Are Voters Rational?
Evidence from Gubernatorial Elections*

Justin Wolfers†

The Wharton School, University of Pennsylvania,

CEPR; IZA and NBER

jwolfers@wharton.upenn.edu

www.nber.org/~jwolfers

Abstract

Standard agency theory suggests that voters will vote to re-elect politicians who deliver favorable outcomes. A second implication is that voters will not support a politician because of good outcomes unrelated to the politician's actions. Specifically, rational voters should try to filter signal from noise, both in order to avoid electing incompetent, but lucky politicians, and to maximize the link between their votes and optimal incentives. This paper provides insight into the information processing capacities of voters, by measuring the extent to which they (irrationally) reward state governors for economic fluctuations that are plausibly unrelated to gubernatorial actions. Simple tests of relative performance evaluation reveal that voters evaluate their state's economic performance relative to the national economy. However, these tests only provide evidence of rule-of-thumb performance filtering. More sophisticated tests reveal that voters in oil-producing states tend to re-elect incumbent governors during oil price rises, and vote them out of office when the oil price drops. Similarly, voters in pro-cyclical states are consistently fooled into re-electing incumbents during national booms, only to dump them during national recessions. Consistent with an emerging behavioral literature, this suggests that voters make systematic attribution errors and are best characterized as quasi-rational.

This Draft: January 30, 2007

---

* Thanks for helpful insights and data are due to: Alberto Alesina, James Alt, Tim Besley, Olivier Blanchard, Robert Barro, Anne Case, Alan Durell, Chris Foote, Ed Glaeser, Christopher Jencks, Larry Katz, Michael Knetter, David Laibson, Jeff Liebman, Greg Mankiw, Andrew Oswald, Sam Peltzman, Betsey Stevenson, and James Williams. Thanks also to seminar participants at Chicago, Columbia, Harvard, LSE, MIT, NYU, Princeton, Stanford, Sydney, and Wharton.

† Assistant Professor of Business and Public Policy, The Wharton School, Steinberg-Dietrich Hall #1456, 3620 Locust Walk, University of Pennsylvania, PA 19104. Phone: (215) 898-3013; Fax: (215) 873-0679.

I. Introduction

Although political economists have devised models of political business cycles under the contrasting assumptions of voter rationality and voter naïveté, there is not much direct empirical evidence on how well voters process information when making political choices. This paper provides a very simple and direct assessment of the ability of voters to distinguish signal from noise, analyzing data on all gubernatorial elections in all fifty U.S. states since World War II.

My assessment of the information-processing performed by voters derives from the agency problem that exists between voters and their elected officials: Voters are concerned with re-electing competent leaders, but cannot directly observe effort and ability. Standard principal-agent theory suggests that voters should base their vote on observable indicators of competence, including economic performance. If voters efficiently process this information they will reward good economic outcomes that reflect the governor's actions, *but filter from their assessment economic events that reflect influences outside the politician's locus of control*. This paper takes this prediction to the data, examining the extent to which voters (rationally) reward *competence* and (irrationally) reward *luck*. Assessing competence involves a standard signal extraction problem, and imperfect performance filtering reflects inefficient inference about gubernatorial ability. In turn, inefficient inferences raise the probability of re-electing incompetent, but lucky incumbents, and lower the probability of re-electing unlucky, but competent governors.[3] Responding to noise in addition to signal also weakens the link between a politician's actions and their electoral consequences, and thus dilutes political accountability and the power of incentives wielded by voters.

State economic outcomes are subject to important shocks caused by identifiably exogenous influences, and hence state gubernatorial elections provide an ideal testing ground for better understanding the information-processing ability of voters. To preview the results, I find evidence that voters make some attempt to decompose observed economic performance into results reflecting the competence of politicians and those reflecting other causes. Arguably, they are at least as successful at this task as corporate boards are when setting CEO compensation. Tests of relative performance evaluation reveal that voters appear to focus on the performance of the state economy relative to the national norm. However, instrumental variables tests of performance filtering reveal that certain fine-grained distinctions elude voters, and the incumbent's re-election prospects are systematically related to factors outside their control. For instance, governors in oil-producing states are likely to be re-elected following a rise in oil prices, while their counterparts in the rust-belt are likely to be ousted. That is, only approximate rule-of-thumb adjustments are made when evaluating whether to re-elect incumbents. This account of the limited ability of voters to discern signal from noise stands as a useful qualification to

3. In the words of Alesina, Roubini, and Cohen (1997, p.35), "a naïve voter would punish an unlucky incumbent, whereas a rational voter would, at least in part, take bad luck into consideration."

both purely behavioral theories in which voters are simply persuaded by campaigns, and the higher-order rationality assumed in rational expectations models. In related work, Achen and Bartels (2004) find that voters are more likely to oust incumbents following natural disasters, arguing that this is evidence that "retrospection is blind"; while my results are consistent in pointing to some irrational attribution in retrospective voting, voters also show partial success in not attributing blame for economic downturns to governors during national recessions.

Of course, the finding that voters do not efficiently process information may simply be a reflection of the costs of gathering and processing information. Even so, the question of whether voters make efficient inferences based on publicly available information is central to distinguishing between first and second-generation models that explain the incentive for pre-election policy manipulation.[4] In first generation political business cycle models, voters interpret a strong economy as evidence of competent leadership, but fail to distinguish signal from noise. This naïveté creates an incentive for incumbents to overheat the economy before an election in an attempt to fool voters into inferring that the pre-election boom reflects the incumbent government's competence. By contrast, in second-generation models, efficient inference leads voters to distinguish between good outcomes reflecting competence and good outcomes reflecting pre-election policy manipulation. Typically, it is assumed that competent politicians can manipulate policy relatively cheaply; thus, while a pre-election boom fools neither voters nor politicians, it may facilitate credible revelation that the incumbent is of high quality in a Spence-style signaling game.

Note that both types of models yield substantially similar reduced form implications—there may be an expansionary boom before an election, and incumbents are more likely to be re-elected following a pre-election boom. However, the welfare and policy implications are radically different. In first generation models, the incumbent's ability to generate a pre-election boom simply generates wasteful output volatility and insulates incumbents from electoral competition. Further, politicians motivated by self-interest are more likely to manipulate the electorate, and hence are more likely to be re-elected. By contrast, in second-generation models, "good" (competent) politicians generate pre-electoral booms, and because political business cycles facilitate credible revelation of the incumbent's quality, these political cycles may even enhance efficiency. Thus naïve voter models suggest policy prescriptions—such as budget rules and central bank independence—that may be entirely inappropriate if voters use efficiently exploit available information about candidate quality, thereby reducing the reasons for pre-election manipulation.

Rather than focusing on reduced form implications of specific models, this paper directly

---

4. These models attempt to explain the incentive for pre-election policy manipulation. First generation models include Nordhaus 1975 and Lindbeck 1976. Second generation models include Cukierman and Meltzer 1986,

implements a test of the ability of voters to distinguish signal from noise when assessing political leaders.

II. The Political Agency Problem

Our starting point is the political principal-agent problem, in which voters are principals and politicians are their agents. Agency problems arise because voters cannot observe all of the politician's actions and effort choice. This gives the politician scope to pursue his or her personal agenda rather than that of the constituents.

Results in the political agency literature largely shadow those in the standard principal-agent literature. Just as greater productivity can be elicited from workers by linking their wages with indicators that are correlated with their effort, voters can motivate a politician to better serve their will by linking votes to observable outcomes. Further, just as employers are concerned with retaining only high quality workers, so too voters are interested in re-electing competent governors. Poor outcomes are evidence of incompetence, and hence suggest that the agent (the governor) should be fired.

To provide a simple formalization, consider the attempts of voters to maximize the competence of their governors.[5] Political candidates have some ability (or exert some effort level), $a \sim N(\mu_a, \sigma_a^2)$. Voters observe state economic conditions, s, which is a noisy indicator of the incumbent's ability, since state economic conditions reflect not only the governor's efforts, but also three other orthogonal shocks: national economic conditions, n; a set of other observable shocks, x, which affect each state differently, and unobservable factors, $e \sim N(0, \sigma_e^2)$. Thus, having observed state economic conditions, $s=a+n+\beta_s x+e$, Bayes' Rule yields a simple expression for a voters' inference about the effort or ability of the incumbent, which is a precision-weighted average of their prior and the signal:

$$E[a|s,n,x] = [\mu_a \sigma_a^{-2} + (s-n-\beta_s x)\, \sigma_e^{-2}] \,/\, (\sigma_a^{-2}+\sigma_e^{-2}). \qquad (1)$$

Expectations about the challenger are simpler: With no track record she represents a fresh draw from the distribution of political candidates, and hence is expected to have ability level μ_a. Thus voters infer an ability differential between the candidates equal to $(s-n-\beta_s x-\mu_a)$ multiplied by the signal-to-noise ratio, $\sigma_a^2/(\sigma_a^2+\sigma_e^2)$. If individual voting decisions reflect both this perceived difference in the ability of the two candidates to generate a strong economy, and some unrelated voter-specific idiosyncratic assessment, $\varepsilon \sim N(0, \sigma_\varepsilon^2)$, then:

$$\textit{Incumbent vote share} = V\left(\frac{\sigma_a^2}{\sigma_a^2+\sigma_2^2}\left(s-n-\beta_s x-\mu_a\right)+\varepsilon \right); \quad V'>0 \qquad (2)$$

Rogoff and Sibert 1988, Persson and Tabellini 1990, Rogoff 1990, Alesina and Rosenthal 1995. A largely separate literature explores partisan models that yield different electoral business cycles under different parties.

5. This simple formalization focuses only on the signal extraction problem; for a fuller principal-agent treatment that yields identical implications, see Patty and Weber (2005).

Yielding three testable implications : $dV/ds > 0; \quad dV/dn\big|_{s-n} = 0; \quad dV/dx\big|_{s-n-\beta x} = 0$.

Thus, voters seeking to maximize the ability of their officials will link their vote to state economic conditions ($dV/ds>0$). In this signal-extraction model voters reward good economic conditions because they probably reflect competence; in alternative formalizations, this rule also has the benefit of inducing greater effort from election-motivated politicians. This first implication of the theory is typically called *retrospective economic voting*. There is substantial evidence of a robust correlation between strong economic outcomes and the re-election of incumbent politicians in national elections.[6] However, influential papers in both political science (Chubb, 1988) and economics (Peltzman, 1987) have suggested this relationship does not hold at the state level. Section 3 reviews these results, and concludes that they reflect inadequate data, and that retrospective economic voting is an important factor in gubernatorial elections.

Sections 4 and 5 then turn to evaluating the extent to which voters reward luck (exogenous factors) as compared with competence. Section 4 implements tests of *relative performance evaluation*, testing the second implication that $dV/dn|_{s-n}=0$. Finally, the model implies that <u>*events unrelated to a governor's competence should have no effect on the voting decisions of rational agents*</u> ($dV/dx|_{s-n-\beta x}=0$).[7] Thus Section 5 turns to more sophisticated instrumental variables tests of performance filtering. Section 6 discusses remaining questions regarding endogeneity, and section 7 concludes.

Three related literatures are worth mentioning. First, Alesina and Rosenthal (1995) reject the restrictions imposed by voter rationality in a four-equation structural model relating congressional and presidential elections, competence, and output growth. Their test is similar to mine in that they ask whether voters distinguish between variation in aggregate output that reflects competence and variation attributed to luck (as identified by their structural model). But their rejection of a rational voter model could reflect either inefficient inference by voters, or a rejection of an excessively restrictive and highly parameterized structural model. By contrast, the present paper assesses the information processing abilities of voters in a framework that minimizes assumptions about the structure of the economy.

My assessment shadows a related literature on CEO pay which tests whether corporate boards filter out the effects of aggregate market conditions in determining CEO compensation. It typically finds little evidence of performance filtering. In a recent example, Bertrand and Mullainathan (2001, p.1) find that CEO pay "responds about as much to a lucky dollar as to a general [earned] dollar." In a similar vein, di Tella and Fisman (2001) find evidence that the pay of state governors (rather than their electoral fortunes) is linked to both policy variables and exogenous factors, leading them to argue for evidence of

---

6. Fair (1978) documents the correlation between strong aggregate growth and re-election of the incumbent party in presidential races, Kramer (1971) analyzes congressional races, and Lewis-Beck (1988) documents similar findings for other OECD countries.

both pay-for-performance and rent extraction.

A third literature draws on social psychology experiments that find evidence of a "*fundamental attribution error*." These experiments typically reveal that subjects aiming to assess competence systematically fail to take sufficient account of background or environmental factors. Patty and Weber (2005) provide laboratory experiments in which they simulate elections under imperfect information, yielding results complementary to the field evidence assembled below.

III. Economic Voting in Gubernatorial Contests

I start by examining whether voters in state elections systematically reward superior economic outcomes. Table 1 analyzes the relationship between state economic outcomes and the electoral fate of the governing party. The basic political data incorporate the results of all post-war gubernatorial elections in all fifty U.S. states from 1947 to 1997.[8] As is standard in the literature, this paper looks only at those elections that approximate a two-party contest.[9] Thus, the largest sample analyzed includes 636 elections. In keeping with the two-party metaphor, only votes for the top two candidates are analyzed.[10] (Appendix A provides basic descriptive statistics.) Table 1 regresses the change in the incumbent party's vote share between consecutive elections on various indicators of economic activity; each column represents a different regression that analyzes a different indicator. That is, Table 1 examines regressions of the following form (*s* denotes state and *t* denotes year):

$$\Delta Incumbent\ Party's\ Vote\ Share_{s,t} = c + \gamma \Delta State\ economic\ performance_{s,t} + \varepsilon_{s,t} \qquad (3)$$

Throughout this paper, independent variables are coded as changes over the two years leading up to the election.[11] The first column of Table 1 analyzes real personal income per capita data, which are included mainly to provide a link to the existing literature.[12] Peltzman (1987) found that these state income data were a poor predictor of gubernatorial election results, which he interpreted as evidence that voters believe that state governors have little effect on the local economy. Chubb (1988, p.150) analyzed similar data, concluding that this weak relationship suggests that "gubernatorial elections remain contests of party and personality and not of performance." (Furthermore, while Besley and Case (1995a) are

---

7. Unlike many other applications of the principal-agent framework, risk aversion is not necessary for performance filtering to be optimal.
8. Data up until 1990 were extracted from ICPSR (1995). From 1990 to 1997, data were hand-coded from Congressional Quarterly (1998).
9. Thus, special elections were eliminated from the sample, as were both contests in which a third party gained more than 20% of the vote, and those that were largely uncontested (only one candidate, or a winning vote share in excess of 80%). Further, because first differences will be analyzed, the elections following any of these eliminated cases were also removed from the sample.
10. Where it was obvious, candidates were labeled as either Democrats or Republicans. The constituency represented by independent and other party candidates was inferred from the affiliation of the candidate they were opposing.
11. This is standard in the existing literature, presumably reflecting the presence of two-year election cycles in several states.

concerned with analyzing the effects of tax policy on election outcomes, they also include state income growth as a control, finding it to be unrelated to the re-election prospects of incumbent governors.) Like all these authors, I find little evidence of a link between growth in state income and electoral outcomes.

There are two key reasons to be cautious about such findings. First, these state income data are subject to significant measurement error (especially in the earlier years), which biases estimated coefficients towards zero. Second, because the BEA publishes income per capita separately for each state but only compiles national deflators, it is impossible to disentangle whether measured income growth reflects a rise in real incomes (which voters would reward), or a rise in local prices (which they would punish).[13] Thus columns 2 and 3 examine alternative, and arguably preferable, economic indices.

The unemployment rate provides a natural measure of shifts in the real economy. Unfortunately, state unemployment data are only available since the mid-1960s, and even then, are unreliable for small states. I therefore constructed a proxy indicator for the unemployment rate, the *employment gap*. This is constructed from the more comprehensive non-farm payrolls data, and measures the deviation of log employment from its trend level. (To see the analogy, note that the unemployment rate is approximately equal to the deviation of log employment from another slow-moving trend, the log labor force.) The trend estimate is derived using a Hodrick-Prescott filter. (See Appendix B for further details about the construction of the employment gap and evidence showing close co-movement with the unemployment rate.)

Column two shows a strong and statistically significant relationship between movements in the employment gap and the re-election prospects of the incumbent party. To simplify the presentation, a positive number denotes an improving economy. Regressions (not shown) based on the deviation of employment growth from state averages yield largely similar results.

Finally, housing prices are the regressor in column three. House prices are especially relevant because they represent a forward-looking market valuation of the full range of local amenities, including public goods, crime, job opportunities, and so on. The house price indices published by the Office of Federal Housing Enterprise Oversight are particularly informative because they are based on repeated sales of the same houses over time, thereby minimizing distortions induced by compositional changes.[14] Although these data are only available since 1980, they still reveal a strong and statistically significant relationship between state performance and the re-election prospects of the incumbent party.

The magnitudes in Table 1 are directly interpretable. For instance, a 1 percentage point fall in

---

12. Although Gross state product may be a preferable aggregate, it is only available since 1977.
13. It is striking how much less noisy recent state income data are compared with data for the earlier period. Running analogous regressions on data since 1980 reveals that a 1 percentage point rise in income growth over the two years leading up to an election yields a 0.44 percentage point swing to the incumbent (se=0.22). Stronger results in this sample probably reflect higher quality data, and lower and less variable inflation, both of which lead measured income growth to be an increasingly better proxy for real income growth.
14. These data are available at: *www.ofheo.gov*, and were deflated using the chain GDP deflator.

the employment gap is associated with a 0.27 percentage point rise in the governing party's share of the vote. While this is both larger and more significant than analogous estimates in the existing literature on state elections, it is somewhat smaller than its counterpart in the literature on presidential elections.[15]

In sum, Table 1 provides reasonably strong evidence that the state of the economy affects the re-election prospects of the incumbent party. Only one indicator—growth in state income—fails to reveal this correlation; this is also the indicator that has received the most attention in the existing literature.[16]

IV. Relative Performance Evaluation

The simplest test of the principal-agent predictions follows from tests of "relative performance evaluation" common in the literature on the determinants of CEO pay (Gibbons and Murphy, 1990). In both cases the basic idea is that principals should not reward agents for national trends but only for how their enterprise performs relative to this baseline. Adapting the tests for CEOs to the present context involves regressing the change in the vote share won by the incumbent party against both changes in the national economy and changes in the state economy relative to the change in the national economy. The identifying assumption is that the national average provides no information about the governor's competence, whereas the state's deviation from the national outcome reflects gubernatorial competence and effort. That is:

$$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda\ \Delta National\ performance_t + \varepsilon_{s,t} \quad (4)$$

The estimated coefficient α is characterized as the reward for gubernatorial competence, and λ measures the reward for "luck." Rational voters should be unmoved by false signals, suggesting that efficient inference implies $\lambda=0$. Conversely, the psychology literature on attribution errors suggests that environmental or background factors are under-weighed in assessments of competence, leading to $\lambda>0$. At the other extreme, Gibbons and Murphy test (and reject) the null that principals do not distinguish at all between signal and noise when assessing CEOs ($\alpha=\lambda$).[17]

The effect of national and state-specific economic performance on the fate of incumbents is

---

15. For example, Alesina and Rosenthal (1995) find that an additional percentage point of economic growth adds 0.8 percentage points to the incumbent president's vote share. To compare this with the employment gap estimates in Table 1, note that each percentage point of annual income growth in excess of 2¼ percent causes the employment gap to fall by about 0.7 percentage points. Thus the estimate in Table 1 that an 0.27 percentage point vote swing follows a 1 percentage point decline in the employment gap suggests that a 1 percentage point increase in income growth yields a vote swing of 0.2 percentage points. Note that this does not imply that economic performance is less salient in state elections; it may be that a greater share of the variation in state output reflects idiosyncratic factors that voters discount when assessing the performance of their governor. Table 2 shows further evidence in support of this interpretation.

16. However, Lowry, Alt, and Ferree (1998) include a range of fiscal and political variables in their analysis, finding that state income growth relative to the national average is significantly correlated with the incumbent's vote share in gubernatorial, but not state legislature, elections.

17. The CEO pay literature contains papers that reject both the rationality null (Bertrand and Mullainathan, 2001), and this latter irrationality null (Gibbons and Murphy, 1990).

shown in Table 2.[18] Each column shows results for a different indicator of the state's economic health. The dependent variable in Panel A is the change in the share of the two-party vote gained by the incumbent party. The unit of analysis in Panel B is the individual governor rather than the party, and this analysis focuses on probit regressions explaining whether the incumbent governor was re-elected.[19] These data, covering a shorter sample period, were kindly provided by Anne Case.

Table 2 shows strong evidence of performance filtering by voters. The coefficients reported in the first row of each panel show that a state's relative economic health is an important factor in gubernatorial elections; local economic factors were found to be highly significant (and correctly signed) in five of the six regressions. (Again, data on personal income per capita yield inconclusive results.[20]) The second row asks whether voters are fooled by developments at the national level. In no case can the null that voters completely filter out these false signals be rejected.[21] The third row tests the opposite null of no performance filtering. I reject this null for the employment gap. The tests using housing price and state income data yield little statistical power.

Thus Table 2 suggests both that the state economy has important effects on electoral outcomes, and that performance is evaluated relative to a national yardstick. This is consistent with Besley and Case's evidence of yardstick competition in tax setting. Besley and Case (1995a, p.35) also analyze unemployment and income growth, and conclude that "while it is possible for citizens to give the governor a relative grade based on these criteria, it does not appear that voters are judging governors in this way." This conclusion was based on only 85 elections, which yielded fairly imprecise estimates. Thus, a preferable interpretation is that they lacked the statistical power to answer this question.[22]

Table 3 explores the possibility that confounding factors are driving the results in Table 2. The

---

18. Separate estimates of the national levels of income per capita and the employment gap were computed for each state, removing the influence of that state on the national aggregate.

19. In examining the re-election of incumbent governors, probit regressions are necessary to minimize problems arising from the unobservable vote share of incumbents who choose not to stand for re-election. Specifically, incumbents who did not run, despite being eligible, are coded as losing their re-election bids; elections in which the incumbent was either barred by virtue of term limits or chose to run for higher office were dropped from the sample.

20. The significant effect of income per capita in Panel B but not Panel A appears not to reflect different sample periods. Regressing the change in the incumbent party's vote share on state and national performance over the 356 elections in Panel B yields a coefficient on the national cycle of 0.16 (se=0.24), and a coefficient of 0.60 (se=0.17) on the state-specific component.

21. A natural question when assessing the effects of the national cycle is whether to estimate standard errors as though there are 636 independent experiments in the data, or whether there is effectively only one independent experiment each year. Fortunately theory resolves the issue: under the null of voter rationality there will be no cross-state correlation in anti-incumbent sentiment, and hence each observation is an independent experiment. (Empirically, the results presented are largely unaffected by this choice.)

22. Lowry, Alt and Ferree (1998, p.763) appear to have stumbled upon evidence of relative performance evaluation in their work assessing the incumbent's vote share. When discussing specifications that they discarded, they state "when we include separate variables for national and state income growth… coefficients on the national income growth variable are insignificant *or perversely negative*" (emphasis added). Of course, the finding that voters subtract the national growth rate from state growth when assessing the incumbent is not perverse, but rather, evidence of performance filtering.

first row of Table 3 simply repeats the central estimates from Table 2. Confounding variation might arise if voters use gubernatorial elections as an opportunity to "send a message to Washington." That is, if voters treat gubernatorial elections as a referendum on the president's performance, and if the political affiliation of the typical governor is correlated with that of the president, then this signaling may create a link between the performance of the national economy, and pro- or anti-incumbent swings in gubernatorial elections unrelated to attribution errors. The regression in the second row addresses this by controlling for a variable, *President's Party*$_{s,t}$, that takes the value +1 if the incumbent governor is of the president's party and –1 otherwise, and an interaction of this variable with the performance of the national economy. (This symmetric treatment arises because one party's gain is the other party's loss.) The α and λ coefficients retain their interpretation as the returns to competence and luck, respectively. My central estimates appear roughly unchanged by this modification. The coefficients on control variables are not reported in Table 3, but across all three specifications the president's party fares about 4 percentage points worse than the opposition party. This magnitude is similar to the mid-term slump suffered by the president's party in congressional elections. The interaction term captures the extent to which the incumbent governor is helped or harmed by voters expressing (dis)satisfaction with the president's handling of the economy, and is statistically significant only in the real income data.

A less parameterized set of controls is included in the third row of Table 3, which interacts the *President's Party* variable with a full set of year fixed effects. These interaction terms control for all national swings in partisan sentiment, not simply those reflecting national economic performance, and again α and λ are directly comparable with the estimates in Table 2. The interaction terms are strongly significant, suggesting that the national partisan mood is at least partly expressed in state elections. (Alternatively, partisan preferences may vary with economic conditions.) The fourth row also adds state fixed effects. Political preferences may also vary over the business cycle. While the interaction of *President's Party* with year fixed effects partials out much of this variation, the fifth row also controls for the interaction of a *Party* variable with state conditions. This interaction is both small and statistically insignificant. The sixth row also adds direct year fixed effects. These year fixed effects render the national cycle unidentified, and hence I report only the coefficient on relative state performance.[23] Year fixed effects partial out swings in pro- or anti-incumbency sentiment, while the interacted year fixed effects control for swings in partisan sentiment. Both sets of controls are extremely statistically significant, yet the estimated effects of the relative performance measure remain approximately unchanged. A similar investigation of the probit specification yields analogous evidence that the results from the simple specification in Table 2 are quite robust.

---

23. Because the national performance variable for each state is, in fact, the performance of the United States *excluding that state*, strictly speaking this variable remains identified. However, it is only identified to the extent

In sum, the main findings of Table 2—that voters perform relative performance evaluation—appear to be robust to the inclusion of a fairly wide set of further control variables.

A further set of concerns arises from the possibility that basic parameters of the voting decision vary with economic conditions. Two examples seem salient. First, it may be that voters are fooled by a national downturn into inferring that the incumbent is of low quality, but feel that a recession is an inappropriate time to experiment with a new governor. This fear may be motivated by the possibility that either the advantages of electing an experienced governor, or voters' risk aversion, may vary with the state of the economy. Second, information revelation may vary with the cycle. Consider a simple asymmetry: it may be that the potential for a governor to distinguish herself as either competent or incompetent only occurs during a boom, and that unless convinced of gubernatorial incompetence, most voters choose to re-elect the incumbent. This asymmetry would lead to anti-incumbent swings only during a national boom, confounding the pro-incumbent swing caused by voters making attribution errors. (To adopt the earlier formalization, if $\varepsilon=f(s)$ or $\sigma_e^{-2}=f(s)$ or then $dV/dn|_{s-n}$ will be correlated with s.)

Central to both of these examples is the possibility that my results are driven by an asymmetry. Table 4 shows that the estimated returns to competence and luck appear robust even after dropping data from either periods of robust growth or recession.

In sum, the findings in Table 1 suggest that economic voting is relevant in state elections. Table 2 reinforces this finding, showing that the performance of the state economy, relative to the national economy, appears to be the relevant variable for voters. Consistent with efficient inference, voters appear not to be fooled by improvements in the state economy that are due to the national cycle. Table 3 shows that this result appears not to be driven by confounding variation. Table 4 demonstrates that these results are not identified by an asymmetry. The next section provides more sophisticated instrumental variables tests of whether voters filter a range of other exogenous shocks from their voting decisions. The ensuing section will return to some of the subtler questions regarding the identification and interpretation of these results.

V. Instrumental Variables Tests of Efficient Filtering

The basic approach pursued in this section is to find sources of variation in economic performance that are unrelated to a politician's actions, and then test whether these false signals fool voters into inferring competence. Whereas the previous section tested whether variation in the state economy due to the national cycle affected voting behavior, this section extends this analysis to a range of other shocks. Specifically, I aim to isolate specific driving forces of state fluctuations that are

---

that the state's performance differs from the national norm, and hence including it would only confound interpretation of the coefficient on relative state performance.

plausibly unrelated to a governor's actions. Rather than simply regress the change in the incumbent's vote share on these exogenous variables, this section applies instrumental variables (IV) methods that facilitate direct comparison of the returns to luck and the returns to competence.

The IV approach provides a direct generalization of standard tests of relative performance evaluation to examine whether a range of factors exogenous to the governor's decisions affect voting behavior.[24] A first stage regression partitions variation in state economic performance into that which clearly does not reflect gubernatorial competence, and a residual component that may, in part, reflect competence. The second stage regression assesses the link between re-election prospects and these indicators of gubernatorial competence and luck.

To understand the link between this and the previous section, it is useful to re-conceptualize the simple tests of relative performance evaluation in Tables 2 to 4 in an IV framework. Tests of relative performance evaluation implicitly treat *National performance* as a plausibly exogenous instrument for *State performance*. Specifically, it is an instrument for the "luck" that has aided the governor's management of the state economy. The residual variation in state outcomes, *State performance minus National performance*, is an imperfect proxy for gubernatorial competence. Implicitly, this first stage regression was imposed (with a coefficient of one) rather than estimated, and these two instruments were then used as regressors to assess the role of luck and competence in voting decisions.

This section simply extends this approach to also test for the effects of other sources of luck:

$$\begin{aligned}\Delta \text{Incumbent vote}_{s,t} = {} & \lambda\, \Delta \text{National employment gap}_t + \delta(\beta_s * \text{Observable Shock}_t) \\ & + \alpha(\Delta \text{State employment gap}_{s,t} - \Delta \text{National employment gap}_t - \beta_s * \text{Observable Shock}_{s,t}) + \varepsilon_{s,t}\end{aligned} \quad (5)$$

As before, the signal of the state of the economy is the two-year change in the national employment gap. Thus, λ is an estimate of the returns to luck. Specifically, λ reflects the impact of the national economy on the incumbent governor's vote share, and hence indicates voters' capacity to perform simple "rule of thumb" performance filtering, comparing the performance of the state economy with the national average. A finding that $\lambda>0$ suggests that voters systematically underweigh national trends when assessing their governor's performance. $\lambda=0$ suggests that at least "rule-of-thumb" performance filtering occurs.

Of the remaining variation in state outcomes after subtracting off the national change in the employment gap, at least some is attributable to an observable shock, *x*, which is both known to be outside the governor's control, and may have differential effects across states. In order to isolate this variation (and hence identify the β_s coefficients in equation 5), I run a first stage regression:

$$\Delta \text{State employment gap}_{s,t} - \Delta \text{National employment gap}_t = \sum_{s}^{states} \beta_s * \text{Observable Shock}_{s,t} + \varepsilon_{s,t} \quad (6)$$

24. Bertrand and Mullainathan (2001) provide a related approach.

The predicted values from this equation are attributed to $x_{s,t}$, and hence reflect factors outside the governor's control. Thus δ (in equation 5) reflects the capacity of voters to make more fine-grained assessments, filtering variation due to the observable shock from their assessment of competence. A finding that $\delta>0$ is evidence of imperfect performance filtering.

The interpretation of α, the coefficient on the residual from the first stage equation, is subtler, and will help in assessing the relative magnitudes of the returns to luck and the returns to competence. Recalling the driving forces of state economic conditions, $s=a+n+\beta_s x+e$, the remaining variation in state outcomes after taking account of the national economy and these observable shocks reflects both gubernatorial abililty and effort, *a*, and other unobservable (to the econometrician) shocks, *e*. Thus, the estimated α coefficient in the voting equation is identified off this amalgam of ability and unobserved shocks and will be a weighted average of the returns to competence and the returns to these unobserved shocks. The extent to which voters observe *e* is unknown, complicating any intuition regarding the extent to which α accurately reflects the returns to competence.

Three central cases will clarify the relationship between α and the true causal effects of competence on re-election prospects. First, if these unobserved shocks, *e*, are not observed by voters, then they cannot distinguish between the noise and signal elements of $s-n-\beta_s x$, and hence they reward both *a* and *e* at the same rate. Thus, in this case, α accurately represents the returns to the governor of making competent decisions. Second, if "unobserved luck," is observed by rational voters, they will fully discount it, and hence α reflects an attenuated estimate of the returns to competence. The third case is a generalization of this result. If *e* is observed by voters, and yet they continue to make attribution errors, then α will reflect a weighted average of the returns to luck and the returns to ability. Weaker first stage regressions exacerbate this bias by increasing the relative contribution of unexplained economic shocks versus ability to the residual variation in economic outcomes used to estimate α.[25]

In sum, an appropriately exogenous instrumental variable isolates a source of variation in economic performance that reflects only "luck." Hence the estimated λ and δ coefficients accurately reflect the returns to luck (albeit different sources of luck). However, the difficulty in fully separating variation due to ability from that due to luck means that estimated returns to ability are probably biased towards the returns to luck. Thus, assuming that the true returns to luck do not exceed the true returns to ability, α is a downward biased estimate of the causal effect of ability on re-election prospects.

The regressions in each column of Table 5 and 6 reflect different instrument sets. In Table 5, I test whether voters filter out state-level economic fluctuations that are attributable to oil shocks. Oil

---

25. $\hat{\alpha}$ is a biased estimate of the returns to competence for reasons analogous to the weak instruments problem (Bound, Jaeger, and Baker, 1997). *State performance-National performance-*β_s*Observable shocks* is only a weak instrument for ability (*a*) and, the weaker the first stage regression, the more likely it is that $\hat{\alpha}$ will be biased towards the coefficient from an OLS regression of *ΔIncumbent's vote* on *State performance*.

prices are an ideal instrument both because they are unlikely to be affected by gubernatorial competence and because they cause a boom in oil producing states, and a downturn in rust-belt states. In the first column, the (lagged) two-year ended change in the log of the real oil price is interacted with state fixed effects.[26] That is, fifty separate instruments are employed, and oil prices are estimated to have a different effect on each state:

$$\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t = \sum_{s \in states} \left(\mu_s + \beta_s \Delta Log\ Real\ Oil\ Price_{t-1} \right) \quad (7)$$

Panel A shows the first stage regression. This first stage regression yields plausible estimates with remarkable explanatory power, suggesting that states like Alaska, Wyoming, and Texas are the major beneficiaries of price rises, and Michigan and Indiana the major casualties.

Panels B and C analyze the effects of these shocks on the vote share of the incumbent party and the probability of re-election of the incumbent governor, respectively. The first row in each panel shows that voters are not fooled by the national cycle into favoring incumbents. The second row asks whether voters are also cognizant of the extent to which their state's relative economic performance has been helped or harmed by oil price shocks. The large and statistically significant effect on election outcomes of the economic fluctuations caused by oil prices suggests that voters in oil-producing states are systematically fooled into re-electing their governors when the oil price has shot up, while their counterparts in oil-dependent states vote their incumbents out. The third row confirms that the residual variation, labeled "competence," continues to have explanatory power. Comparing the coefficients in rows two and three is not particularly informative, given that measurement error is likely attenuating the returns to the unexplained component of measured state performance.

The second column imposes more structure on the instrument for luck. Rather than simply allow the effects of oil prices to vary independently across states, this column imposes the restriction that these different sensitivities reflect different industry mixes. Thus, the fifty state-specific instruments are replaced with nine instruments, each reflecting the interaction of the national cycle with the share of each industry's state earnings. Thus the first stage equation is:

$$\begin{aligned} &\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t \\ &= \sum_{i}^{9\ industries} \theta_i \left(Industry\ share_{i,s}^{1940s} * \Delta Log\ Real\ Oil\ Price_{t-1} \right) + \varepsilon_{s,t} \end{aligned} \quad (8)$$

State industry shares are measured using BEA data on Earnings by Industry.[27] These industry shares are measured using historical data (the average share of state earnings in the 1940s), so that it is clear that these do not represent the legacy of any incumbent governor in this sample. The first stage

26. Oil price data are annual averages of the Crude Petroleum PPI, deflated using the chain GDP deflator. Alternative oil price data yield similar results.
27. State Earnings data were downloaded from *www.bea.doc.gov*. This national accounting aggregate includes wage and salary disbursements, other labor income, and proprietors' income.

regression reveals that states with historically large manufacturing and service sectors are particularly vulnerable to oil price rises, while states with large mining, finance, construction, and farm sectors have a more muted sensitivity; transport and wholesale trade are approximately neutral.[28] These first stage results yield predicted state sensitivities that are both plausible and comparable with those in column one (the range of predicted state sensitivities is shown in Panel A).

Turning to the returns to luck and competence in Panels B and C, column two largely reinforces the results from column one. Voters appear to perform simple rule-of-thumb adjustments, but by failing to take account of the effects of oil prices on state economic outcomes resulting from the state's historical industry mix, they make systematic attribution errors.

Note that the regressions in Panels B and C report standard (heteroscedasticity consistent) OLS standard errors. These standard errors are biased downward because they do not account of the extra uncertainty associated with the analysis of generated regressors. Rather than attempting to analytically derive the relevant standard errors for the case in which both a predicted value and its residual are included as regressors, I also report bootstrapped standard errors. These yield standard error estimates [shown in square brackets] are remarkably close to the White standard errors.

Table 6 examines an alternative source of exogenous shocks to state economies. Specifically, I examine economic fluctuations that reflect the differential sensitivity of each state to aggregate shocks. Column one employs fifty instruments—one for each state—so that the aggregate shocks reflected in the national employment gap are estimated to have a different effect on each state:

$$\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t = \sum_{s}^{50\ states} (\mu_s + \beta_s \Delta National\ employment\ gap_t) + \varepsilon_{s,t} \quad (9)$$

Shadowing results in Blanchard and Katz (1992) and Davis, Loungani, and Mahidhara (1997), Panel A shows statistically and economically significant differences in cyclical sensitivities, ranging from Michigan, which is extremely pro-cyclical, to Alaska, which is in fact counter-cyclical.

Consistent with the message in Tables 2 to 4, the first rows in Panels B and C of Table 6 suggest that voters at least perform rule-of-thumb signal extraction, comparing their state outcomes with the national norm, and there is no (statistically discernible) systematic tendency for national booms to lead to the re-election of incumbents. However, the second row strongly rejects the notion that voters are using a fully efficient signal extraction rule. That is, in pro-cyclical states a national boom leads to an even stronger local boom, and voters are fooled by the relative strength of their state compared with the national average into voting for the incumbent, presumably based upon a mistaken inference of

28. The logic behind the result for agriculture is slightly altered by the fact that the employment data are *non-farm payrolls*. To be clear, the regression finds that the non-farm sector is less sensitive to oil prices in farming-intensive states. Similar interpretation issues arise in Table 6.

gubernatorial competence. On the flipside, these same voters are too quick to oust the incumbent in a downturn. In those states with a more muted business cycle, a symmetric, but opposite pattern occurs, with excessive turnover in national booms and a (relative) status quo bias in downturns. The third row continues to confirm that the residual variation in economic performance – a rough proxy for competence – has explanatory power.

As with Table 5, the second column of Table 6 imposes more structure on the instrument for luck, constraining each state's estimated cyclical sensitivity to reflect its historical industry mix. Thus, the fifty individual instruments are replaced with nine instruments, each reflecting the interaction of the national cycle with the share of each industry in State Earnings:

$$\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t = \sum_{i}^{9\ industries} \theta_i (\ Industry\ share_{i,s}^{1940s} * \Delta National\ employment\ gap_t\ ) + \varepsilon_{s,t} \quad (10)$$

The first stage regression reveals that states with historically large manufacturing, wholesale trade, and service sectors tend to be more sensitive to the national cycle, while agriculture, construction, finance, and government sectors yield a more muted sensitivity; transport and mining are approximately neutral.

The results in column two reinforce those in column one. Again, there is no systematic national tendency to vote for incumbents in a national upswing (or against them in a downturn). However, an upswing leads those in pro-cyclical states to vote for the incumbent, while their peers in counter-cyclical states vote against the incumbent. A downturn leads to the opposite conclusion.

A full battery of checks shadowing those in Tables 3 and 4 suggest that the basic results in Tables 5 and 6 are robust. Appendix C also explores questions regarding the potential endogeneity of oil prices, and yields fairly similar results. I also examined a variety of other sources of luck, including Canadian/US dollar exchange rate movements interacted with proximity to the border, a Bartik (1991) instrument based on the interaction of industry shares and industry shocks, and weather-based instruments interacted with agricultural intensity. In each case, the first stage regressions yielded statistically significant relationships between the instrument and economic outcomes, but the instrument lacked sufficient power to offer much insight into the relationship with voting patterns.

In sum, the evidence accumulated above suggests that voters perform simple rule-of-thumb adjustments when evaluating gubernatorial competence. However, they still appear to make systematic attribution errors, and voters are systematically fooled into re-electing lucky but less competent governors.

VI. Endogeneity Issues

The instrumental variables analyzed in this paper are valid if they explain state economic

performance, but do not have a direct effect on the voting decision of rational voters. In the related literature testing for performance filtering in CEO compensation, Himmelberg and Hubbard (2000) argue that aggregate instruments reflecting "luck" are also shocks to the demand for inelastically supplied CEO talent. Thus, even in the absence of attribution error, positive aggregate shocks will lead to a rise in CEO pay. However, while wages are set taking account of the outside option of the principal, this argument lacks a clear parallel in the voting context. Thus, elections arguably provide a cleaner test of attribution error than the CEO labor market.[29]

However, there are other possible reasons to doubt the IV exclusion restriction. Politicians have thus far been characterized as motivated solely by the desire for votes. Real governors also have other conflicting objectives, such as a desire to implement an ideological platform, or simply to extract rents. The potential extra votes gained from a stellar economic performance may be spent in pursuing these alternative agendas.[30] Politicians who care only about partisan or rent-stealing objectives will spend all of their excess political capital on these objectives, and hence if all of the extra votes delivered by a booming economy are "spent" on these alternative, vote-losing activities, no correlation between economic performance and the vote for the incumbent would be observed.

Most likely is an intermediate case, in which politicians "spend" some of their extra votes pursuing unpopular projects. This endogenous response will attenuate my estimates of the returns to both luck and competence. However, the estimated returns to luck will be wrongly signed only in the perverse case when for every hundred extra votes delivered to the incumbent governor by a booming economy, this renewed job security induces them to undertake so many unpopular projects that they lose over a hundred votes. Thus, as long as the marginal propensity to consume political capital is less than one, the estimated returns to luck are properly signed, albeit biased towards zero. Despite this bias, Tables 5 and 6 report significant and large returns to luck.

This paper also provides evidence on the extent to which the returns to luck differ from the returns to competence. If politicians have an equal propensity to "spend" the political capital gained through either competence or luck, then the estimated returns to both luck and competence will be biased towards zero. Indeed, this bias will affect them both in equal proportion, and hence their ratio will be unaffected, although their difference will also be biased towards zero.[31]

While controlling for observable indicators of a governor's policy choices may help offset the

---

29. Further, in the CEO context, strategic incentives to commit to softer product competition (Aggarwal and Samwick, 1999), or the possibility of sabotage, collusion, or production externalities, may yield optimal contracts at odds with the standard of relative performance evaluation, undermining the inference of attribution error (Gibbons and Murphy, 1990). These arguments also seem less applicable in the political domain.

30. In related evidence, Besley and Case (1995b) find that governors barred by term limits from running for re-election make systematically different decisions from those who have an eye on their re-election prospects.

extent to which this endogenous feedback biases my results, variables describing state government revenue, levels of spending, and types of spending were found to be jointly and individually insignificant, and their inclusion did not perceptibly change the estimated returns to competence and luck.[32] Even so, if this feedback is expressed at other (non-fiscal) margins, it is likely to cause Type II but not Type I errors.

VII. Discussion

The principal-agent model suggests that voters will use information implicit in observable indicators—such as the state of the economy—to assess the competence of incumbent governors. This is strongly confirmed in the data. The second implication is that voters will discount variation in the state economy that is due to factors that bear no relation to gubernatorial competence. The data suggest that voters evaluate their state's economy relative to the national economy. However, these results suggest only rule-of-thumb performance filtering. Instrumental variables regressions indicate that voters in pro-cyclical states are systematically fooled into re-electing incumbents during national booms, only to dump them during national recessions. Similarly, voters in oil-producing states tend to re-elect incumbent governors during oil price rises, while voters in oil-dependent states oust their incumbents.

While I interpret these results as rejecting efficient inference, evidence that voters do not act in strict accordance with the predictions of the principal-agent model does not *necessarily* infer irrationality. There may exist rational deviations from the principal-agent model, such as the preference of voters for stable leadership during periods of war. This is a specific example of a more general issue: that almost any set of choices can be rationalized by appealing to a particular formulation of preferences. Ultimately the plausibility of such a counter-argument rests upon the plausibility of the required preferences. The simplest explanation of my results is that voters are quasi-rational—while they assess state performance relative to a national yardstick, more fine-grained assessments remain elusive, and they make systematic attribution errors.

One potentially useful benchmark is available from the literature on the determinants of CEO pay. Gibbons and Murphy (1990) analyze the sensitivity of CEO pay to a firm's rate of return on common stock. They find some evidence of performance filtering, and CEO pay is roughly unaffected by those aggregate shocks evident in a broad index of stock returns. However, shocks that are idiosyncratic to a firm's industry are only weakly filtered, if at all. Bertrand and Mullainathan (2001) provide recent confirmatory evidence that CEOs are as likely to be rewarded for industry-specific shocks as for effort. Further, they also analyze a sample of oil companies, and despite the obvious importance

---

31. Again, one can construct a case in which this bias might lead one to wrongly reject the null that $\alpha=\lambda$. This would only occur if politicians spend a greater proportion of the extra votes that they win through "luck," than those they win through competent governance.

of oil prices to firm performance, they find little evidence of performance filtering. Measured against this yardstick, the rather similar performance of voters in assessing their agents is striking, especially given that voting typically involves small stakes.[33]

This paper has assessed the ability of voters to distinguish signal from noise when assessing the competence of their incumbent governor. The results suggest that voters do a reasonable, albeit imperfect job of disentangling variation in the economy due to competence from variation from other factors. As such, the results in this paper can help inform debates about the relative merits of rational and naïve voter models of political business cycles. Recall that the ability to distinguish signal from noise was central to distinguishing between models of rational and naïve voters. However, the question asked by these models differs subtly from the question examined in this paper. This paper analyzes whether voters respond to variation in state economic outcomes that reflects clearly exogenous and easily observed shocks. By contrast, in political business cycle models, the question is whether voters respond to variation in state economic outcomes that reflects pre-electoral policy manipulation by opportunistic politicians. While a statistically efficient signal-extraction rule dictates that voters should not respond to noise in either context, the costs and benefits of inefficient inference may differ across these domains. Specifically, note that politicians are likely to manipulate policy instruments in as opaque a manner as possible, raising the cost of making accurate inferences about their ability. However, the benefits of filtering out the effects of pre-election manipulation are also higher, as filtering—in addition to lowering the probability of occasionally re-electing an incompetent incumbent—also removes any incentive for incumbents to create wasteful output volatility. More direct evidence on how voting principals assess their elected political agents can help resolve these issues.

---

32. Sam Peltzman generously shared the data on state government spending from his 1992 QJE paper.
33. Note that Bertrand and Mullainathan show that corporate boards with greater interest in good governance yield CEO contracts closer to the principal-agent predictions.

References

Achen, Christopher and Larry Bartels, "Blind Retrospection: Electoral Responses to Droughts, Flu and Shark Attacks", *mimeo*, Princeton University.

Alesina, Alberto and Howard Rosenthal, *Partisan Politics, Divided Government, and the Economy* (Cambridge: Cambridge University Press, 1995).

Alesina, Alberto, Nouriel Roubini, and Gerald D. Cohen, *Political Cycles and the Macroeconomy* (Cambridge: MIT Press, 1997).

Bertrand, Marianne and Sendhil Mullainathan, "Do CEOs Set Their Own Pay? The Ones Without Principals Do," *Quarterly Journal of Economics* 116(4) (2001).

Besley, Timothy and Anne Case, "Incumbent Behavior: Vote-Seeking, Tax-Setting, and Yardstick Competition," *American Economic Review* 85(1) (1995a)

_____ and _____,"Does Electoral Accountability Affect Economic Policy Choices? Evidence from Gubernatorial Term Limits," *Quarterly Journal of Economics* 110(3) (1995b)

Blanchard, Olivier Jean and Lawrence F. Katz, "Regional Evolutions," *Brookings Papers on Economic Activity* 1 (1992)

Bound, John, David A. Jaeger, and Regina Baker, "Problems with Instrumental Variables Estimation when the Correlation Between the Instruments and the Endogenous Explanatory Variable Is Weak," *Journal of the American Statistical Association* 90, no. 430 (1997)

Chubb, John E., "Institutions, The Economy, and the Dynamics of State Elections," *American Political Science Review* 82(1) (1988)

Congressional Quarterly, *Gubernatorial Elections, 1787–1997* (Washington, D.C.: Congressional Quarterly Inc., 1998).

Cukierman, Alex. and Allan Meltzer, "A Positive Theory of Discretionary Policy, the Cost of Democratic Government, and the Benefits of a Constitution," *Economic Inquiry* 24 (1986)

Davis, Steven J., Prakash Loungani, and Ramamohan Mahidhara, "Regional Labor Fluctuations: Oil Shocks, Military Spending, and Other Driving Forces," *International Finance Discussion Papers, Federal Reserve Board* no. 578 (1997)

Di Tella, Rafael and Ray Fisman, "Are Politicians Really Paid Like Bureaucrats?"*, Journal of Law and Economics*, 47(2) (2004)

Fair, Ray C., "The Effect of Economic Effects on Votes for President," *Review of Economics and Statistics* 60(2) (1978)

Gibbons, Robert and Kevin Murphy, "Relative Performance Evaluation for Chief Executive Officers," *Industrial and Labor Relations Review* (1990).

Hamilton, James D., "What is an Oil Shock?," *mimeo University of California, San Diego* (2000).

Himmelberg, Charles and R. Glenn Hubbard, "Incentive Pay and the Market for CEOs: An Analysis of Pay-for-Performance Sensitivity," *mimeo* Columbia University (2000).

Inter-university Consortium for Political and Social Research, *Candidate and Constituency Statistics of Elections in the United States, 1788–1990*, electronic file: www.icpsr.umich.edu (1995).

Kramer, Gerald H., "Short Term Fluctuations in U.S. Voting Behavior, 1896–1964," *American Political Science Review* 65 (1971)

Lewis-Beck, Michael S., *Economics and Elections: The Major Western Democracies* (Ann Arbor: University of Michigan Press, 1988).

Lindbeck, Assar, "Stabilization Policies in Open Economies with Endogenous Politicians," *American Economic Review* 66(2) (1976)

Lowry, Robert C., James E. Alt, and Karen E. Ferree, "Fiscal Policy Outcomes and Electoral Accountability in American States," *American Political Science Review* 92(4) (1998)

Nordhaus, William D., "The Political Business Cycle," *Review of Economic Studies* 42(2) (1975)

Peltzman, Sam, "Economic Conditions and Gubernatorial Elections," *American Economic Review* 77(2) (1987)

_____, "Voters as Fiscal Conservatives," *Quarterly Journal of Economics* 107(2) (1992)

Persson, Torsten and Guido Tabellini, *Macroeconomic Policy, Credibility, and Politics* (Switzerland: Harcourt, 1990).

Rogoff, Ken, "Equilibrium Political Budget Cycles," *American Economic Review* 80(1) (1990)

Patty, John and Roberto Weber, "Letting the Good Times Roll: A Theory of Voter Inference and Experimental Evidence," *mimeo*, Carnegie-Mellon University (2005).

Rogoff, Ken and Anne Sibert, "Elections and Macroeconomic Policy Cycles," *Review of Economic Studies* 55(1) (1988)

Table 1: Effects of State Economic Conditions on the Vote for the Incumbent Party in State Gubernatorial Elections

Dependent Variable: Change in incumbent party's share of two-party vote

	Real Income per capita[a]	Employment gap[b]	Real Housing Prices[c]
Coefficient	.06	.27***	.19**
	(.07)	(.10)	(.08)
Adjusted R^2	-.001	.011	.025
Sample	636 elections 1947-97	636 elections 1947-97	185 elections 1982-97

Each column shows a separate regression of the change in the incumbent party's share against a specific indicator of the state's economic performance over the two years leading up to the election.
(Robust standard errors in parentheses.)
***, **, * denote statistical significance at 1%, 5%, and 10%, respectively.
[a] Two-year ended percentage change in state income per capita. Data from the Commerce Department, deflated using the Chain GDP deflator.
[b] Two-year ended change in the employment gap, measured in percentage points. Constructed from BLS non-farm payrolls data (see Appendix B for details). Results are coded so that a positive number denotes an improving economy (analogous to a declining unemployment rate).
[c] Two-year ended percentage change in real housing prices. Housing prices are measured from a repeat sales index provided by the Office of Federal Housing Enterprise Oversight, deflated using the Chain GDP deflator.

Table 2: Do Voters Filter Out the Performance of the National Economy?

	Independent Variable (2 year-ended changes) (Each column is a separate regression)		
	ΔEmployment gap (%points) [a]	**%ΔHouse Prices**	**%ΔReal income per capita**
Panel A: OLS regression **Dependent Variable: ΔIncumbent *Party* Two Party Preferred Vote Share (OLS)**			
	(1)	**(2)**	**(3)**
α: Effects of competence ($\Delta State_{s,t}$-$\Delta National_t$)	.42***	.22***	.08
	(.13)	(.08)	(.08)
λ: Effects of luck ($\Delta National_t$)	.14	.09	.01
	(.12)	(.14)	(.12)
Test: α=λ	F(1,633) = 2.84*	F(1,182)=0.73	F(1,633)=0.33
Adjusted R^2	.013	.022	-.001
n (elections)	636 (*1947-97*)	185 (*1982-97*)	636 (*1947-97*)
Panel B: Probit regression[b] **Dependent Variable: Indicator =1 if incumbent *Governor* was re-elected; =0 otherwise**[c]			
	(4)	**(5)**	**(6)**
α: Effects of competence ($\Delta State_{s,t}$-$\Delta National_t$)	3.0***	1.2**	1.4**
	(1.0)	(0.5)	(0.6)
λ: Effects of luck ($\Delta National_t$)	1.6*	0.8	1.6
	(0.9)	(0.8)	(1.0)
Test: α=λ	$\chi^2(1)$=1.16*	$\chi^2(1)$=0.18	$\chi^2(1)$=0.02
Pseudo R^2	.026	.059	.016
n (elections)	356 (*1950-88*)	69 (*1982-88*)	356 (*1950-88*)

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively.
(Robust standard errors in parentheses.)

[a] Employment gap is coded so that a positive number denotes an improving economy (analogous to a declining unemployment rate).

[b] Probit coefficients report the marginal change in the probability of re-election for a marginal change in economic outcomes, evaluated at cell means.

[c] Re-elected incumbents are coded to a value of 1. Incumbents who lost a primary race, a general election, or who decided not to run again are coded as 0. Incumbents who ran for higher office, or were barred from re-election by term limits are dropped from the sample.

Table 3: Robustness: Relative Performance Evaluation

Dependent Variable: ΔIncumbent *Party* Two Party Preferred Vote Share (OLS)

	Employment gap[a] *(2-year ended Δ, %pts)*		**House Prices**[a] *(2-year ended %Δ)*		**Real income per capita**[a] *(2-year ended %Δ)*	
	α: ΔState-ΔNational *Competence*	**λ: ΔNational** *Luck*	**α: ΔState-ΔNational** *Competence*	**λ: ΔNational** *Luck*	**α: ΔState-ΔNational** *Competence*	**λ: ΔNational** *Luck*
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \varepsilon_{s,t}$						
Basic Specification (from Table 2)	.42***	.14	.22***	.09	.08	.01
	(.13)	(.12)	(.08)	(.14)	(.08)	(.12)
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \phi President's\ Party_{s,t} + \varphi(President's\ Party_{s,t} * \Delta National\ performance_t) + \varepsilon_{s,t}$						
+ Control for president's ec. performance (*President's Party*[b] * *ΔNational Performance*)	.41***	.13	.20**	.10	.09	.04
	(.13)	(.12)	(.08)	(.14)	(.07)	(.11)
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \phi President's\ Party_{s,t} + \sum_{t \in years} \varphi_t (President's\ Party_{s,t} * Year_t) + \varepsilon_{s,t}$						
+ Control for national partisan swings (*President's Party*[b] * *Year Fixed Effects*)[c]	.42***	.15	.21**	.06	.07	-.03
	(.14)	(.12)	(.08)	(.15)	(.07)	(.11)
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \phi President's\ Party_{s,t} + \sum_{t \in years} \varphi_t (President's\ Party_{s,t} * Year_t) + \sum_{s \in states} \mu_s State_s + \varepsilon_{s,t}$						
+ Control for *state fixed effects*	.41***	.15	.20*	.10	.09	-.04
	(.14)	(.13)	(.12)	(.21)	(.08)	(.12)
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \phi President's\ Party_{s,t} + \sum_{t \in years} \varphi_t (President's\ Party_{s,t} * Year_t) + \sum_{s \in states} \mu_s State_s + \kappa Party_{s,t} + \pi(Party_{s,t} * \Delta State\ performance_{s,t}) + \varepsilon_{s,t}$						
+ Control for economy-contingent preferences (*Party*[b] * *ΔState Performance*)	.40***	.15	.20*	.13	.09	-.04
	(.14)	(.13)	(.11)	(.21)	(.08)	(.12)
$\Delta Incumbent\ vote_{s,t} = \alpha(\Delta State\ performance_{s,t} - \Delta National\ performance_t) + \lambda \Delta National\ performance_t + \phi President's\ Party_{s,t} + \sum_{t \in years} \varphi_t (President's\ Party_{s,t} * Year_t) + \sum_{s \in states} \mu_s State_s + \kappa Party_{s,t} + \pi(Party_{s,t} * \Delta State\ performance_{s,t}) + \sum_{t \in years} \tau_t Year_t + \varepsilon_{s,t}$						
+ Control for national swings in pro- or anti-incumbent sentiment (*Year fixed effects*)	.47***	n.a.	.18	n.a.	.10	n.a.
	(.15)		(.12)		(.09)	
Sample (*n*)	*1947-97 (636 elections)*		*1982-97 (185 elections)*		*1947-97 (636 elections)*	

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively. (Robust standard errors in parentheses.)

[a] Economic variables are measured as a percentage change over the two years leading up to the election. (Employment gap is a simple difference, measured in percentage points.)

[b] *President's Party* takes a value of +1 if the incumbent governor is of the same party as the President, and –1 otherwise. (Symmetric treatment reflects the fact that one party's loss is the other party's gain.) Similar logic underlies the coding of $Party_{s,t}$, which takes a value of +1 for Democrat incumbents, and –1 for Republicans.

[c] Note that the interaction of *President's* $Party_{s,t}$ with year fixed effects could equivalently be expressed as the interaction of $Party_{s,t}$ with year fixed effects.

Table 4: Relative Performance Evaluation – Testing for Asymmetries

Dependent Variable: ΔIncumbent *Party* Two Party Preferred Vote Share (OLS)

$$\Delta Incumbent\ vote_{s,t} = c + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t) + \lambda \Delta National\ employment\ gap_t + \varepsilon_{s,t}$$

	α: ΔState-ΔNational *Competence*	λ: ΔNational *Luck*	Adj. R^2	*n*
Whole Sample	**.42*** (.13)**	**.14 (.12)**	.013	636
Sample restriction	**Sample excluding Downturns**			
ΔState Emp. Gap > -10%	.47*** (.13)	.18 (.12)	.018	632
ΔState Emp. Gap > -8%	.51*** (.14)	.21 (.12)	.019	622
ΔState Emp. Gap > -6%	.54*** (.14)	.23 (.13)	.020	602
ΔState Emp. Gap > -4%	.52*** (.15)	.25 (.15)	.017	544
ΔState Emp. Gap > -2%	.41** (.19)	.20 (.18)	.006	434
ΔState Emp. Gap > 0%	.41** (.22)	.20 (.24)	.004	283
Sample restriction	**Sample excluding Booms**			
ΔState Emp. Gap <+10%	.44*** (.14)	.15 (.13)	.014	631
ΔState Emp. Gap <+8%	.47*** (.15)	.17 (.13)	.014	615
ΔState Emp. Gap <+6%	.40** (.16)	.08 (.14)	.008	587
ΔState Emp. Gap <+4%	.36* (.19)	.14 (.17)	.004	535
ΔState Emp. Gap <+2%	.23 (.22)	.00 (.21)	-.001	463
ΔState Emp. Gap <0%	.27 (.29)	-.15 (.31)	.001	353

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively.
(Robust standard errors in parentheses.)
Employment gap measured as the change over the years leading up to the election (percentage points)

Table 5: IV Tests: State Responses to Oil Prices

Panel A: First Stage Regression (2-year ended changes, % points)		
	Instruments for ΔState Employment Gap – ΔUS Employment Gap	
	ΔLog Real Oil Price Interacted with state dummies $\Delta State\ emp\ gap_{s,t} - \Delta US\ emp\ gap_t = \sum_{s \in states} \mu_s + \beta_s (State_s * \Delta P_{t-1}^{Oil})$	**ΔLog Real Oil Price Interacted with state industry shares (measured in the 1940s)** $\Delta State\ emp\ gap_{s,t} - \Delta US\ emp\ gap_t = \sum_{i \in industries} \theta_i (Ind.\ share_{s,i}^{1940s} * \Delta P_{t-1}^{Oil})$
Estimated coefficients	Average: 0 State Range: -.07 to +0.23	Average: 0 State Range: -.06 to +0.20
Most positive effects	AK, LA, OK, TX, WY	Construction, farm, finance, mining
Most negative effects	IN, MI, NV, TN	Manufacturing, services
Adjusted R^2	.15	.12
Sample (n)	2504 (1947-97)	2504 (1947-97)
Panel B: Explaining ΔIncumbent *Party* Vote Share (OLS) $\Delta Incumbent\ vote\ share_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * Oil\ Shock_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * Oil\ Shock_{s,t}) + \varepsilon_{s,t}$		
ΔNational emp. gap ***Rule of thumb signal extraction (λ)***	.15 (.12) [.12]	.15 (.12) [.12]
ΔState emp. gap caused by oil shock ***Sophisticated signal extraction (δ)***	1.13*** (.37) [.40]	.89* (.53) [.57]
Residual ΔState emp. gap ***Returns to competence (α)***	.32** (.14) [.15]	.38*** (.14) [.14]
Adjusted R^2	.017	.014
Sample	636 (1947-97)	636 (1947-97)
Panel C: Explaining Whether Incumbent *Governor* is Re-elected (Probit[a]) $Incumbent\ elected_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * Oil\ Shock_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * Oil\ Shock_{s,t}) + \varepsilon_{s,t}$		
ΔNational emp. gap ***Rule of thumb signal extraction (λ)***	1.7* (0.9) [0.9]	1.6* (0.9) [0.9]
ΔState emp. gap caused by oil shock ***Sophisticated signal extraction (δ)***	4.3* (2.4) [2.7]	3.9 (3.1) [3.5]
Residual ΔState emp. gap ***Returns to competence (α)***	2.8*** (1.0) [1.1]	2.9*** (1.0) [1.0]
Pseudo R^2	.027	.026
Sample	356 (1950-88)	356 (1950-88)

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively.
(Robust standard errors in parentheses.) [Bootstrapped standard errors in square brackets.]
Changes in the employment gaps are measured over the two years leading up to the election.
[a] Probit coefficients report the marginal change in the probability of re-election for a marginal change in the employment gap, evaluated at cell means. Re-elected incumbents are coded to a value of 1. Incumbents who lost a primary race, a general election, or who decided not to run again are coded as 0. Incumbents who ran for higher office, or were barred from re-election by term limits, are dropped from the sample.

Table 6: IV Tests: Differential State Sensitivities to the National Cycle

Panel A: First Stage Regression (2-year ended changes, % points)		
	Instruments for ΔState Employment Gap – ΔUS Employment Gap	
	ΔNational Employment Gap Interacted with state dummies $\Delta State\ emp\ gap_{s,t} - \Delta US\ emp\ gap_t = \sum_{s \in states} \mu_s + \beta_s (State_s * \Delta US\ emp\ gap_t)$	**ΔNational Employment Gap Interacted with state industry shares (measured in the 1940s)** $\Delta State\ emp\ gap_{s,t} - \Delta US\ emp\ gap_t = \sum_{i \in industries} \theta_i (Ind.\ share_{s,i}^{1940s} * \Delta US\ emp\ gap_t)$
Estimated coefficients	Average: 0 State Range: -2.4 to +0.9	Average: 0 State Range: -2.1 to +0.6
Most pro-cyclical	CT, IN, KY, MI, OH	Manufacturing, services, wholesale
Most muted (or counter-cyclical)	AK, HI, ND, NM, OK, SD, WY	Construction, farm, finance, government
Adjusted R^2	.12	.11
Sample (n)	2504 (1947-97)	2504 (1947-97)
Panel B: Explaining ΔIncumbent *Party* Vote Share (OLS) $\Delta Incumbent\ vote\ share_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * \Delta National\ employment\ gap_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * \Delta National\ employment\ gap_t) + \varepsilon_{s,t}$		
ΔNational emp. gap ***Rule of thumb signal extraction (λ)***	.15 (.12) [.12]	.17 (.12) [.12]
β_s*ΔNational emp. gap ***Sophisticated signal extraction (δ)***	.54* (.30) [.34]	.80** (.34) [.36]
Residual ΔState emp. gap ***Returns to competence (α)***	.40*** (.14) [.14]	.39*** (.14) [.14]
Adjusted R^2	.012	.013
Sample	636 (1947-97)	636 (1947-97)
Panel C: Explaining Whether Incumbent *Governor* is Re-elected (Probit[a]) $Incumbent\ elected_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * \Delta National\ employment\ gap_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * \Delta National\ employment\ gap_t) + \varepsilon_{s,t}$		
ΔNational emp. gap ***Rule of thumb signal extraction (λ)***	1.6* (0.9) [0.9]	1.7* (0.9) [0.9]
β_s*ΔNational emp. gap ***Sophisticated signal extraction (δ)***	3.4 (3.0) [3.1]	5.2* (3.2) [3.3]
Residual ΔState emp. gap ***Returns to competence (α)***	2.9** (1.0) [1.1]	2.8*** (1.0) [1.0]
Pseudo R^2	.026	.027
Sample	356 (1950-88)	356 (1950-88)

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively.
(Robust standard errors in parentheses.) [Bootstrapped standard errors in square brackets.]
Changes in the employment gaps are measured over the two years leading up to the election.
[a] Probit coefficients report the marginal change in the probability of re-election for a marginal change in the employment gap, evaluated at cell means. Re-elected incumbents are coded to a value of 1. Incumbents who lost a primary race, a general election, or who decided not to run again are coded as 0. Incumbents who ran for higher office, or were barred from re-election by term limits, are dropped from the sample.

Appendix A: Summary Statistics

	Mean	Std. Dev.	Min.	Max.	Sample
Political Variables					
Incumbent party's vote share	53.9%	8.8%	23.6%	79.8%	636 (1947-97)
ΔIncumbent party's vote share	-3.3%	9.5%	-38.5%	27.8%	636 (1947-97)
Incumbent governor re-elected	56.7%	49.6%	0	1	356 (1950-88)
Democrat incumbent	53.5%	49.9%	0	1	636 (1947-97)
President's party	53.8%	49.9%	0	1	636 (1947-97)
Economic Variables					
ΔReal income per capita; BEA series *(2-year ended % change)*					
State	4.1%	5.5%	-23.3%	53.8%	2504 (1947-97)
State-National	0.2%	4.0%	-22.3%	44.5%	2504 (1947-97)
National	4.0%	4.0%	-13.9%	12.4%	2504 (1947-97)
ΔEmployment Gap, Constructed from BLS Non-farm Payrolls as per Appendix B *(2-year ended change)*					
State	0.0%	4.0%	-27.0%	15.8%	2504 (1947-97)
State-National	0.0%	3.0%	-31.5%	19.7%	2504 (1947-97)
National	0.0%	2.9%	-5.2%	6.4%	2504 (1947-97)
ΔReal Housing Prices; OFHEO series, deflated using chain GDP deflator *(2-year ended % change)*					
State	1.2%	9.8%	-23.3%	49.4%	800 (1982-97)
State-National	0.0%	9.6%	-30.2%	42.1%	800 (1982-97)
National	1.1%	4.9%	-10.0%	12.1%	800 (1982-97)
ΔReal Oil Prices; Crude Petroleum PPI, deflated using chain GDP deflator *(2-year change, log points)*					
ΔLog Real Oil Prices	0.013	0.231	-0.719	0.582	52 yrs (1947-98)

Appendix B: The Employment Gap

The basic data are annual average non-farm payrolls data by state compiled by the BLS. These data provide a consistent and reliable indicator of employment going back to the pre-war period. The employment gap is defined as the deviation of the natural log of employment from its trend level, which is estimated by applying a Hodrick-Prescott filter to the data. The top figure shows the construction of a national employment gap. Next the estimated employment gap and the unemployment rate are shown. The final figure shows that constructed state employment gaps co-move closely with the unemployment rate over the period in which both are available.

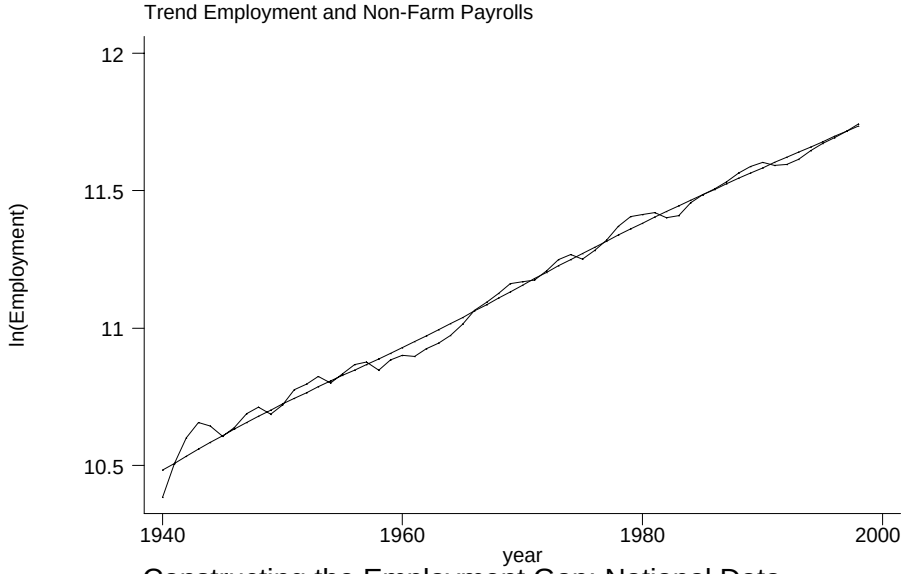

Constructing the Employment Gap: National Data

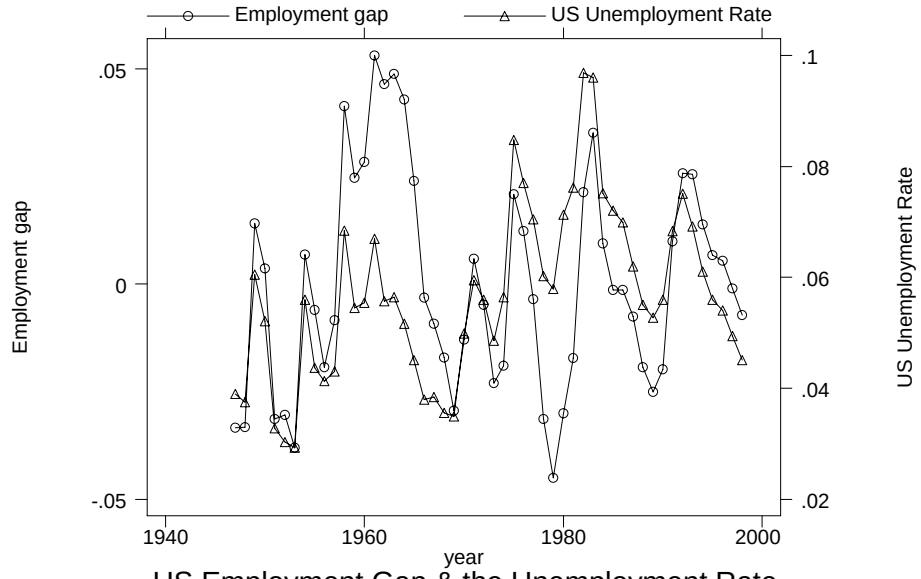

US Employment Gap & the Unemployment Rate

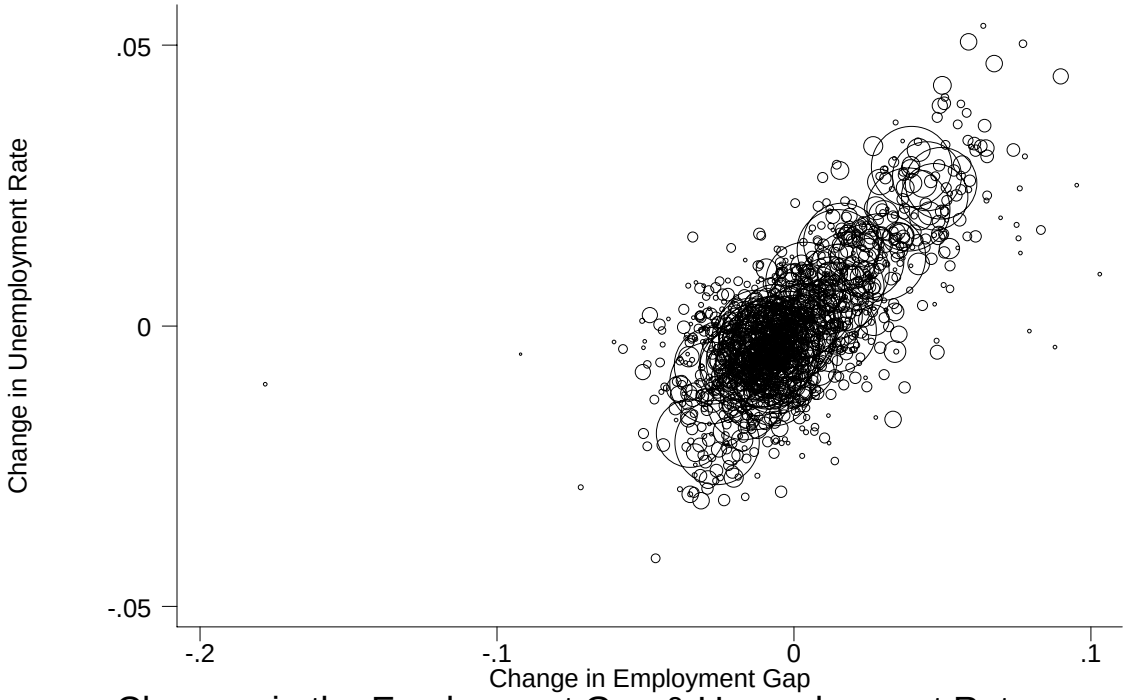

Changes in the Employment Gap & Unemployment Rate

Despite these similarities, there remain important conceptual differences between the employment gap and the unemployment rate. By construction the employment gap has a mean of zero for all states, and is mean-reverting. Further, the employment gap contains only business cycle and higher frequency variation.

Appendix C: Are Oil Shocks Exogenous?

A key assumption in Table 5 is that oil prices are unaffected by state governors, and hence voters should filter outcomes due to oil shocks from their assessments of the state's economic performance. This assumption is contestable. The main market player prior to OPEC was the Texas Railroad Commission, a Texas state agency. The Commission started regulating oil production in Texas in the late 1920s. In 1935, the federal government undertook to enforce in interstate commerce the production quotas set by the Texas Railroad Commission and its sister state agencies. Generally, the commission fixed quotas to keep the oil price steady. The commission lost traction in the 1970s and production quotas were entirely removed in March 1971. This appendix shows that the attribution errors highlighted in Table 5 do not simply reflect Texan voters voicing their displeasure at the decisions of the Texas Railroad Commission. Column one simply reproduces the results from Table 5. The second column drops Texas from the sample. The third column focuses only on the post-OPEC era. The fourth column follows Hamilton (2000) and is identified not by (potentially endogenous) oil prices, but by plausibly exogenous disruptions to the world oil supply.

Panel A: Explaining ΔIncumbent *Party* Vote Share (OLS)

$$\Delta Incumbent\ vote\ share_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * Oil\ Shock_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * Oil\ Shock_{s,t}) + \varepsilon_{s,t}$$

Instrument for Luck	**Oil Shock interacted with state fixed effects**				**Oil Shock interacted with historical industry shares**			
Column No.	*(A1)*	*(A2)*	*(A3)*	*(A4)*	*(B1)*	*(B2)*	*(B3)*	*(B4)*
ΔNational emp. gap	.15	.15	.16	.15	.15	.15	.13	.15
Rule of thumb signal extraction (λ)	(.12)	(.12)	(.23)	(.12)	(.12)	(.12)	(.23)	(.12)
ΔState emp. gap caused by oil shock	1.13***	1.18***	1.40***	.94*	.89*	.89*	1.12*	.97*
Sophisticated signal extraction (δ)	(.37)	(.39)	(.40)	(.53)	(.53)	(.53)	(.62)	(.57)
Residual ΔState emp. gap	.32**	.32**	.32	.37***	.38***	.38***	.49*	.39***
Returns to competence (α)	(.14)	(.14)	(.28)	(.14)	(.14)	(.14)	(.25)	(.14)
Sample	636	624	250	636	636	624	250	636

Panel B: Explaining Whether Incumbent *Governor* is Re-elected (Probit[a])

$$Incumbent\ elected_{s,t} = \lambda\ \Delta National\ employment\ gap_t + \delta(\beta_s * Oil\ Shock_t) + \alpha(\Delta State\ employment\ gap_{s,t} - \Delta National\ employment\ gap_t - \beta_s * Oil\ Shock_{s,t}) + \varepsilon_{s,t}$$

ΔNational emp. gap	1.7*	1.6*	0.7	1.6*	1.6*	0.6*	1.4	1.6*
Rule of thumb signal extraction (λ)	(0.9)	(0.9)	(1.4)	(0.9)	(0.9)	(0.9)	(3.7)	(0.9)
ΔState emp. gap caused by oil shock	4.3*	5.4**	5.4**	0.2	3.9	4.3	11.6	1.5
Sophisticated signal extraction (δ)	(2.4)	(2.4)	(2.8)	(3.1)	(3.1)	(3.2)	(8.7)	(3.9)
Residual ΔState emp. gap	2.8***	2.8***	2.5	3.3***	2.9***	3.0***	8.3*	3.0***
Returns to competence (α)	(1.0)	(1.1)	(1.9)	(1.0)	(1.0)	(1.0)	(4.4)	(1.0)
Sample	356	346	112	356	356	346	112	356

***, **, * denote statistically significant at 1%, 5%, and 10%, respectively.
(Robust standard errors in parentheses.) Further details on this table, see notes to Table 5.
First Stage Regressions: For columns (1)-(3), first stage regressions are as shown in Panel A of Table 5. Column (4) replaces oil price changes with Hamilton's estimates of exogenous disruptions to world oil supply (lagged 2 years). This measures annual world oil production losses due to the following events: 1956 Suez Crisis (10.1%), the 1973 Arab-Israeli War (7.8%), the 1978 Iranian Revolution (8.9%), the 1980 Iran-Iraq war (7.2%), and the 1990 Gulf War (8.8%). See Hamilton (2000) for further details.
Specifications: (1) baseline specifications from Table 5; (2) drops Texas from the sample; (3) drops data from 1947-74; (4) full sample, but replaces oil prices as the instrument for luck with Hamilton's measure of exogenous disruptions to world oil supply.