, who consume less political communication, start the campaign *lagging behind* high information respondents in support for extremists, but partly *catch up* by becoming more supportive during the campaign. If the campaign went on long enough, they might fully catch up, which would entail the complete disappearance of the moderation effect. Thus, what I have been calling the greater sensitivity of low information voters to the ideology of the candidates could reflect something more like a lag in partisan mobilization behind weak (i.e., extremist) candidates. In this account, many high and low information voters dislike extremist candidates, but high information voters, because more partisan and more exposed to campaign communication, do a better job of overcoming their dislike.

How exactly low information voters understand extremist challenges to the status quo is unclear from these data. Perhaps they personalize their perceptions in the form of complaints about back-biting and mudslinging, as Campbell et al's materials suggest. Or perhaps low-information voters take cues from trusted reference groups, as Lupia (1994) and Lupia and McCubbins (1998) would suggest. Or perhaps their "gut-level rationality" enables them to spot the oddball and avoid him, as Popkin (1991) might argue. Or perhaps extremism evokes fear, which motivates a successful search for information, as Marcus, MacKuen, and Neuman (2001) would perhaps argue. My preferred explanation would be that extremist candidates attract more criticism – including criticism from fellow partisans at their nominating conventions – and that today's high intensity campaigns communicate this information to essentially everyone who cares enough to cast a vote. Thus, many different mechanisms might plausibly explain the issue sensitivity of less informed voters in presidential elections.

One should not, however, conclude from the results in this paper that the low information voters who were more centrist and status-quo oriented in the period 1948 to 2000 would be more centrist in every context. Even in the absence of pertinent data, one must, for example, assume that low information voters in the Great Depression were more likely to embrace what would, in the terms of this paper, be called the “extremist” policies of Franklin D. Roosevelt than were high information voters.[57] I would therefore propose to summarize the propensities of the less informed as follows: *They are less strongly attached to, or socialized into, the dominant ideological and partisan structures of society and therefore more apt to abandon them when societal problems arise*. And they are not so ignorant of politics that they fail to recognize problems when they occur. Thus, they are more apt than the highly informed to reject their party when it nominates an ideological oddball in otherwise normal times, and also more apt to reject any sort of incumbent in troubled times, even if that means embracing a relative extremist, such as Roosevelt. The possibility that low information voters might have embraced an even more radical extremist in the Great Depression, as many European voters did, is not, I think, contradicted by the findings reported in this paper if these findings are understood as relative lack of attachment to established structures.

Thus, the propensities of the less informed are still well-captured by the term Daudt used for them 40 years ago, “floaters.” If low information voters are better seen nowadays through the more normatively positive frame of “retrospective voters,” it may be more because the nature of the times has changed since the perilous days of fascism, mass communism, and world depression than because the nature of low information voters has changed.

This set of interpretations departs from the classic account of the American Voter, but only modestly. It embraces the classic arguments that high information voters are more partisan and ideological, and that low information voters are more concerned about “the nature of the times” and candidate qualities. But it departs by infusing a tint of issue coloration into the “No Issue Content” category, arguing, in effect, that the many “candidate-oriented” voters at the low end of the involvement spectrum are better characterized as *low* information than as *no* information actors – and as being quick to act on such information as they do possess.

---

[57] We do know that class voting became much stronger in the Roosevelt elections (Sundquist, 1983).

It should be added that Campbell et al. certainly recognized that citizens in the bottom rung of their classification scheme are capable of political responsiveness. Thus, they wrote that “some proportion of people classified in Level D are placed there not because they are unable to comprehend [politics], but because the times do not focus their attention on the relevant political phenomena” (p. 256). What Campbell et al. may have underestimated, then, is not individuals’ capacity, but the ease by which modern political campaigns can focus everyone’s attention on politics.

The present paper also fills out, without in any way contradicting, the “information flow” argument proposed by Converse (1962) and elaborated by Dryer (1970) and Zaller (1992, 1996). As part of this argument, Converse maintained that voters who paid little attention to politics would, in consequence, have little “inertial” basis for resisting the new information they encountered in presidential campaigns.

What has not been clear in Converse’s argument – or in other work, including my own, that has built on it – has been whether those less involved voters who are prone to instability actually understand the messages that move them. Are they, in other words, simply pushed about by whichever communication flow is more intense, or do they react in meaningful fashion to the communication they receive? The analysis in this paper strongly suggests the latter: All modern campaigns are high intensity affairs, but low information voters discriminate among messages on the basis of each year’s content.

The classic political behavior studies have been called “minimalist” because they emphasize the extent to which information and political substance are absent in the political behavior of the electorate (Sniderman, 1993). The “new look” in political behavior research has, as Sniderman elaborates it, looked for ways in which, despite lack of information, citizens manage to make sense of politics and behave reasonably. Although I have taken pains to show the continuity of the findings in this paper with the arguments of the classic Michigan work, this paper would seem to fit comfortably within the new look perspective.

But in a critique of some of the new scholarship, Bartels (1996) asks whether, despite reliance on heuristics and cues that might help them over come their informational deficiencies, less informed voters really manage to do so. He finds that they do not: All else equal, women who are poorly informed about politics do not vote in presidential elections in the same way that better informed women do. Poorly informed Catholics do not, all else equal, vote the way well-informed Catholics do. For these and other groups, Bartels' evidence is compelling.

In a non-trivial way, this paper supports the views of both Bartels and Sniderman. It obviously finds, as does Bartels, that poorly informed voters differ from well-informed voters. But it also finds that low information voters know enough to respond to their political environment in a meaningful fashion. Indeed, to the extent that it is rational to vote retrospectively on the basis of national conditions and the relative ideological location of candidates, the present paper could be read as demonstrating that low information voters are actually more politically rational than any other segment of the electorate. I do not, however, make this argument, since it is difficult to say what exactly rationality consists of. One can plausibly argue that it is rational to ignore short-term factors and to vote always on the basis of party principle and ideology. Thus, the low information Democrats who defected to Nixon in 1972 because the economy was good, Nixon's foreign policy seemed successful, and McGovern was extreme might actually have been more rational to stick it out with the Democratic party, as their better informed co-partisans more often did. Ditto with respect to the poorly informed Republicans who defected to Johnson in 1964. But, however this may be, it seems clear that the presidential ballots of low information voters are grounded in one plausible form of political rationality.

A final question is whether presidential campaigns "educate" citizen-voters, as Gelman and King (1993) suggest, or "mobilize" partisan voters, as Berleson et al. (1954) have contended. Here my observations are quite tentative. I noted earlier that the Gallup data on vote intentions betray little evidence of education and some evidence of the opposite – voters becoming less sensitive to the ideological location of the candidates over the course of the campaign.

Frail as the data in Figure 12 are, they may shed just a bit more light on this question. For low information voters, economic performance seems to have more impact on vote intentions as the campaign progress.[58] Could this be evidence of an education effect? At the same time, economic performance has less impact on the vote intentions of the highly informed on election day than at the beginning of the campaign.[59] Indeed, all of the aggregate determinants of the vote in Figure 12 seem to have less impact the ballot choices of the highly informed than they do on their vote intentions at the start of the campaign. The odds of this occurring by chance alone are $(.5)^4 = .0625$. Are these changes perhaps evidence that high information voters are undergoing partisan mobilization?

It is not implausible that the campaign could bring about both education for low information voters (who stand in need of education) and partisan mobilization for high information voters (who are more partisan to begin with). But what about the finding, also in Figure 12, that low information voters become less responsive to Candidate Moderation as the campaign progresses? Do low information voters experience "education" with respect to economic conditions and "partisan mobilization" with respect to candidate qualities?

If the present data could readily answer these questions, I would have reported them. The best I can do is to offer methodological prescriptions on how to go about finding the answers, as I do in the next section.

58 This tendency has borderline statistical significance. If the model of vote intentions in Table 4 is run within the low information group – which means dropping the information interactions – an interaction between time of interview and economic performance gets two-tailed p-values of .15, .06, and .05 for the Hibbs, NBI, and RDI measures.

59 Using the same model as in the previous note for the high information group, this same interaction has p-values of .16, .24, and .19 – except with a sign to indicate the decreasing impact of economic conditions over the course of the campaign.

VI. Why Previous Studies Have given Different Results

The reader has reasons to be cautious about the findings reported in this paper. At the head of the list of reasons are over-report for the winner in one of the key dependent variables and the slipperiness of “intentions” in the other; the difficulty of measuring ideology, as exemplified by disagreement over whether Carter is relatively moderate or relatively extreme; and inherent power limitations in estimating interactions in a dataset that, in a real sense, has only 14 observations. It will take more elections – especially elections featuring extremists and salient foreign policy events – before we can be entirely sure about information effects in U.S. presidential elections. Yet, despite good reasons for caution, the available data do point fairly consistently to the conclusion that low information voters are, in important respects, as responsive to the political content of presidential elections as high information voters and quite possibly more so. And certainly the paper has found no evidence that low information voters are, in any general way, less responsive to the political content of elections.

Even the second and more cautious of these conclusions stands in stark contrast to the bulk of existing research, which has repeatedly found high information voters most responsive to the elements of national politics, including certain dynamic elements. Even researchers neither wanting nor expecting not to find this pattern have found it. For example, Fiorina (1981) writes in his revisionist account of voting behavior in presidential and congressional elections:

> Previous studies have consistently found that high education, interest, and information accompany increased reliance on issues when people vote. I had expected that the lack of such civic virtues would predispose people toward a heavy reliance on retrospective evaluations. As shown, the data fall in between but with a noticeable learn toward the traditional findings (p. 54).

In another “new look” study, Sniderman, Glaser and Griffin (1991) find evidence that poorly educated voters rely more on retrospective cues in the 1980 election, but no evidence of policy voting by less educated voters. The latter finding is despite the ideological gulf between Carter and Reagan that, in my analysis, generates support for the more moderate candidate among low information voters.

The most comprehensive study of this matter is perhaps that of Delli Carpini and Keeter (1996). They report a litany of shortcomings of the less informed, including the following finding concerning participation in presidential elections:

> Poorly informed citizens are less likely to participate, and when they do participate, they are less likely to tie their actions effectively to the issues stands and political orientations they profess to hold (p. 265).

In the Delli Carpini and Keeter study, low information voters cannot even link their retrospective revaluations to vote choice.

The question addressed in this final section is why the present study has generated such different results than most of the rest of the literature.

A big part of the answer is that most research examines one election at a time, and that, within any single election study, it is difficult to get a clear fix on the effects of factors, such as economic performance and relative moderation, that vary across elections. Researchers may try, but as I shall show in a moment, it is easy to get misleading results. What comes through instead from single-election studies is a picture of how voters respond to the particular set of candidates in that election. And the main finding in this regard is that high information voters are, indeed, more "ideologically consistent" in their reactions to candidates than low information voters. This persistent finding has given rise to the view that highly informed voters are, in general, more responsive to the elements of national politics than low information voters.

But it is easy to see in cross-election studies a pattern of responsiveness that is quite different: The candidates vary, but the choices of the high information voters remain relatively stable. Now the low information voters are the most responsive. Probably just because these voters are less partisan and especially less ideologically consistent, they can be more responsive to the political factors that change between elections. This greater responsiveness of low information voters is not a logical entailment of their tendency to be less ideologically constrained, but it is a nice complement to it.

Yet reliance on cross-election data to capture effects of aggregate variables is not the whole explanation for the different conclusions reached in this paper. Single election studies often do attempt to capture the effect of such aggregate factors as economic performance on vote choice. They do so by means of questions which ask voters their perception of national economic conditions. Thus, as several studies have shown, voters who perceive the economy to have improved over the last year tend to vote for the incumbent party, while voters who think conditions have worsened tend to vote against the incumbent.

This approach, however, fails to take account that, for many voters, ideological consistency extends to perceptions of economic performance (Wlezien, Franklin, and Twiggs, 1997). That is, many voters take a partisan view of economic performance, such that conditions are more likely to be improving when their party is in power and more likely to be worsening when it is out of power. Table 5 shows how such partisan rationalization looks in two election studies. When, in 1980, a Democratic incumbent was presiding over a weak economy, self-identified Republicans were more likely to say that the economy was getting worse. But when, in 1992, a Republican president was presiding over a weak economy, Democrats were more likely to assert that the economy was getting worse.

INSERT TABLE 5 ABOUT HERE

This form of biased perception appears in each NES presidential study since 1980. Moreover, the size of the bias toward partisan consistency in perception of the economy is, as past research on ideological consistency would suggest, nearly twice as large among voters scoring high on political information as among voters scoring low.

When, as thus described, party and information jointly drive economic perceptions, and when economic perceptions are in turn used to predict vote choice, it makes it seem that high information voters are voting on the basis of economic performance. This impression is false, however. What high information voters are doing is, to a considerable extent, rationalizing their votes in terms of economic perceptions. The tendency toward rationalization cannot be clearly detected within a single election study, but it is hard to miss when results are compared across several studies (see Duch, Palmer and Anderson, 2000).

The evidence for these assertions is set out in Table 6 and Figure 13. In Table 6, the dependent variable is perception that the economy is doing better or worse. The independent variables are party identification (scored in the direction of the incumbent administration), political information, and the interaction between party ID and information. Year dummies have been included to control for correlated errors.[62] In Table 6, the coefficient for the effect of party on perceptions of the economy is .19 and the coefficient for the interactive effect of party X information is .05. The p-values associated with these coefficients leave little doubt that an interaction between party and information affects perceptions of the economy. The magnitude of the effect is, moreover, quite large. Given that information is scored on a scale from 1 to 5, the total effect of party on economic perceptions for low information respondents is .19 + (1 X .05) = .24; for high information respondents, the total impact of party is .19 + (5 X .05) = .44. Thus, party has nearly twice the impact on perceptions of economic performance among high information voters compared to low information voters.

INSERT TABLE 6 AND FIGURE 13 ABOUT HERE

Given that high information voters are more likely to form “ideologically consistent” perceptions of economic performance, it is no surprise that the correlation between perceptions of performance and vote choice is higher for them than for low information voters. Figure 13 provides a graphical illustration of this tendency from voting in the 1992 election.

[62] The dependent variable is scored from –2 to +2. To avoid the possible endogeneity of party identification, I used only the initial party question (and not the follow-ups about strength of attachment and leaning). Partisans of the presidential outparty are scored –1, partisans of the inparty were scored +1, and others were scored 0.

These results, along with results from the previous section, show that the voting behavior of high information voters is not actually more sensitive to economic performance; it is simply more likely to be rationalized in terms of economic performance.[63]

Generalizing this line of argument, I would say that another reason that this paper reaches different conclusions about low information voters is that it avoids potentially endogenous attitude statements as explanatory variables, preferring to use aggregate-level variables in their place. Thus, instead of asking respondents their perceptions of the economy, this paper uses a measure of actual economic performance. Instead of asking respondents how the president has managed foreign affairs, it asks a panel of international relations scholars. Instead of asking individual respondents to offer their perceptions of the spatial location of candidates, the paper uses aggregate estimates of ideological location. The paper does use party identification as a control variable in its individual analysis, but it takes pains to show that the findings hold up in aggregate analysis when partisanship is measured as a mere dummy variable that allows different responses to Democratic incumbents than to Republican ones. The only independent variable in this paper that has always been measured in the traditional, survey-based manner is political information.

Another value of using aggregate-level measurements as explanatory variables is to equalize and probably to reduce random measurement error. It is well-known that measurement error in independent variables attenuates the impact of those variables, and the attitude statements of low information respondents have often been shown to be more prone to random measurement error (Feldman, 1989; Zaller, 1990). Use of external measures of key variables will at least

[63] Information in Figure 13 is measured additively by means of seven dichotomous information items, v5915 to v5921 in the NES codebook. The economic perceptions item is v3531. The information X perceptions interaction term is statistically significant at $p < .01$ in a logit analysis that also includes main effect terms for information and perceptions. If party identification is added to this model, the magnitude of the interaction drops by about 15 percent and its p-value rises to about $p = .03$, two-tailed. Additional controls control variables would further weaken the impact of the interaction term, but not my point, which is that high information respondents are more apt to rationalize vote their decisions in the politically conventional manner.

equalize measurement error across information groups and probably reduce its absolute magnitude as well.

I doubt that many readers would at this point challenge the value of cross-election data for understanding the effects of aggregate determinates of the presidential vote. But the same problems – endogenous independent variables and higher levels of measurement error among the less informed -- arise in a great many contexts in which research on public opinion is done, thereby creating the same need for aggregate-level independent variables. Experiments, in which all respondents get the same treatment, are one solution to these problems. Time series analysis, in which all respondents experience the same temporal events, is another. Quasi-experimental designs, as in the fine Johnston et al. (1992) study of the 1987 Canadian election, are also helpful. And when out-of-survey independent variables are simply unavailable, instrumental variables and panel data can be used both to reduce endogeneity and to equalize random measurement error. All of this is well-known. But for some 25 years, endogenous, error-laden survey measures have been giving what now seems a misleading if not wrong answer to the important and well-researched question I have examined in this paper. In light of this, the standard prescription should be taken more seriously.

Altogether, then, the results presented in this paper challenge more than the widespread notion that low information voters are less responsive to national political forces than high information voters. They also raise questions about the methodology on which this conventional view is based. More specifically, the results suggest that survey researchers should use more caution than they customarily do in deploying individual-level opinion statements as independent variables. The analyst may simply be unable to tell what is causing what. The danger in ignoring this admonishment is that our research may become better at explaining "survey behavior" than "political behavior".

REFERENCES

Aldrich, John H. 1995. *Why Parties?* University of Chicago Press.

Aldrich , John H., John L. Sullivan, Eugene Borgida 1989. “Foreign Affairs and Issue Voting: Do Presidential Candidates "Waltz Before A Blind Audience?" *American Political Science Review*, 83: 123-141.

Bartels, Larry M. 1996 “Uninformed Votes: Information Effects in Presidential Elections.” *American Journal of Political Science*, 40: 194-230.

Bartels, Larry M. 1997. “Specification Uncertainty and Model Averaging.” *American Journal of Political Science* 41: 641-74.

Bartels, Larry M. and John Zaller 2001. "Presidential Vote Models: A Recount." *PS.*

Berelson, Bernard, Paul F. Lazarsfeld, William McFee 1954*. Voting: A Study of Opinion Formation in a Presidential Campaign*. University of Chicago Press.

Box-Steffensmeier, Janet M. and Suzanna De Boef 2001. “Macropartisanship and Macroideology in the Sophisticated Electorate.” *Journal of Politcs*. 63:232-48.

Brady, Henry E. and Paul S. Sniderman 1985. “Attitude Attribution: A Group Basis for Political Reasoning.” *American Political Science Review*. 79: 1061-78.

Brody, Richard A. 1991. “Stability and Change in Party Identification: Presidential of Off-years. P. 179-205, in Sniderman et al., 1991.

Budge, Ian, Hans-Dieter Klingeman, Andrea Volkins, and Jdith Bara 2001. *Mapping Policy Preferences – Estimates for Parties, Electors, and Governments, 1945-1998*. New York: Cambridge University Press.

Converse, Philip 1964. Nature of Belief Systems in Mass Publics. In David Apter (ed.) *Ideology and Discontent*. New York: Free Press.

Converse , Philip 1962. “Information Flow and the Stability of Partisan Attitudes.” *Public Opinion Quarterly*. 26: 578-89.

Campbell, Angus, Philip Converse, Warren Miller, Donald E. Stokes 1960. *The American Voter*. New York: Wiley.

Daudt, H. 1961. *Floating Voters and the Floating Vote: A Critical Analysis of American and English Election Studies*. Leiden, Holland: Kroese.

Dryer, Edward C. 1971. “Media Use and Electoral Choices: Some Political Consequences of Information Exposure.” *Public Opinion Quarterly*. 35: 544-53.

Duch, Raymond M., Harvey D. Palmer, Christopher J. Anderson 2000. “Heterogeneity in Percpetions of Ntional Economic Conditions.” *American Journal of Political Science*. 44:635:52.

Delli Carpini, Michael X, and Scott Keeter 1996. *What Americans Know About Politics and Why It Matters.”* New Haven: Yale University Press.

Downs, Anthony 1957. *An Economic Theory of Democracy*. New York: Harper and Row.

Erikson, Robert S., Michael B. MacKuen, and James A. Stimson 2001. *The Macro Polity*. Cambridge UK: Cambridge University Press.

Erikson, Robert S. and Christoper Wlezien, 1999. Presidential Polls as a Time Series: The Case of 1996." *Public Opinion Quarterly*, 63: 163-177.

Fearon, James D. 1994. "Domestic Political Audiences and the Escalation of International Conflict." *American Political Science Review*, 88: 577-592.

Geddes, Barbara (in press). *Sandcastles and Paradigms*. Berkeley. University of California Press.

Free, Lloyd S and Hadley Cantril 1967. *The Political Beliefs of Americans*. New Brunswick: Rutgers University Press.

Fiorina, Morris P. 1981. *Retrospective Voting In American National Elections*. New Haven: Yale University Press.

Gelman, Andrew and Gary King 1993. "Why Are American Presidential Election Campaign Polls So Variable When Elections Are So Predictable?" *British Journal of Political Science*. 23: 409-51.

Hibbs, Douglas A., Jr. 1987. *The American Political Economy: Macroeconomics and Electoral Politics.* Cambridge, MA: Harvard University Press.

Hibbs, Douglas A., Jr. 2000. "Bread and Peace Voting in U.S. Presidential Elections." *Public Choice 104*: 149-80.

Holsti, Ole R. 1996. *Public Opinion and American Foreign Policy*. Ann Arbor: University of Michigan Press.

Iyengar, Shanto and Donald R. Kinder, 1987. *News That Matters*. Chicago: University of Chicago Press.

Johnston Richard, Andre Blais, Henry Brady, and Jean Crete 1993. *Letting the People Decide*. Palo Alto: Stanford University Press.

Key, V. O., Jr., with Milton C. Cummings 1966. *The Responsible Electorate*. New York, Vintage.

Kramer, Gerald H. 1971 "Short-term Fluctuations in Voting Behavior, 1896-1964." *American Political Science Review*. 65:131:43.

Krause, George A. 1997. "Voters, Information Heterogeneity, and the Dynamics of Aggregate Economic Expectations." *American Journal of Political Science*. 41: 1170-1200.

Lewis-Beck, Michael S. 1990. *Economics and Elections: The Major Western Democracies*. Ann Arbor: University of Michigan Press.

Lupia, Arthur 1994. Shortcuts vs. Encyclopedias: Voting Behavior in California Insurance Reform Elections. *American Political Science Review*. 88:63-76.

Luskin, Robert C. "Measuring Political Sophistication." *American Journal of Political Science*. 31: 856-99.

MacKuen, Michael B., Robert S. Erikson, and James A. Stimson 1989. "Macropartisanship." *American Political Science Review*. 83: 1125-42.

Markus, Gregory B. 1988. "The Impact of Personal and National Economic Conditions on the Presidential Vote: A Pooled Cross-Sectional Analysis." *American Journal of Political Science.* 32: 137-54.

McCubbins, Mathew D. and Arthur Lupa 1998. *The Democratic Dilemma.* Cambridge UK.Cambridge University Press.

Moon, Woojin 2002. "Resource Constrained Electoral Competition: Theory, Tests, and Applications." Unpublished Ph.D. dissertation, UCLA.

Nadeau, Richard and Michael S. Lewis-Beck 2001. "National Economic Voting in U.S. Presidential Elections." *Journal of Politics*. 63: 159-81.

Popkin, Samuel L. 1991. *The Reasoning Voter*. Chicago: University of Chicago Press.

Poole, Keith T. 1998. "Recovering a Basic Space From a Set of Issue Scales." *American Journal of Political Science*. 42: 954-93.

Price, Vincent and John Zaller 1992. "Who Gets the News: Measuring Individual Differences in Likelihood of News Reception," *Public Opinion Quarterly.* 57: 133-64.

Rosenstone, Steven J. 1983. *Forecasting Presidential Elections*. New Haven, CT: Yale University Press.

Schuman, Howard, Charlotte Steeh, Lawrence Bobo, and Mria Krysan 1997. *Racial Attitudes in America*. Cambridge: Harvard University Press.

Schultz, Kenneth 1998. "Domestic Opposition and Signaling in International Crises." *American Political Science Review*, 92: 829-844.

Smith, Eric R. A. N. 1989. *The Unchanging American Voter*. Berkeley: University of California Press.

Sniderman, Paul M., Richard A. Brody, Philip E. Tetlock 1991. *Reasoning and Choice: Explorations in Political Psycholgy.* Cambridge UK*:* Cambridge University Press.

Sniderman, Paul M., James M. Glaser, Robert Griffin 1991. "Information and Electoral Choice," P. 164-78 in Sniderman et al., 1991.

Sniderman, Paul M. 1993. "The New Look in Public Opinion Research." In Ada W. Finifter, ed., *Political Science: The State of the Discipline*. Washington D.C.: American Political Science Association.

Sobel, Richard 2001. The Impact of Public Opinion on U.S. Foreign Policy since Vietnam: Restraining the Collusus. Oxford: Oxford University Press.

Stanley, Harold and Richard Niemi 1995. *Vital Statistics on American Politics*. Fifth Edition. Washington D.C.: CQ Press.

Stimson, James 1975. "Belief Systems: Constraint, Complexity, and the 1972." P. 138-59 in Richard Niemi and Herbert Weisberg, *Controversies in American Voting Behavior*. San Francisco: Freeman.

Stimson, James A. 1999. *Public Opinion in America*. Boulder: Westview.

Wlezien, Christoper 1995. "The Public as Thermostat: Dynamics of Preferences for Spending." *American Journal of Political Science*, 39: 981-1000.

Wlezien, Christopher, Mark N. Franklin, and Daniel Twiggs 1997. "Economic Perceptions and Vote Choice: Disentangling the Endogeneity." *Political Behavior*, 19: 7-17.

Wlezien, Christopher 2001. "On Forecasting the Presidential Vote," *PS* (March).

Wright, Gerald C. 1984. Convention paper.

Zaller, John 1992. *Nature and Origins of Mass Opinion.* Cambridge UK: Cambridge University Press.

Zaller, John 1996. "The Myth of Massive Media Impact Revived: New Support for a Discredited Idea." P. 17-79 in Diana Mutz, Richard Brody, and Paul Sniderman (eds.), *Political Persuasion and Attitude Change*. Ann Arbor: University of Michigan Press.

Zaller, John 1990. "Political Awareness, Elite Opinion Leadership, and the Mass Survey Response," *Social Cognition*. 8: 125-153.

Appendix A

Instructions to Expert Foreign Policy Coders

I would like your help in a study involving the conduct of American foreign policy in presidential election years. Specifically, I would like you to evaluate success or failure of the administration's conduct of foreign policy in defending or advancing American interests, as seen from the perspective of a person who follows events in the newspapers and judges them in light of mainstream values of his own day. Evaluations should be based on events in the 12 months or so preceding each presidential election, up to the time of the election.

Your evaluations should reflect the gravity of the American interests that were engaged by our foreign policy. To take a simplistic example: A president who brilliantly represents the U.S. at diplomatic ceremonies should be judged less successful than a president who does a merely "good" job of saving the country from nuclear annihilation.

If a presidential administration has been more successful in some domains of foreign policy than others, your judgments should reflect a balance based on two things: Matters seen as more important at the time should be given somewhat more weight, and matters occurring in closer proximity to election day should be given somewhat more weight.

I am providing you with a 7-point scale, running from -3 to +3, for the ratings; however, if you wish to make in-between ratings, such as +2.5, please do so.

Appendix B
Some Regression Coefficients for Figure 11

	Level of Political Information				
	Low	IM	Middle	HM	High
RDI Q12-Q15	3.15	3.89	3.07	2.51	-0.43
one-tailed p=	.01	.00	.00	.00	.34
Consensus foreign policy	6.82	1.29	4.03	3.17	3.14
p=	.01	.30	.04	.03	.10
Party fatigue	-0.72	3.06	0.05	-0.02	-5.88
p=	.34	.05	.49	.49	.00
Candidate moderation	8.50	13.36	2.88	3.90	-10.12
p=	.04	.01	.22	.09	.02
deminc	10.09	4.22	1.26	-4.35	-5.92
constant	48.2	36.5	44.5	45.4	63.4
adj r2	.80	.62	.60	.86	.81
r2	.87	.77	.75	.91	.89
Bayes	2806	38	23	33609	5487
Hibbs measure	2.32	3.30	2.74	2.09	-0.13
one-tailed p=	.02	.00	.00	.00	.44
Expert foreign policy	7.01	1.45	4.13	3.28	3.08
p=	.01	.26	.02	.02	.10
Party fatigue	-2.55	1.01	-1.51	-1.37	-5.54
p=	.08	.24	.09	.08	.00
Candidate moderation	3.54	7.48	-1.70	0.07	-9.34
p=	.20	.03	.26	.49	.01
deminc	9.43	3.36	0.56	-4.90	-5.85
constant	54.2	42.5	48.7	49.4	61.9
adj r2	.75	.69	.73	.87	.81
r2	.85	.81	.83	.92	.88
Bayes	731	139	385	66723	4739
NBI-instrument	1.53	2.71	2.68	2.21	0.74
one-tailed p=	.09	.01	.00	.00	.18
Expert foreign policy	3.73	0.14	1.55	0.89	2.27
p=	.03	.46	.09	.13	.05
Party fatigue	-3.89	0.20	-2.30	-1.66	-5.79
p=	.03	.45	.04	.04	.00
Candidate moderation	5.04	15.10	2.44	5.41	-11.82
p=	.23	.01	.29	.05	.02
deminc	8.77	2.10	-0.65	-6.19	-5.71
constant	58.8	46.4	50.2	49.7	59.1
adj r2	.70	.60	.67	.90	.83
r2	.81	.75	.80	.94	.90
Bayes	173	24	92	408225	11361

Table 1. *Determinants of the Presidential Vote, 1948 to 2000*

	Hibbs economy			RDI Q12-Q15			NBI (instrument)		
Foreign policy scale:	Wars	Expert	Consensus	Wars	Expert	Consensus	Wars	Expert	Consensus
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Econ. performance	2.7	2.6	2.6	2.8	2.4	2.5	2.59	2.05	2.14
	1.08	.59	.76	.75	.52	.57	.50	.39	.45
one-sided p=	.02	.00	.01	.00	.00	.00	.00	.00	.00
Weight parameter	.74	.83	.80						
	.18	.09	.12						
p=	.00	.00	.00						
Foreign Policy	-5.1	2.2	3.4	-3.7	2.0	3.2	-5.6	1.9	2.9
	4.10	.79	1.76	3.27	.58	1.13	2.64	.52	1.08
p=	.13	.01	.05	.15	.00	.01	.03	.00	.01
Cand. Moderation	-1.4	-0.7	-0.6	3.5	3.1	3.3	2.8	2.4	2.6
	2.44	1.58	1.88	2.95	2.03	2.22	2.24	1.74	2.04
p=	.29	.34	.37	.14	.09	.09	.13	.10	.12
Party Moderation	2.6	2.0	2.4	2.7	2.6	3.0	-1.3	-0.9	-0.7
	2.10	1.50	1.78	2.52	1.72	1.91	1.92	1.50	1.78
p=	.12	.11	.11	.15	.09	.08	.27	.28	.34
Party incumbency	-1.3	-2.0	-1.8	0.5	-0.6	-0.2	-2.1	-3.0	-2.8
	1.58	.63	.83	1.33	.86	.93	.83	.59	.69
p=	.21	.01	.03	.36	.27	.42	.02	.00	.00
Constant	48.6	50.0	49.6	44.8	47.2	46.3	51.0	53.1	52.7
adjusted r-square	.73	.86	.80	.61	.82	.78	.76	.86	.80
Bayes factor	273	28908	2403	31	5823	1556	877	30714	3256
r-square	.86	.93	.89	.76	.89	.86	.85	.91	.88

Note: Dependent variable is vote for the incumbent presidential party in n=14 elections. The units of the Hibbs and RDI measures are "percent change in RDI per capita," with SDs of 1.4 and 1.8. The NBI, a survey-based measure, has been re-scaled to the units of RDI and has the same SD. The two moderation variables are scored in the metric of the NES 7-point ideology scales and have the same SD, 0.47. In columns 1, 4, and 7, Foreign Policy performance as measured as a dummy taking the values of 1 in 1952 and 1968, the election years following the start of the Korean and Vietnam wars, and 0 otherwise. In columns 2, 5, and 8, foreign policy performance is measured by means of the "Expert" scale, which is the average of nine judge ratings, as described in the text. This variable has a standard deviation of 1.2 In columns 3, 6, and 9, the foreign policy performance measure is an abbreviated version of the Expert scale, which takes values of -1 in 1952, 1968, and 1980, +1 in 1948, 1988, and 1992, and zero otherwise (SD = .68). Its SD is .68. Party Incumbency is scored 1 to 5 for the number of terms the incumbent party has held the presidency; its SD is 1.2.

Table 2. Bivariate correlations between incumbent vote and aggregate causes of the vote, by level of information, 1948 to 2000.

	Level of Political Information				
	Low		Middle		High
	(1)	(2)	(3)	(4)	(5)
Hibbs performance	0.57	0.76	0.81	0.77	0.18
one-tailed p =	.02	.00	.00	.00	.27
NBI	0.30	0.50	0.70	0.64	0.26
p=	.16	.04	.00	.01	.20
Candidate Moderation	0.43	0.42	0.09	0.23	-0.08
p=	.07	.08	.38	.23	.40
Party Moderation	0.15	0.15	-0.01	0.08	0.40
p=	.31	.31	.49	.40	.09
Expert foreign	0.48	0.28	0.38	0.34	0.14
policy	.05	.17	.10	.13	.33
Consensus	0.46	0.17	0.36	0.35	0.13
foreign policy	.06	.29	.11	.12	.33
Wars only	-0.51	-0.10	-0.09	-0.24	-0.34
	.04	.38	.39	.22	.13
Party fatigue	-0.51	-0.28	-0.35	-0.43	-0.59
	.04	.18	.12	.07	.02

Note: Upper cell entries are partial bivariate correlations, within each of five information groups, between incumbent vote and aggregate level determinates of the vote. Lower cell entries are one-tailed p-values. The partial correlation controls for incumbent party.

Table 3. Interactive effects of information on presidential vote, 1948 to 2000.

	Simple vote correction			"Over-corrected" vote		
	Hibbs (1)	NBI (2)	RDI (3)	Hibbs (4)	NBI (5)	RDI (6)
Economy X Information	-1.00	-1.25	-0.72	-0.87	-1.00	-0.46
SE	0.20	0.26	0.22	0.20	0.26	0.22
Consensus Foreign	-0.78	-0.98	-0.49	-0.29	-0.46	-0.14
X Information	0.64	0.64	0.66	0.64	0.64	0.66
Cand. Moderation X Information	-4.24	-6.34	-5.02	-3.56	-5.24	-4.01
se	0.90	1.01	0.94	0.90	1.01	0.94
Party fatigue X Information	-0.92	-1.67	-0.61	-1.28	-1.86	-1.00
se	0.34	0.41	0.34	0.35	0.41	0.34
Dem. Incumbent X Information	-4.30	-4.68	-4.02	-3.79	-4.09	-3.61
se	0.47	0.48	0.48	0.47	0.48	0.48
Incumbent Party ID (-1, 0, +1)	33.81	33.81	33.8	33.83	33.83	33.83
se	0.4	0.4	0.41	0.4	0.4	0.4
Political information	3.56	5.76	2.31	4.11	5.66	2.52
t-ratio	1.06	1.40	1.03	1.06	1.40	1.03
(Average = .22; SD = .94)						
Adjusted r-square	.39	.39	0.4	.39	.39	.39
SE of estimate	38.7	38.7	39	38.7	38.7	38.8
n=	1601 9	1601 9	1601 9	1601 9	1601 9	1601 9

Note: Dependent variable is "corrected" incumbent party vote; the scale before correction took values of 0 and 100. The party variable was on a scale of +1, -1, and 0 prior to the endogeneity correction. All regressions contain "fixed effect" dummy variables for each election year except one and other control variables, as explained in the text.

Table 4. Interactive effects of information on vote intentions in second half of political presidential campaigns, 1948 to 2000.

	Hibbs (1)	RDI (2)	NBI (3)
Economy X Information	-0.91	-1.03	-0.68
SE	0.27	0.33	0.28
Consensus Foreign	-0.81	-0.86	-0.55
X Information	0.76	0.76	0.78
Cand. Moderation X Information	-2.57	-4.26	-3.31
se	1.14	1.29	1.20
Party fatigue X Information	-0.24	-0.80	0.06
se	0.45	0.54	0.44
Dem. Incumbent X Information	-4.56	-4.81	-4.32
se	0.58	0.58	0.59
Incumbent Party ID (-1, 0, +1)	32.22	32.22	32.23
se	0.5	0.5	0.5
Political information	1.53	3.05	0.36
t-ratio	1.42	1.87	1.37
(Average = .22; SD = .94)			
Adjusted r-square	.42	.42	.42
SE of estimate	35.6	35.6	35.6
n=	8274	8274	8274

Note: Dependent variable is pre-election intention to vote for the incumbent party candidate, which takes value of 0, 50, and 100. Regressions are identical in form to those in Table 3.

Table 5. Perceptions that the economy has gotten "much worse" over last year by party attachment, 1980 and 1992

	Strong Republican	Weak & Leaning Republican	Pure Independent	Weak & Leaning Democrat	Strong Democrat
1980	61.3 (124)	49.1 (271)	38.3 (81)	35.6 (306)	31.2 (186)
1992	13.2 (220)	25.0 (444)	47.7 (149)	42.9 (508)	59.1 (337)

Note: Cell entries are percent of voters who say economy has gotten "much worse" over last year. All respondents, including DK respondents, are included in the denominator from which cell entries are calculated.

Source: National Election Study surveys.

Table 6. The effect of party bias on perception of economic performance, 1980 to 1996

	Slope (b)	S.E.	Beta	t-ratio	Two-tailed p-value
Incumbent Party ID (values: -1, 0, 1)	0.19	0.03	0.14	5.7	.000
Political information (range: 1 to 5)	.061	.008	0.08	7.4	.000
Information X Incumbent Party ID	0.05	0.01	0.12	5.05	.000
Year 1984 (range 0-1)	1.54	0.04	.58	40.4	.000
Year 1988	1.11	0.04	0.40	28.3	.000
Year 1992	.23	0.04	0.08	6.1	.000
Year 1996	1.60	0.04	0.54	39.8	.000

Adjusted r-square 0.40

N = 5844

Note: The dependent variable is an item asking whether the performance of the economy has gotten better or worse over the past year. It is scored from –2 (much worse) to +2 (much better). Results are based on major party voters only.

Source: National Election Studies

Figure 1. *Actual versus survey measure of incumbent party vote share, 1948 - 2000*

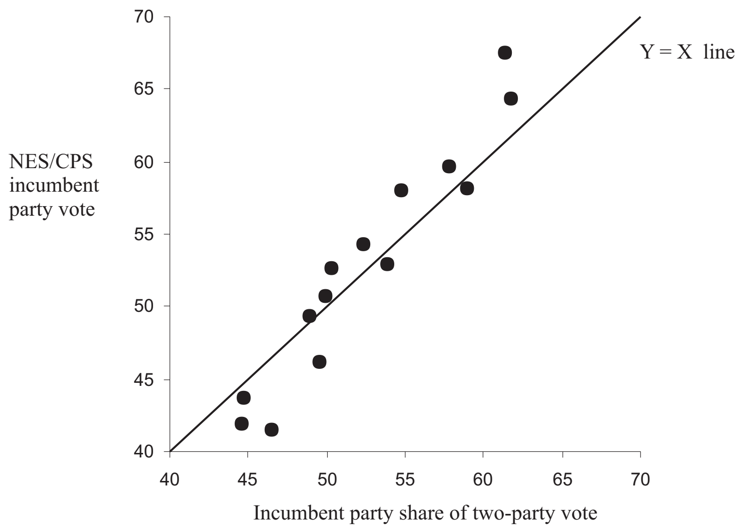

Figure 2. *Gallup versus CPS/NES measure of incumbent party vote share, 1948 - 2000*

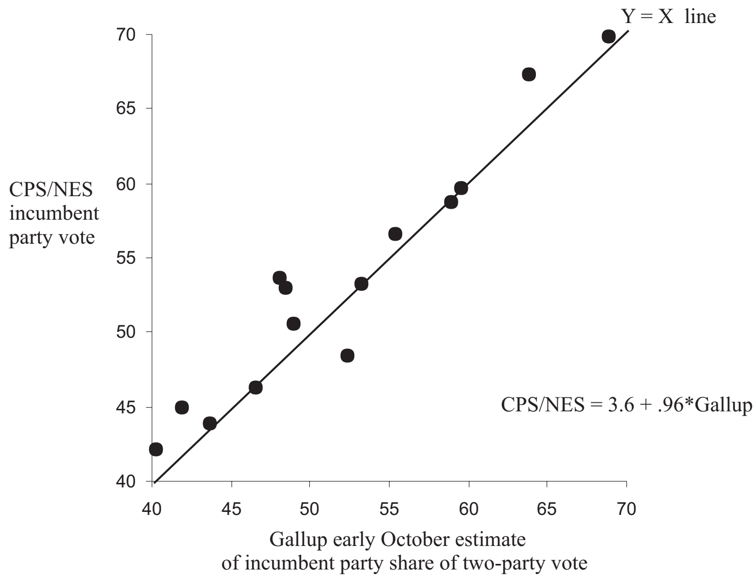

Figure 3. Spatial locations of candidates and parties, 1948 to 2000

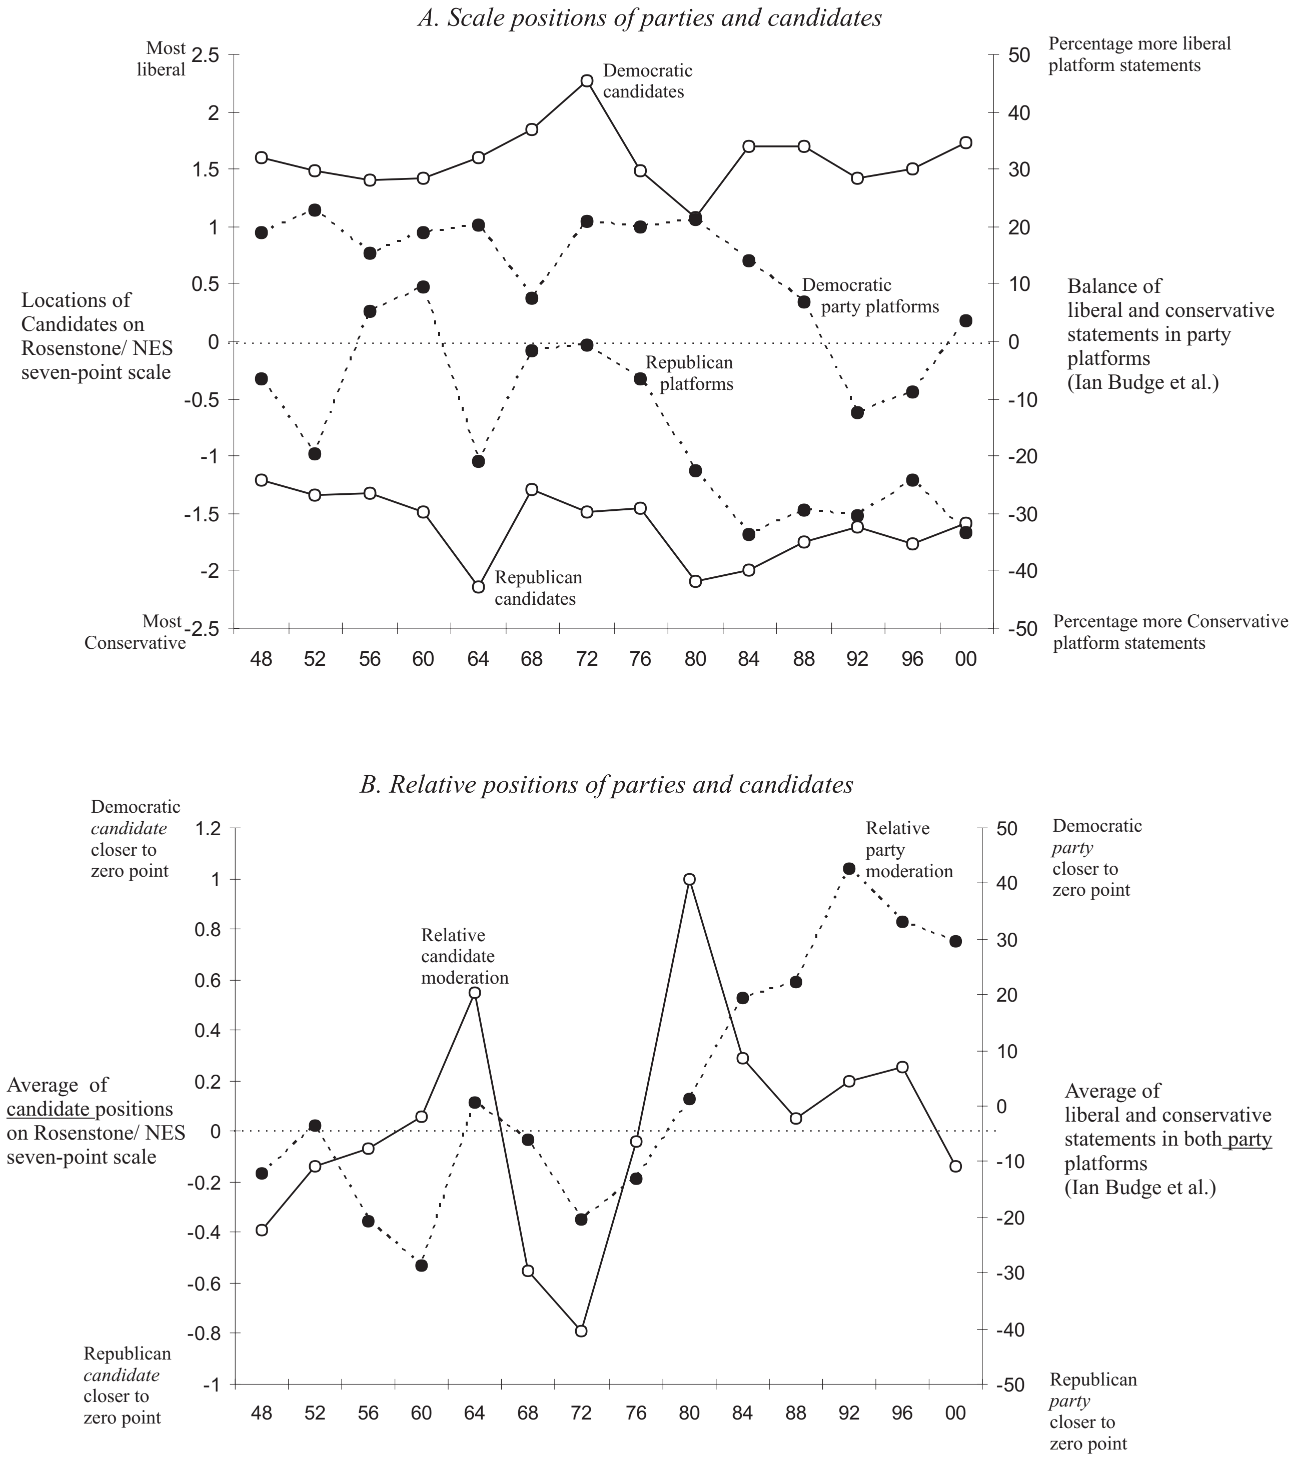

Figure 4. Two estimates of the ideological center-point of opinion, 1952 to 2000

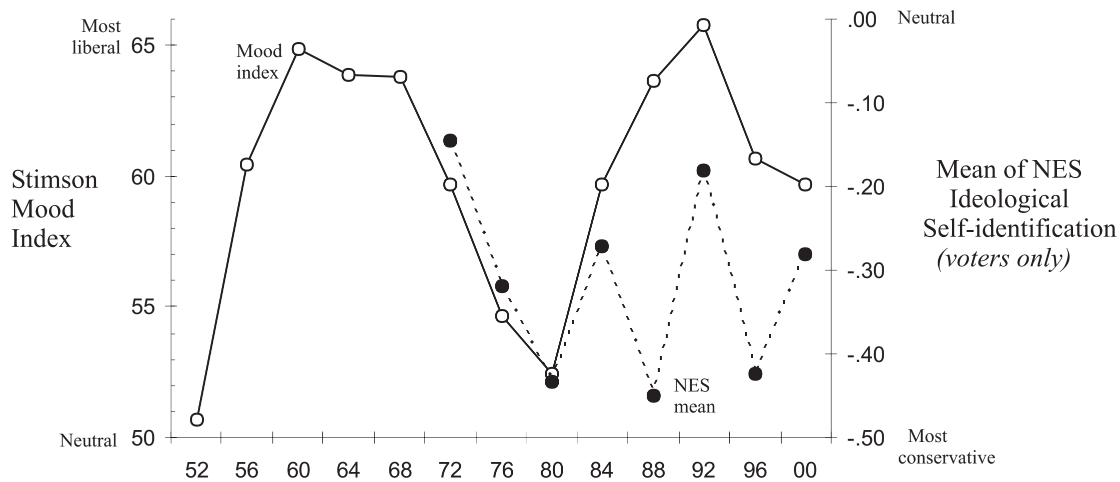

Figure 5. Expert ratings of presidential foreign policy success, 1948 to 2000

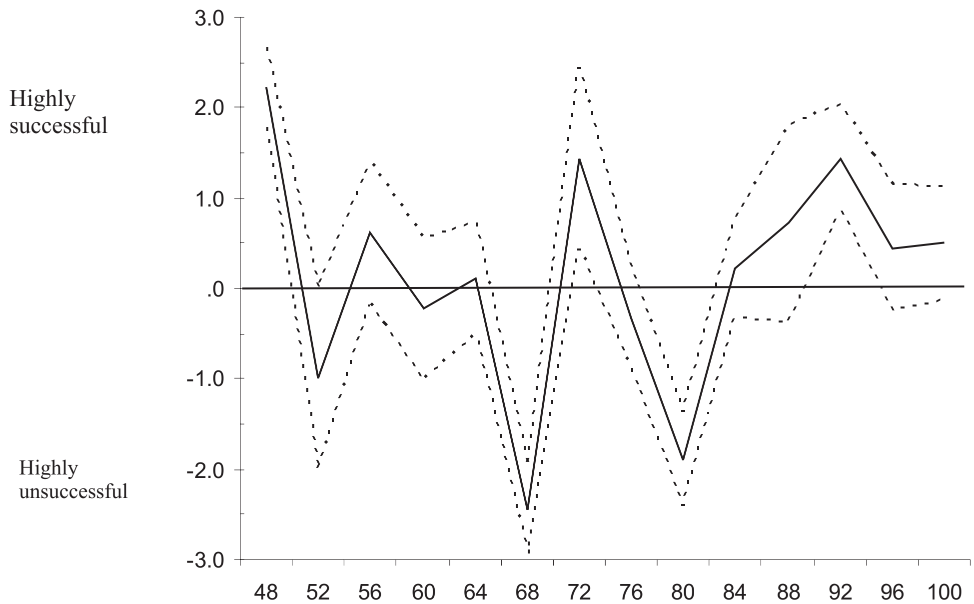

Note: Solid lines shows mean rating of nine judges; dotted line show +/- two standard errors from mean.

Figure 6. Swings in the presidential vote by political information, 1948 to 2000

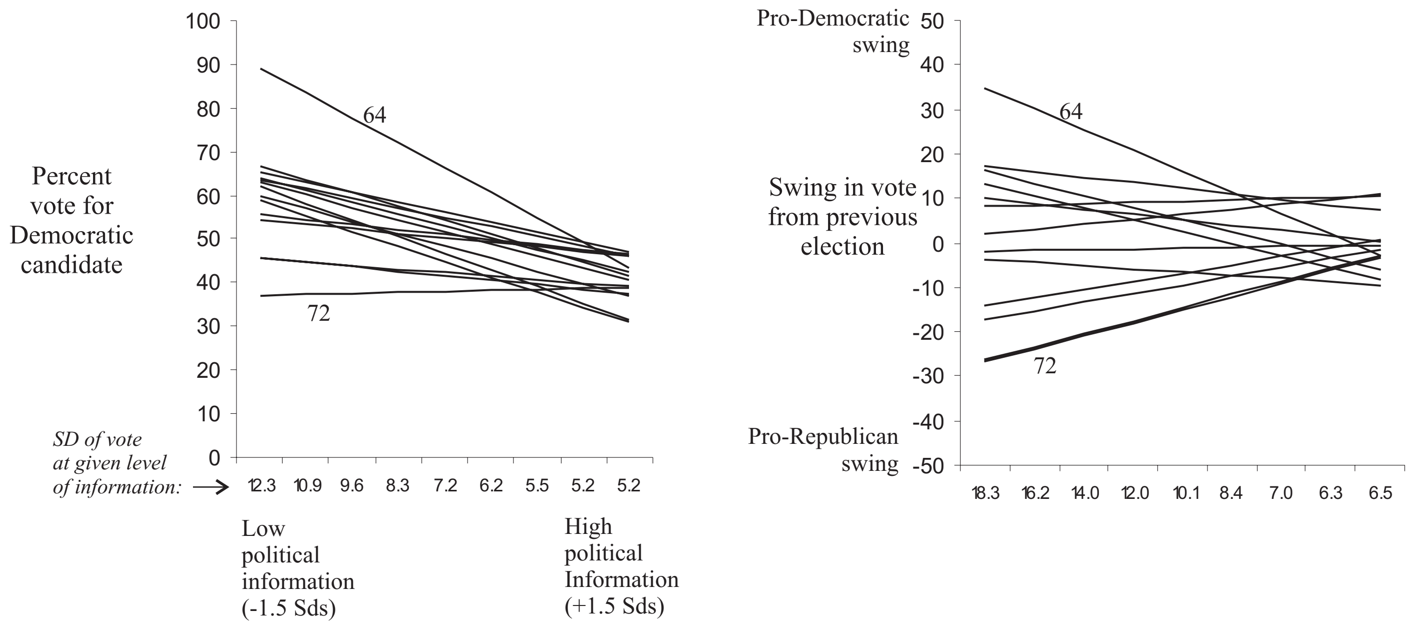

Figure 7. Length of party incumbent and out-party moderation

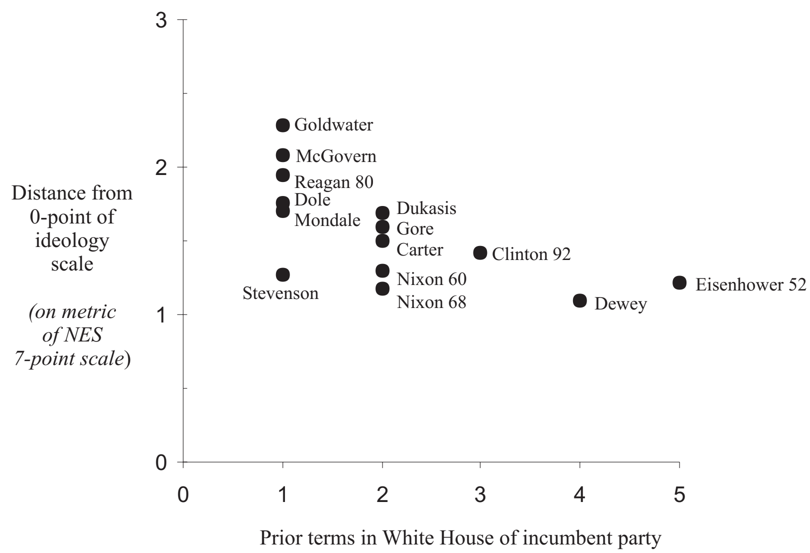

Figure 8. Changes in determinants of the vote during political campaign

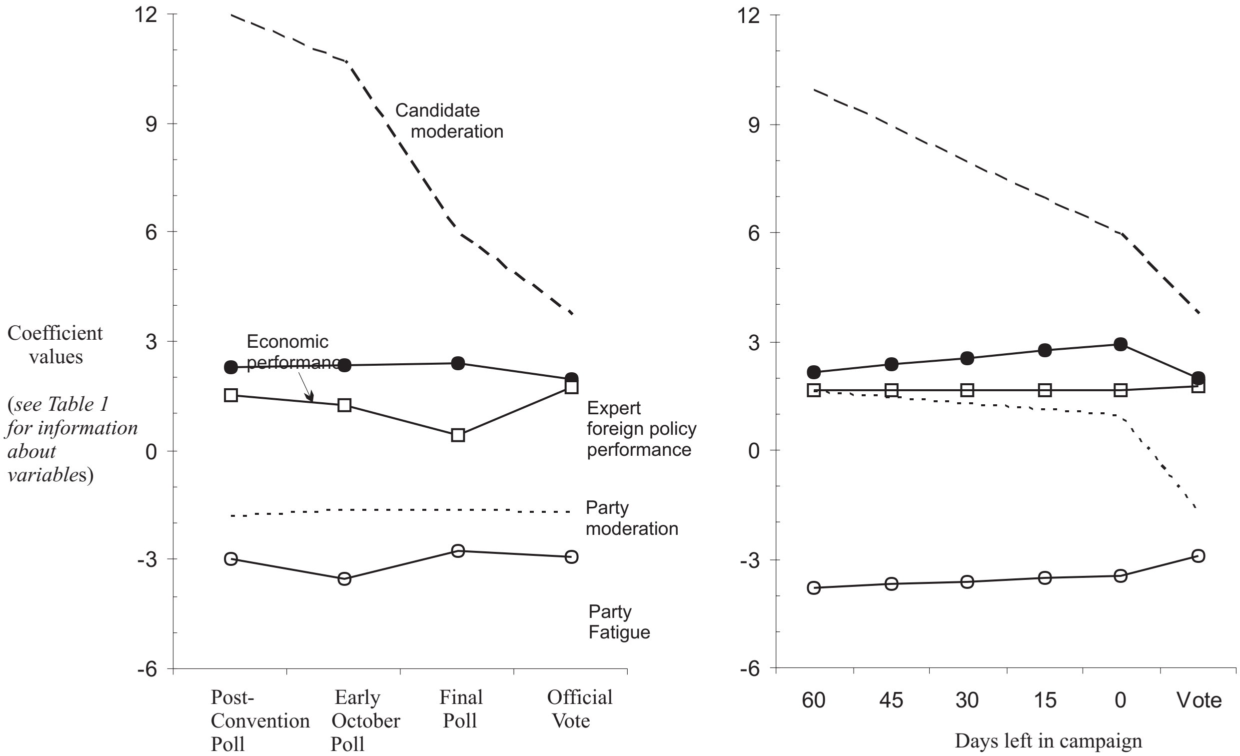

Note: All polls are Gallup polls, as compiled by Stanely and Niemi (1995) through 1992 and augmented by author in recent elections..

Figure 9. Scatterplot of economic performance versus incumbent vote, 1948 to 2000

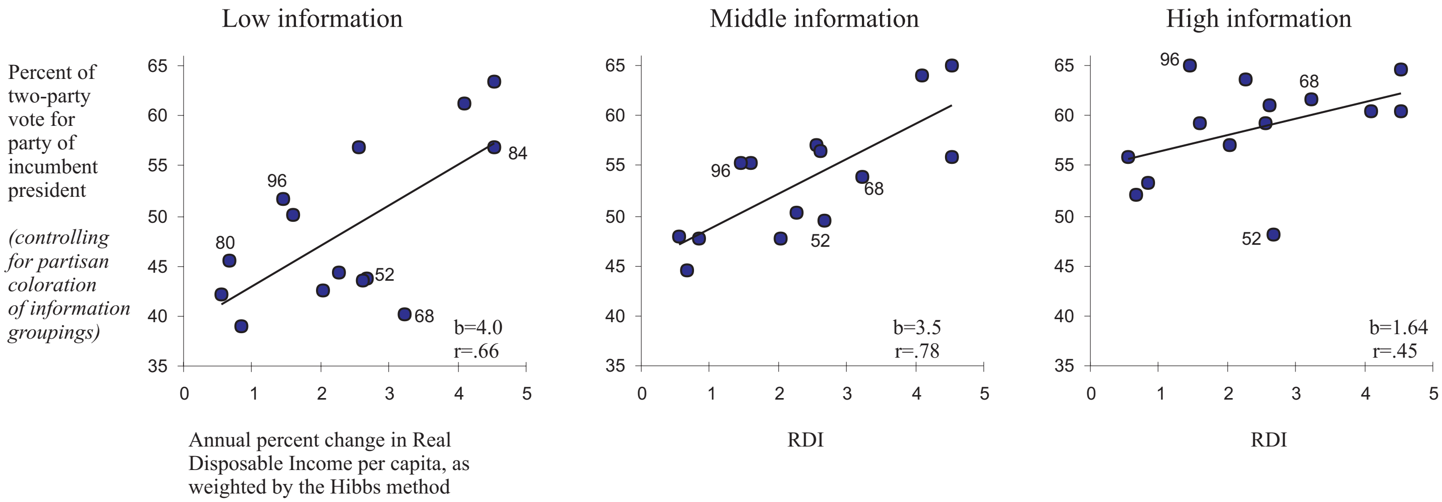

Figure10. Long-term partisan vote bias and political information, 1948 - 2000

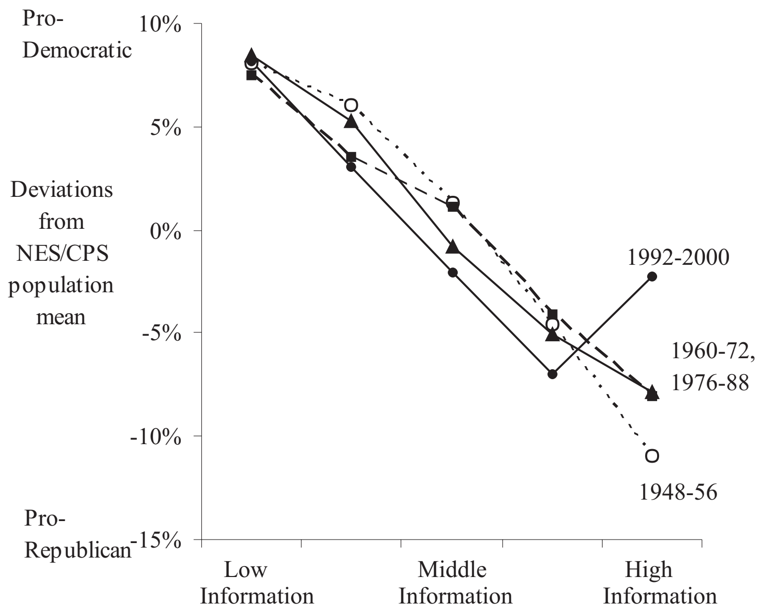

Figure 11. Determinates of the presidential vote by level of information, 1948 to 2000.

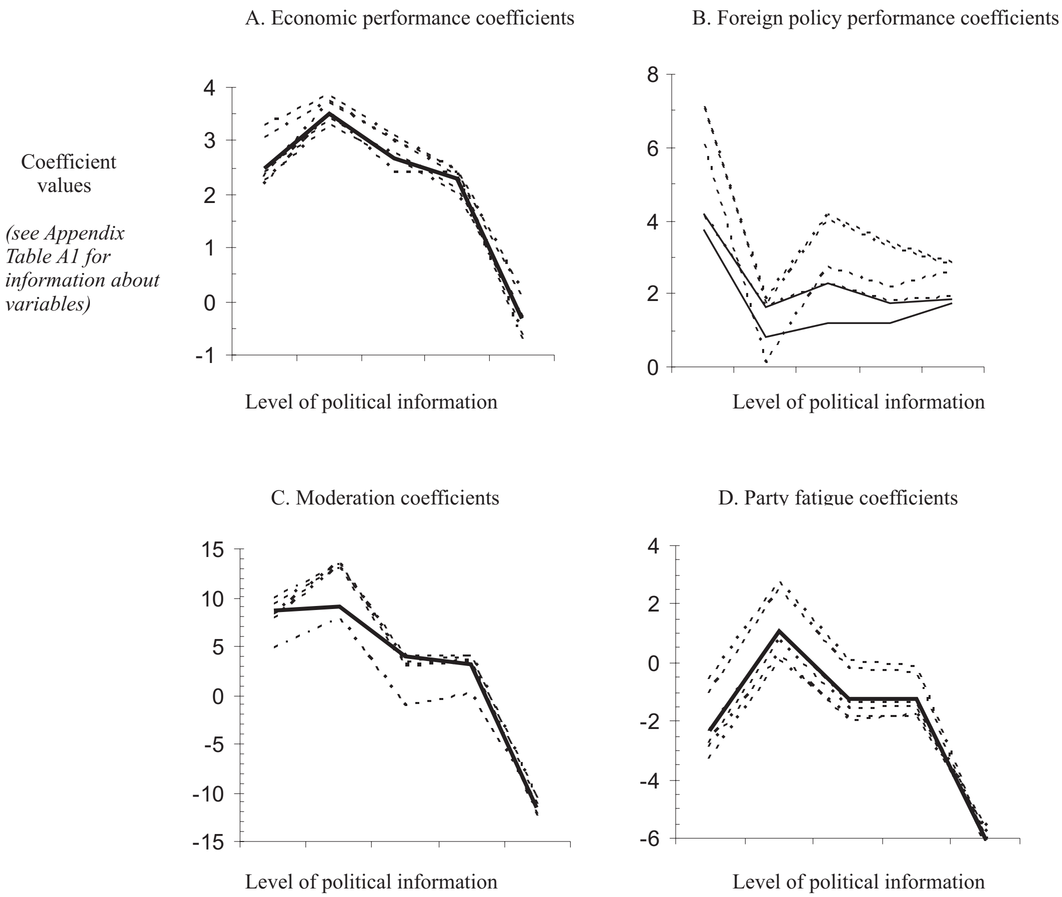

Figure 12. Impacts on the vote by level of information and time of campaign

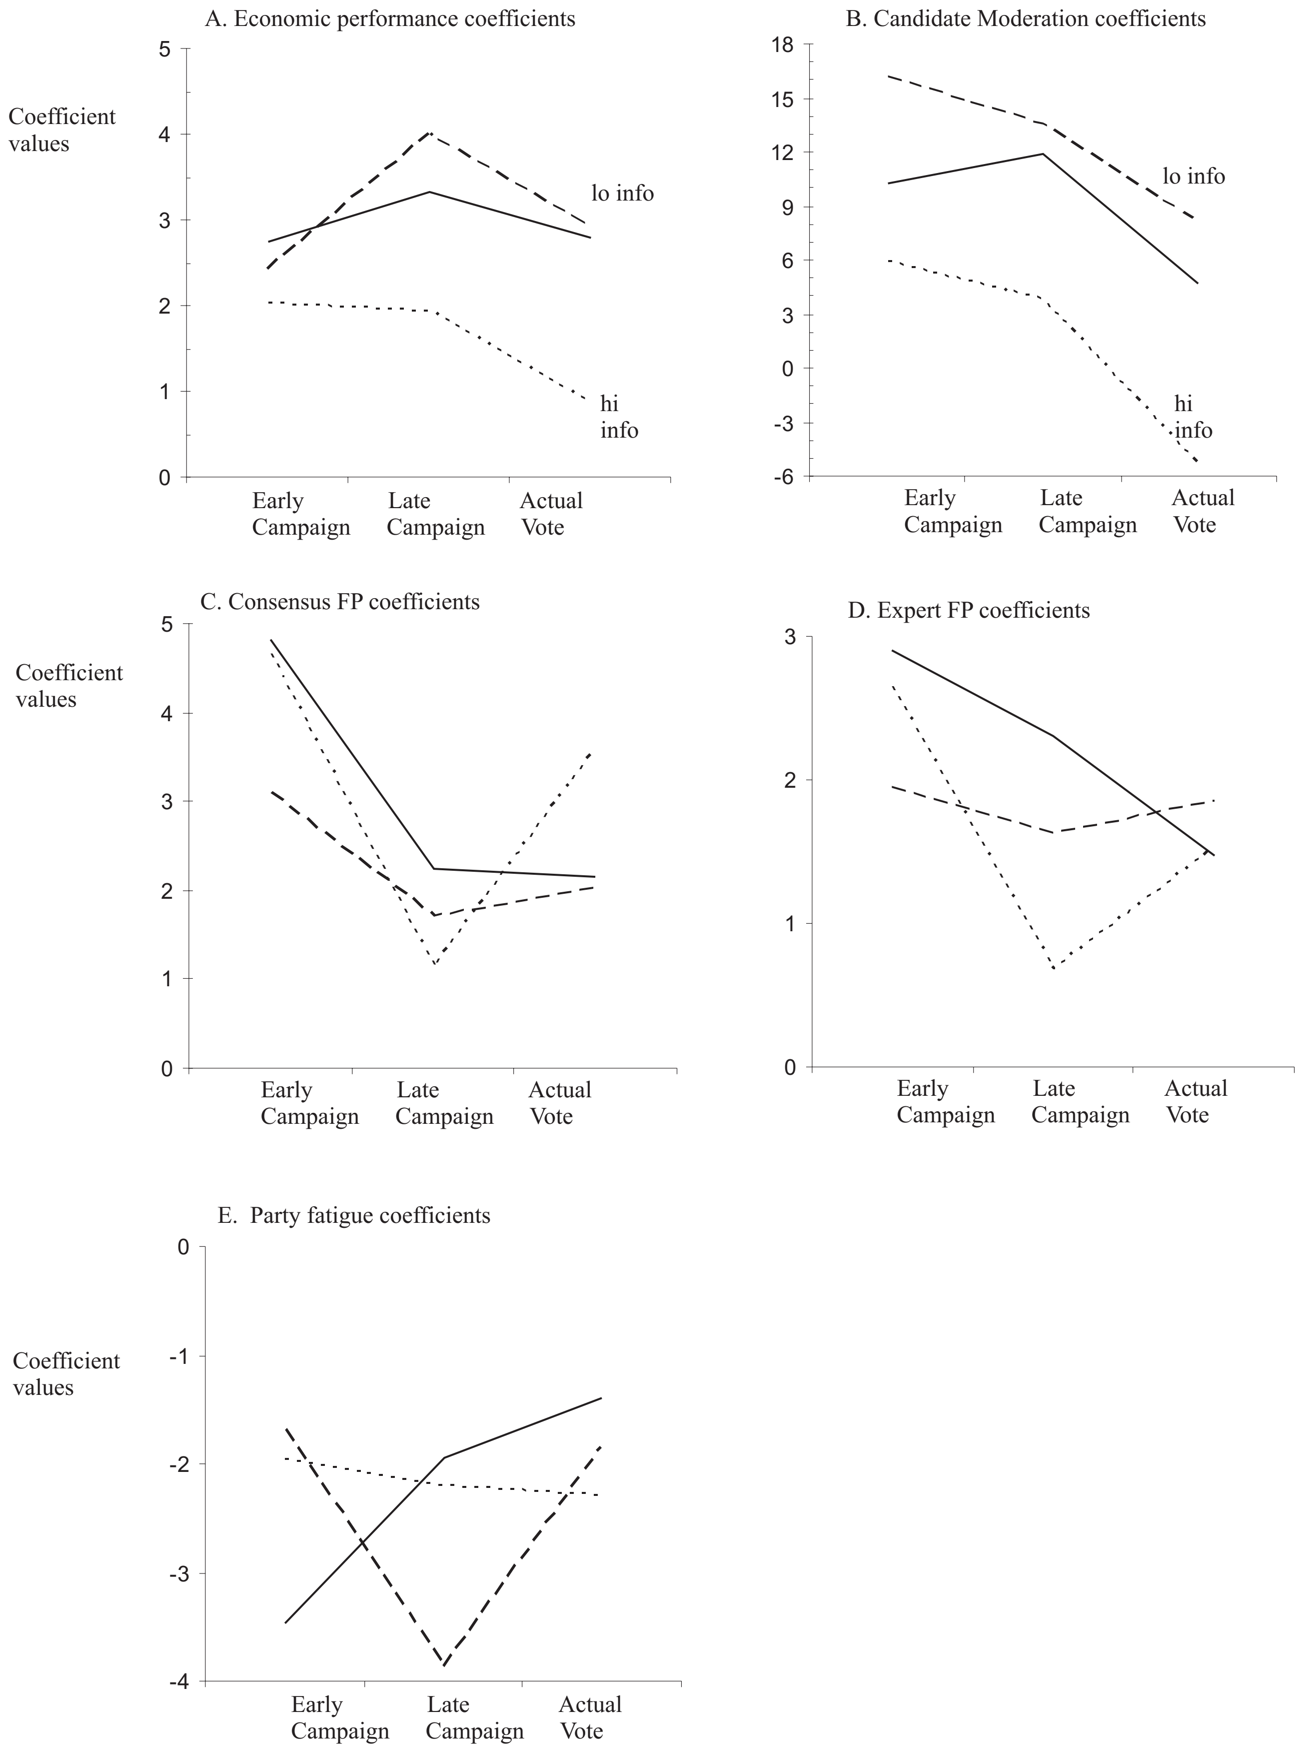

Figure 13. The effect of information and economic perceptions on vote for president in 1992

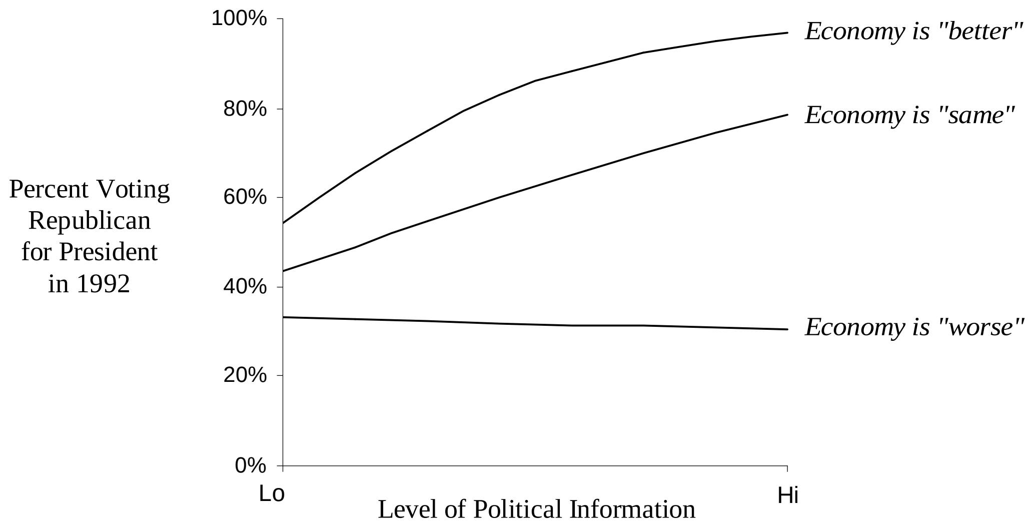

Note: Methodological information in note 17 of text.

Source: 1992 NES study