Letting the Good Times Roll: A Theory of Voter Inference and Experimental Evidence*

John W. Patty† Roberto A. Weber

November 11, 2002

Abstract

This paper explores a psychologically motivated model of belief formation in a political context. Using a retrospective voting framework, we specifically examine the implications of a common inference bias in which voters overweight the effect of an incumbent's unobserved effort on realized outcomes. This bias is motivated by and consistent with the fundamental attribution error in social psychology, whereby people over-attribute the cause of observed outcomes to personal and dispositional causes and underweigh situational causes. We provide experimental evidence of this bias and show that it leads to reduced incentives for the politician to exert effort on the voters' behalf.

Preliminary and incomplete. Please do not cite or quote.

*We would like to acknowledge the comments of Robyn Dawes, Richard McKelvey, John Miller, Tom Palfrey, and the audiences at Carnegie Mellon University and the 2000 meetings of the Economic Science Association in New York City.

†Corresponding author: *jwpatty@cmu.edu*. Patty gratefully acknowledges the financial support of the Alfred P. Sloan Foundation and the Division of Humanities and Social Sciences at the California Institute of Technology. Weber gratefully acknowledges the financial support of Carnegie Mellon University. The experiments in this paper were funded by a faculty development grant from Carnegie Mellon.

1 Introduction

There has been a great deal of debate about how voters determine their preferences between candidates. A large part of this debate has to do with the extent to which voters behave rationally when forming attitudes and casting votes. For instance, one view of voter attitude formation, retrospective voting, holds that voters reward positive prior performance without taking into account expectations of future performance (Kramer, 1971; Norpoth, 1996). The opposite view, prospective voting, holds that voters are more sophisticated and take into account information on past outcomes only to the extent that it has predictive value for future performance (Chappell, Jr. and Keech, 1985; MacKuen, Erikson and Stimson, 1992; Suzuki and Chappell, Jr., 1996). Related work examines the extent to which voters vote myopically out of concern for personal economic outcomes (Fiorina, 1978) or whether they care about the economy as a whole (Kinder and Kiewiet, 1979). Finally, a stream of work explores how voters process information to form attitudes (Gant and Davis, 1984; Lodge, Steenbergen and Brau, 1995).

This paper adds to the political science literature by considering belief formation within a simple, experimental setting that replicates situations where voters make inferences about the effort of elected officials, and developing a model that comports well with the behavior observed in the experiments. More specifically, this paper explores how individuals attribute outcomes to the unobservable efforts of a leader. We examine a classic “principal-agent” problem in which voters’ monetary rewards are determined both by a random process and by the effort exerted by a leader. The leader seeks to maximize the number of voters who believe that he or she exerted effort. However, it is costly to the leader to exert the effort, and the voters are not able to directly observe the level of effort expended by the leader. After observing an outcome – which is determined by both the random process and the leader’s effort – the voters make guesses about whether or not the leader exerted effort. The voters are rewarded for guessing correctly whether the leader exerted effort or not.

When inferring whether or not the leader exerted costly effort on the voters’ behalf, a substantial proportion of the subjects in our experiments reveal beliefs that are “strategically naive” in the sense that they neglect the incentives of the leader. Specifically, they over-emphasize the importance of the observed outcome, neglecting the fact that the leader’s actual effort is a very small component of this outcome. We develop a formal model of belief formation that is consistent with this behavior.[1] Both our experimental findings and the model presented in this paper are consistent with a body of existing research in social and political psychology. In particular, the behavior we observe and model is reminiscent of the so-called "fundamental attribution error" or the "correspondence bias," both of which involve an observer over-estimating the effect that an actor had on observed outcomes.

1.1 Related Social and Political Psychology Research

One motivation of this research is to explore the effect of two inferential biases – the fundamental attribution error and the correspondence bias – on voter belief formation and politicians' incentives. The fundamental attribution error refers to the tendency to over-attribute observed behavior and outcomes to personal rather than situational causes (see Ross and Nisbett, 1991; Ross, Amabile, and Steinmetz, 1977) and the correspondence bias refers to the tendency to assume stable, personal qualities about individuals based on observed outcomes or behaviors that are entirely situation driven (see Gilbert and Malone, 1995; Voonk, 1999). One classic study demonstrating this type of bias was conducted by Ross, Amabile, and Steinmetz (1977). In their study, subjects watched while two people played a simple trivia quiz game in one of two roles: questioners and responders. The questioners had been instructed to come up with challenging, but not impossible, questions to ask the respondent. After observing questioners generate some questions which the responders were unable to answer correctly, subjects were asked to rate the general knowledge of the people in both roles. Even though the role presents questioners with a large situational advantage in revealing knowledge, subjects rated questioners as significantly more knowledgeable than respondents. Subjects ignored the situational advantage and instead attributed the fact that questioners had knowledge that responders did not to a personal difference.

More closely related to the topic of this paper, several experimental studies have found that the fundamental attribution error affects perceptions of leadership quality (Weber, et al., 1999; Mitchell, Larson and Green, 1977; Lord, et al., 1978; Staw and Ross, 1980) in organizational settings. Across these studies, leaders associated with good outcomes are perceived as effective leaders, even when they

[1] As discussed below when we report the experimental results in detail, most subjects whose behavior was not consistent with this model were of two types: voters who essentially always believed that the leader exerted effort and voter who essentially never believed that the leader exerted effort.

have little or no effect on outcomes. Leaders associated with poor outcomes, however, are typically blamed for being responsible even when they similarly had no effect on determining the outcomes. In most of these experiments, subjects briefly observe a leader in action, and then observe the outcomes that result. Typically, the situations and processes generating the outcomes are very simple.[2] That observers make strong attribution errors concerning the effects of even limited leadership in simple situations, suggests the possibility that these errors might become greater when leadership is more visible and observers have even less knowledge of the process by which outcomes are generated. An example of such a situation is the evaluation of political officeholders on complex issues by voters.

The application of biases such as these to voter evaluations is not new. There is some research in political psychology on voter attributions resulting from processes similar to the fundamental attribution error and correspondence bias. For instance, Rapoport, Metcalf and Hartman (1989) report experiments in which voters made attributions about candidate's personal traits from stated issue positions. In recent, intriguing experimental work, Kanthak and Morton (2000) (2000) examine how individuals apply demographic characteristics of canddiates to infer the candidates' ideological positions. The problem they study is closely related to that studied in this paper in an information-processing sense: attribution in both settings is based on information which is costly to gather and/or process. However, the motivations behind the two studies are very different. Kanthak and Morton's results apply specifically to the evaluation of likely position stands, which is necessary for voters to vote informatively in a prospective sense, while our results apply most directly to ex post evaluation of the performance of incumbents, which is necessary for voters to vote informatively in a retrospective sense. Seen from this viewpoint, our work complements the research of Kanthak and Morton, as it seems highly likely that voters in real elections use a combination of prospective and retrospective evaluations in determining their vote choices.

In addition, a few empirical studies have acknowledged and dealt with the possibility that voters may differentially credit or blame elected officials depending on the voters' perceptions of the extent to which that official has control over outcomes. For instance, Hibbing and Alford (1981) find that the prevalence of

[2]For instance, in the Weber, *et al.*, study, leaders were subjects in an experiment randomly selected to give a short speech urging other subjects to coordinate in a game in which coordiantion was either very easy or very hard. The treatment variable ("leadability" of the situation) was uncorrelated with initial leadership perceptions elicited immediately after the leader spoke. After the groups played the game for several periods and groups in one treatment succeeded and groups in the other treatment failed, however, there was a significant difference in leadership ratings.

economic voting is affected by the ability of candidates to have affected economic outcomes. Powell, Jr. and Whitten (1993) conducted a country level study of economic voting, measuring the extent to which incumbent parties' votes are affected by economic conditions. They recognize that voters need to also take into account the extent to which the party in power is responsible for those outcomes. Using several measures of clarity of responsibility of the ruling party, they find an interaction between economic voting and responsibility: voters reward incumbent parties for good economic performance and punish them for poor economic performance only in countries with high degrees of clarity of responsibility. However, due to the difficulty in measuring the actual efforts of real-world politicians, none of these studies directly address the possibility that voters commit the error of overestimating the effect officials have on outcomes.

Research that directly explores this kind of voter error using empirical data is rare. In the only such paper we are aware of, Wolfers (2002) shows that voters in state gubernatorial elections credit or blame incumbent governors for outcomes that they do not control. Specifically, Wolfers finds that voters reward or punish incumbent governors for economic outcomes (such as oil price and national booms and recessions) that are largely independent of state policies.

The above studies all indicate that voters use information about the extent to which elected officials are responsible for outcomes. However, as Wolfers (2002) finds, such kinds of evaluations might be biased in the direction of overemphasizing the impact of officials and under-emphasizing the importance of situational factors in determining outcomes. The presence of this bias in voter inference is also supported by two additional pieces of evidence. First, the fundamental attribution error and correspondence bias are robust phenomena in decision making that hold across a wide variety of contexts. As mentioned above, several studies have shown that perceptions of a person's leadership quality are strongly affected by outcomes even when the leader has little or no control over them. Therefore, in any situation where evaluators are observing outcomes that may be affected by the actions of another individual, we expect individuals to over-attribute the effect to the person at the expense of situational variables. Second, it is quite possible that the actual effect of elected officials on outcomes of interest to voters (such as the state of the economy) is limited. Stigler (1973) makes this point, arguing that economic results are often beyond the control of elected officials. In addition, Salancik and Pfeffer (1977) find that the identity of city mayors has almost no effect on city budgets.

In this paper, we use experiments to directly test the hypothesis that voters are subject to the above bias in a controlled laboratory setting. The use of experiments allows us to eliminate any uncertainty that may exist about the control the leader has on outcomes, thereby allowing us to gauge whether the revealed beliefs of the other subjects are correct.

We also examine the effect of biased voters' attributions of outcomes on the actions of elected officials. To do so, we begin with a model in the next section based on the classic principal-agent problem in which a politician (the agent) must make a decision regarding whether or not to exert costly effort that increases the well-being of a voter (the principal). The voter is able to observe only her own well-being and then must infer whether or not the politician shirked his responsibility. Beginning with such a simple, well-known kind of model allows us a starting point from which we can model our experiments and add biased voter inference.

The paper proceeds as follows. In section 2 we describe a behavioral model of belief formation and apply it to the environment studied in this paper. We then describe and report our experimental results in Section 3. Section 4 contains a discussion of the implications of our model and experimental data for the study of electoral behavior. Finally, Section 5 concludes.

2 A Model of Strategically Naive Inference

This section describes a theoretical model of belief formation within a simple environment. We consider a single voter, i, who is faced with inferring what action (if any) was taken by a leader. The leader, denoted by L, observes a random shock (or state of nature), s, drawn from a set finite S according to a cumulative distribution function F and probability density (or mass) function f, and then chooses an action e from a finite set E.[3] The pair (e, s) results in an outcome, x, in a set X, according to a function $v : E \times S \rightarrow X$. The voter observes x and forms a posterior belief $\mu(\cdot|x)$, which is simply a probability measure on the set of possible actions E.

In our experiments, we do not ask the subjects to reveal their beliefs, μ, *per se*. Instead, we ask subjects simply to guess whether the leader exerted effort or not. We refer to this guess as the *stated inference* and denote it by $a(\mu(\cdot|x))$, or simply $a(x)$ when there is no chance of confusion. We suppose that, for each individual, $a(\mu(\cdot|x))$ is simply a maximum likelihood estimate of the action based upon his or

[3]Defining the model for general sets of actions and shocks is not difficult, but the notation is more awkward. In addition, there are no substantive differences between the finite and infinite cases.

her beliefs, μ. The stated inference of a voter is any selection[4] from the following correspondence.

$$a(\mu(\cdot|\tilde{x})) = \left\{ \tilde{e} \in E : e \in \arg\max_{\alpha \in E} \mu(e = \alpha | x = \tilde{x}) \right\}. \tag{1}$$

We are now in a position to define a behavioral model of belief formation. What we term *strategically naive beliefs* are defined by the following equation:

$$\mu(e = \tilde{e} | x = \tilde{x}) = \frac{\sum_{s \in S} f(s) 1[v(\tilde{e}, s) = \tilde{x}]}{\sum_{e' \in E} \sum_{s \in S} f(s) 1[v(e', s) = \tilde{x}]}. \tag{2}$$

We refer to such beliefs as strategically naive because they do not pay attention to the leaders' incentives. In particular, the validity of strategically naive beliefs depends on the behavior of the politician. It is straightforward to show that Equation 2 is correct if and only if the politician chooses e according to a uniform distribution on E (i.e., all actions are equally likely to have been chosen). We denote the behavior of the politician by $\sigma : S \to \Delta(E)$, where $\Delta(E)$ represents the set of probability distributions over E.

Fact 1 *The voter beliefs defined in Equation 2 are correct (*i.e., *μ is consistent with Bayes rule) if and only if $\sigma(s)$ is the uniform distribution over E for all $s \in S$: that is, $\sigma(e; s) = \sigma(e'; s') = \sigma = |E|^{-1}$ for all $e, e' \in E$ and all $s, s' \in S$.*

Proof: We prove sufficiency and leave necessity to the interested reader. Suppose that $\sigma(s)$ is the uniform distribution over E for all $s \in S$. Let μ' denote beliefs that are in accord with Bayes rule. These beliefs are given by:

$$\begin{aligned}
\mu'(e = \tilde{e} | x = \tilde{x}) &= \Pr[x = \tilde{x} \cap e = \tilde{e}] / \Pr[x = \tilde{x}] \\
&= \frac{\sum_{s \in S} f(s) 1[v(\tilde{e}, s) = \tilde{x}] \sigma(\tilde{e}; s)}{\sum_{e' \in E} \sum_{s \in S} f(s) 1[v(e', s) = \tilde{x}] \sigma(e'; s)} \\
&= \frac{\sigma \sum_{s \in S} f(s) 1[v(\tilde{e}, s) = \tilde{x}]}{\sigma \sum_{e' \in E} \sum_{s \in S} f(s) 1[v(e', s) = \tilde{x}]} \\
&= \frac{\sum_{s \in S} f(s) 1[v(\tilde{e}, s) = \tilde{x}]}{\sum_{e' \in E} \sum_{s \in S} f(s) 1[v(e', s) = \tilde{x}]} \\
&= \mu(e = \tilde{e} | x = \tilde{x}).
\end{aligned}$$

∎

[4] A *selection* from a correspondence $\Gamma : X \mapsto E$ is any function $g : X \to E$ such that for all $x \in X$, $g(x) \in \Gamma(x)$. This technicality is unnecessary if $\arg\max_{\alpha \in E} \mu(e = \alpha | x = \tilde{x})$ is single valued for each $\tilde{x} \in X$.

Put another way, strategically naive beliefs are correct *only if the leader's choice of action is invariant to the realization of* s. If the leader conditions his or her choice of action on the random shock (for example, working harder when "times are tough"), then strategically naive beliefs are necessarily incorrect at least some of the time.

2.1 Shirking Or Working?

We now apply strategically naive beliefs to the setting utilized in our experiments.[5] Suppose that $E = \{0, 1\}$ (*i.e.*, the leader can either "shirk" ($e = 0$) or "work" ($e = 1$)), $S = \{2, 3, \ldots, 40\}$, $v(e, s) = \beta e + s$ (where $\beta > 0$, $e \in E$, $s \in S$), and

$$f(s) = \begin{cases} (s-1)/400 \text{ for } s \leq 21 \\ (41-s) \text{ for } s \geq 22 \end{cases} .$$

(This distribution is sometimes referred to as the "tent distribution," which is, in this case, the distribution of the sum of two fair 20-sided dice.) In this setting the leader privately observes a shock s and then decides whether or not to add β units to this shock in generating the outcome, $x = v(e, s)$. Voters are assumed to know β.[6]

A voter with strategically naive beliefs (i.e., beliefs and stated inference defined by Equations 2 and 1) will infer that $e = 1$ (i.e., the leader "worked") if $x > 21 + \beta/2$ and only if $x \geq 21 + \beta/2$.[7] That is, in this simple environment, the stated inferences of a voter whose beliefs are generated by Equation 2 will appear to be generated by a "cutpoint" rule:

- $a(x) = 1$ if $x > 21 + \beta/2$.
- $a(x) = 0$ if $x < 21 + \beta/2$.
- Either $a(x) = 0$ or $a(x) = 1$ if $x = 21 + \beta/2$.

[5] The only difference between this setting and the experimental design is the number of voters. Here, to reduce notation, we consider one leader and one voter. In each experimental session, there were several voters and one leader.

[6] Alternatively, β might represent the voter's subjective belief about the effect of politician effort. For our present purposes, this distinction is immaterial, so long as the belief about β is not a function of the observed outcome.

[7] With F as defined, 21 is the median, mode, and mean of the distribution of s. The reason 21 is the "baseline" of the cutpoint rule is that 21 is the mode and the distribution is single-peaked. Nonsingle-peaked distributions will generate rules that have "multiple" cutpoints, (or, in other words, nonconnected regions of stated inference).

It is possible to prove that, given E and v as defined above and any $S \subset \mathbf{R}$, stated inferences that are consistent with a similar cutpoint rule will follow from Equations 2 and 1 whenever f is strictly quasi-concave (or, in other words, unimodal or single-peaked).

In order to complete the model, we now define the leader's incentives and examine when the stated inference of a voter with strategically naive beliefs will be incorrect. The leader receives a payoff according to the following function:

$$\pi(e, s, a, v) = 2a(v(s, e)) - e.$$

Thus, a strategic, payoff maximizing politician wishes to set $e = 0$ and have the voter infer that she set $e = 1$. However, setting $e = 1$ is a best response by the politician for any s such that $a(v(s, 0)) = v(s) = 0$ and $a(v(s, 1)) = a(s+\beta) = 1$. Supposing that the voter's stated inference is compatible with Equations 2 and 1 and that politician knows this, the best response for the politician is to set $e = 1$ for all s such that $-\beta/2 \leq s < \beta/2$ and $e = 0$ otherwise.

Thus if the voter forms her beliefs naively and the politician best responds to these beliefs, then we have that for all $x \geq 3\beta/2$, $a(x) = 1$ while $e = 0$, implying that the voter is incorrectly drawing an incorrect inference. Note that $x < \beta/2$ does imply that $e = 0$, so that the voter's beliefs are incorrect, but her stated inference about e is correct. Similarly, while $\beta/2 \leq x < 3\beta/2x$ does not imply that the politician set $e = 1$ with certainty, the voter's beliefs are correct and her stated inference is the maximum likelihood estimate of e.

Interpreting this example, the voter's inference rule as derived from Equations 2 and 1 is consistent with ascribing more credit to the politician for good outcomes than she deserves. This is because a strategically naive voter does not engage in second-order reasoning: the voter's inference rule does not allow for the possibility of strategic decision-making by the politician. Hence, we are confronted with a case in which the voter, though forming beliefs in a way such that the politician is modeled "as if" she were hardly an actor at all, makes stated inferences that overstate the politician's effect on observed outcomes.

We show in our empirical analysis that such rules characterize a significant proportion of our subjects' behaviors. We allow for heterogeneity in rules, which can thought of as allowing for heterogeneity in the subjects' perceptions of the leader's effect on outcomes. As derived above, a voters' cutpoint is an increasing function of β: voters who form strategically naive beliefs and believe that the effect of the leader (the size of the effect of effort) is larger will be more demanding of the leader by inferring that the leader exerted effort only when the outcome is very high. We now describe the experimental setting used to study voter inference.

3 Experiments

We conducted experiments in which we re-created the above game between a politician (labelled the "leader") and a voter in the laboratory. The experiment consisted of several rounds, in each of which the leader simultaneously and independently played the above game with several different voters.

Specifically, in each round, a random process determined an outcome of value to the voters. The leader observed this random outcome (but the voters did not) and then decided whether to add to the value of voters at a cost to herself. The voters then observed the final outcome, which included both the random component and the leader's action, and attempted to determine whether or not the leader had added to their value. These experiments recreate the political economy environment in our model and also real-world situations in which voters, who do not fully observe the actions of elected officials, care only about whether these leader exerted costly effort on their behalf.

3.1 Experimental design

We conducted two types of sessions. In one kind, there was one leader and seven voters. There were three of these sessions. The other session was collected in a large undergraduate classroom and consisted of one leader and 45 voters. In the classroom session, 7 of the 45 voters were selected at random after the experiment to be the 7 voters who would receive actual earnings and whose actions would influence the payoffs of the leader. In this way, the payoffs and basic structure of the game remained constant between the two different types of sessions (in all four sessions, the leader's payoffs were based on the actions of seven voters). In the first type of session, subjects were paid their earnings in cash at the end of the experiment. In the other session, subjects were told that seven of the voters would be randomly selected by the experimenter and these seven voters, plus the leader, would be paid at the beginning of the following class.

The random process in our model consisted of a roll of two twenty-sided dice. The outcome of interest to voters was the sum of the two dice (a number between 2 and 40), which represented the amount (in cents) that both voters and the leader received in that period. The distribution of this outcome is a "tent" distribution with mode at 21 (probability = 0.05) and least mass on the endpoints (2 and 40, probability = 0.0025). The leader rolled the die and was able to observe the outcome. The voters were not able to directly observe the outcome. This outcome can be thought of as part of some unidimensional measure of societal well-being

– such as an index of the economy, crime, etc. – where the value represents the outcome in the absence of any effects of policy actions by the elected official.

After rolling the die and recording the outcome, the leader decided whether or not to add to the total received by voters. If the leader added to this total, then every voter received 5 additional cents and the leader incurred a cost of 20 cents. If the leader did not add to the total, then she did not incur any cost.

The voters did not observe whether or not the leader added to the roll of the dice. Following the leader's decision, the experimenter announced the total outcome (the roll, plus 5 if the leader added). Voters recorded this outcome, which corresponded to the amount in cents that they received for that round. Voters then attempted to determine whether or not the leader added to the total. Voters guessed either "Yes" or "No" and received 50 additional cents if they correctly identified whether or not the leader had added to the total. Voters did not find out whether they had guessed correctly in any round until the end of the experiment.[8]

At the end of each round, the leader was informed of how many voters had voted "Yes," but not of their identities. The leader received 20 cents for every voter that voted "Yes." The incentives for the voters are such that voters want the leader to add to the total and want to be able to determine when the leader has added to this total. The incentives for leaders are such that they prefer every voter to believe that the leader worked without actually having to work. The payoff for leaders and voters are given by:

$$\pi_L = r - 20a + 20\sum_i g_i \pi_{V_i} = r + 5a + 50|g_i - a|,$$

where r is equal to the outcome of the roll of the two dice, a is a binary variable equal to 1 if the leader adds to the total and 0 otherwise, and g_i is a binary variable equal to 1 if the voter guesses "Yes" and 0 if the voter guesses "No."

Each session consisted of 23 rounds of the above game, plus two practice rounds for which players were not paid. In the practice rounds, the experimenter

[8]The first session we conducted, as well as a pilot session not reported in this paper, included a second "opposition" leader who rolled the dice in each round and then announced whatever number he wanted to claim was the role of the dice. At the end of each round, voters voted whether to keep the leader from the previous round or replace this leader with the opposition leader. Voting to replace was costly in that, if the leader was replaced, then every voter who voted to replace was charged 5 cents. This was intended as a second, behavioral, measure of voter beliefs. However, while some voters voted to replace the leader and using this variable instead of voter guesses yields similar results, the leader was never actually replaced. Therefore, this was dropped from the remaining sessions.

instructed the leader whether or not to add and voters were informed of whether or not they had guessed correctly at the end of each of the rounds. The experiments were conducted using undergraduates at Carnegie Mellon University.

3.2 Results: Correlation Analysis

We now present a traditional analysis of our data, examining correlations, conditional statistics, and logistic analyses of the leaders' and voters' behavior in our experiments. As opposed to our later analysis in which we allow for voters' inference rules to be heterogeneous, this means of analysis implicitly treats the voters' strategies as being homogenous.

While based on an admittedly strong homogeneity assumption, our analysis of descriptive statistics provide several means by which a bias in the voters' inferences can be observed. In particular, the correlation of voters' inferences that effort was exerted with the observed totals is strongly positive, while the correlation between the totals and effort exertion is negative in 3 out of the 4 sessions. In the one session in which effort exertion is positive, this correlation is smaller than the correlation of voters' inferences and the observed total (0.35 vs. 0.79). Table 1 presents descriptive statistics for each session.[9] The third column of the table contains the frequency with which the leader added in a particular session, while the fourth column contains the frequency with which voters guessed "Yes." The next column contains the correlation between total number of voters guessing "Yes" in a round and the actual action of the leader. The final two columns contain the correlations between the roll of the dice and the number of voters guessing "Yes" (column 6) and the actual behavior of the leader (column 7). These last two columns indicate that voters are not doing a particularly good job of guessing whether or not the leaders actually added – in two of the four sessions the correlations are negative – but that they are relying heavily on the value of the roll. This provides the first piece of support for our hypothesis that voters over-emphasize the role of outcomes relative to the actual ability of the leader to influence outcomes. Recall that the leader can only add 5 to the final outcome, meaning that the effect on the number that voters observe is not likely to be large. Since leaders are not more likely to add when the outcome of the die is higher (column 7), then the outcome observed by voters is due mostly to this random process and much less so to the actions of the leader. However, voters behave as if the value of this outcome is more diagnostic of whether the leader actually worked than it

[9] The entire dataset is available at http://www.andrew.cmu.edu/user/rweber/data.htm.

is (column 6).

Session	n	mean Effort	mean Guesses	corr(Guesses, Effort)
1	8	0.65	0.57	0.52
2	8	0.61	0.73	-0.13
3	8	0.57	0.66	-0.44
4	46	0.48	0.53	0.31
Session	n	corr(Guesses, Total)	corr(Guesses, Roll)	corr(Effort, Roll)
1	8	0.83	0.79	0.35
2	8	0.69	0.63	-0.59
3	8	0.88	0.86	-0.69
4	46	0.92	0.85	-0.04

Table 1: Summary of results by session

In Table 2, the summary statistics (averages of leader behavior and voter guesses) are divided for each session into rounds with low rolls (21 or below) and high rolls (above 21). As the results in the table indicate, leaders are less likely to add to the total (43% vs. 72%) when the roll is higher. Voters, however, are more likely to guess that the leader added to the total (77% vs. 47%) when the initial roll is higher.

Session	Roll	Number of rounds	mean a	mean g_i
1	Low	13	0.54	0.43
	High	10	0.80	0.76
2	Low	10	1.00	0.66
	High	13	0.31	0.78
3	Low	11	0.91	0.39
	High	12	0.25	0.90
4	Low	12	0.50	0.42
	High	11	0.45	0.63
Aggregate	Low	46	0.72	0.47
	High	46	0.43	0.77

Table 2: Summary of leader and voter behavior by roll

Another way to examine the bias in observed voter behavior can be seen from the results of several logistic regressions (in which the data from all 4 sessions has been pooled), shown in Table 3. The first regression examines the effect of the dice roll on the leaders' binary decision regarding whether to add to the outcome or not. These regression results indicate that leaders add to the total less when the initial roll is higher. The second regression indicates that voter guesses about whether or not the leader added to the total are predicted both by whether or not the leader actually added to the total and by the roll of the dice. In the final two regressions, the variables *Wrong Yes* and Wrong No are dummy variables equal to 1 if the voter incorrectly guessed Yes or incorrectly guessed No, respectively, and 0 otherwise. These two regressions indicate that voters are more likely to incorrectly believe that the leader added to the total when the total they observe is high and they are more likely to incorrectly believe that the leader did not add to the total when the total is low. All three regressions are consistent with the result that voters are overattributing high numbers to the leader having added and low numbers to the leader not having added.

Dependent variable:	Add[a]	Guess[ab]	Wrong Yes	Wrong No
Roll	-0.074^{***}	0.115^{***}		
	(0.284)	(0.010)		
Add		0.486^{***}		
		(0.131)		
Total			0.041^{***}	-0.024^{***}
			(0.009)	(0.009)
Constant	1.839^{**}	-4.620^{***}	-4.107^{***}	0.734
	(0.765)	(0.654)	(1.056)	(0.504)
N	92	1310	1445	1311
Log likelihood	-58.172	-737.513	-774.08	-668.43
Pseudo R^2	0.072	0.186	0.078	0.080

[a] - Regression includes fixed effects for session.

[b] - Regression omits observations where total < 7 or total > 40

* - $p < 0.1$; ** - $p < 0.05$; *** - $p < 0.01$

Table 3: Logistic regression results

3.3 Strategically Naive Beliefs and Cutpoint Rules

This section attempts to identify the rules applied by the subjects in our experiments. In so doing, we explicitly estimate cutoffs for voters that are consistent with the model of strategically naive belief formation presented above. The correlation analysis reported in the last section implicitly assumes that every voter (controlling for the experimental session) is generating observations from the same "population," or inference rule. We now relax this assumption and examine a class of inference rules that we refer to as c*-cutpoint rules*, of which the rule derived in Section 2 is an example. For any number c, a voter using a c-cutpoint rule infers that the politician worked if and only if the outcome is strictly greater than c.

For each subject in our experiments, we report the cutpoint that maximizes the percent of observed data in which correctly predicts the subject's behavior. Formally, for a subject i, let r_i denote the number of rounds in i's experimental session, x_i^t denote the outcome observed by subject i in round t , and $y_i^t = 1$ if voter inferred that the politician exerted effort in round t and $y_i^t = 0$ otherwise. Finally, let X_i and Y_i denote the $r_i \times 1$ vectors of x_i^t and y_i^t, respectively. Then, for a voter i, we define the *score* of a cutpoint c, given a vector of outcomes X and a vector of inferences Y , as

$$S_i(c) = \sum_{t=1}^{r_i} \left| \mathbf{1}[x_i^t > c] - y_i^t \right|.$$

Simply put, the score of a cutpoint rule for voter i is the number of times it "misses" when applied to i's behavior. After constructing this measure, we then find, for each subject i, the cutpoint c_i^* that minimizes S_i. When a set of two or more cutpoints yield the minimum score for a given voter, we arbitrarily take the minimum of this set.[10] A histogram of the estimates for the 66 subjects is displayed as Figure 2.

[Figure 2 about here.]

The mean cutpoint estimate is 20.84 and the modal estimate is 20. Several other characteristics of Figure 2 are also of interest. First, the cluster of 15 subjects whose minimum score cutpoint estimates are less than or equal to 11 are subjects who (nearly) always inferred that the leader exerted effort. We think of this type of subject as being either altruistic, since inferring that the leader worked amounts increases the leader's payoff, or simply naive. Similarly, the estimates that are in the upper 30s represent individuals who (nearly) always inferred that the leader did *not* exert effort. This behavior is consistent with a Nash equilibrium in which the politician never exerts effort and the voters never infer that she did.[11]

The next figure, Figure 3, displays the average accuracies of each cutpoint rule.[12]

[Figure 3 about here.]

The minimum score cutpoint rules perform well on average.[13] The total percent correctly predicted by the heterogeneous cutpoint estimates is 82.8%. Combined with the regression analysis above, we conclude from these results that subjects are using a heterogeneous set of systematic inference rules. In addition,

[10] In some sense, this choice biases our estimates away from the fundamental attribution error, insofar as the fundamental attribution error in this setting is thought of as a tendency by subjects to overattribute the outcome of the process to the effort of the politician.

[11] Note that, given the fact that we take the minimum of a set of optimal estimates, the estimate resulting from this "Nash" behavior will be equal to the highest observed outcome in the data set, which will in general not be 40.

[12] The reported average for a cutpoint c is the average percent correctly predicted by the c rule for all subjects whose minimum score cutpoint estimate is equal to c.

[13] Comparing Figures 2 and 3, the relatively low accuracies (for example, 17 and 34) correspond to small numbers of subjects (1 and 2, respectively).

	Percent Correctly Guessed[14]
Cutpoint	-0.00229
	(0.00164)
Distance from Mean Cutpoint [15]	-0.00471^{*}
	(0.00256)
Constant	0.621^{**}
	(0.0473)
df	63
F	2.145
R^2	0.0637

[16] [17]

$^{*}-p<0.1$; $^{**}-p<0.01$

Table 4: Regression of Performance on Estimated Cutpoint

many voters "demand too much" of their leader. We now examine the relationship between performance and estimated cutpoints. This relationship is displayed in Figure 4.

[Figure 4 about here.]

Figure 4 illustrates a vague nonlinear relationship between cutpoint estimate and performance. Namely, less extreme cutpoints seem to have higher average performance. The results of a regression of performance on both cutpoint and distance of the cutpoint from the mean cutpoint, 20.84, is reported in Table 4. The results indicate that voters estimated to have more extreme cutpoint strategies performed worse on average than those estimated to have moderate cutpoints.

Finally, we treat the estimated frequencies of cutpoint rules as the distribution in the population and examine the implications for an expected payoff maximizing leader facing a group of randomly drawn voters within the parameters used in our experimental sessions. Recall that the best response of the leader is to exert effort only if by such exertion the expected number of voters who will infer effort increases by 1/7 of the population. Using the frequencies displayed in Figure 2, we can form the cumulative distribution function of voter cutpoints and calculate what the leader's best response to each dice roll would be. Figure 5 displays both the cumulative frequencies of the cutpoint estimates and the best response region for a leader facing a population of voters whose c-cutpoint rules are drawn independently from our estimated distribution.

[Figure 5 about here.]

The most striking characteristic of Figure 5 is the width of the region in which the leader should exert effort when facing a random group of voters. For all dice rolls with total x where $15 \leq x \leq 26$, the proportion of our subjects who will infer that the leader exerted effort conditional upon observing $x + 5$ exceeds the expected proportion who will infer that the leader exerted effort after observing x by a margin large enough to justify the exertion of effort by a payoff-maximizing leader. Therefore, *it appears that the presence of heterogeneity in c-cutpoint rules allows the voters to increase the likelihood that an expected payoff maximizing leader will exert effort on their behalf*. Specifically, the probability of a dice roll summing to $15 \leq x \leq 26$ is 0.51, while homogeneous adoption of the optimal c-cutpoint rule (i.e., $c = 23.5$) results in the leader having an incentive to work when the sum of the dice is $19 \leq x \leq 23$, which occurs with probability 0.235 in our experimental setting. The heterogeneity of rules offers an incentive for the leader to work more than twice as often as he or she would if the voters adopted a homogeneous rule.[18]

4 Implications of the Voter Bias

While we have provided one stylized model of belief formation that we find compelling as the basis of a "types" explanation of the observed behavior that we report above, we think that the corrlelation analysis provide additional evidence that an overall bias is present in individual decision-making in this simple environment. In particular, we would succinctly define the bias we observe as follows: *individuals are increasingly likely to believe that the politician exerted effort as the observed outcome increases*.

It is easily shown that, unless the bias is fortuitously distributed across voters in a heterogeneous fashion (as allowed for in the estimation of heterogeneous cutpoints, above), the voter bias can reduce the incentive for a support-seeking politician to exert effort. Furthermore, if the voter's beliefs are not sufficiently sensitive to small changes in the observed outcome, then the politician may never possess an incentive to exert any effort at all. That is, voters simply giving more credit to the politician when outcomes are good than when they are bad does not necessarily imply that the politician has any incentive to improve the outcome.

Also, while not identical, the bias we observe in the experimental data is

[18] Note that, by optimal, we mean the homogeneous cutpoint rule that maximizes the proportion of dice rolls in which the leader has an incentive to exert effort. Other homogeneous cutpoint rules would lead to the leader having an incentive to work less frequently.

evocative of economic voting (in either the "pocketbook" or the "sociotropic" sense (Kinder and Kiewiet (1979))), in which voters are more likely to give credit to the politician when times are good than when they are bad. In this setting, the outcome is more dependent on the random process than with the effort exerted by the politician. Even in our experiments, this tendency was noted by at least one of the leaders. Leaders did not work very often when outcomes were very good (with a few exceptions) and did not work when outcomes were very bad (again, with a few exceptions). Effort was most likely to be exerted when initial outcomes were slightly above average, implying that the politicians were somewhat sophisticated in expending effort.

The results reported in Table 3 directly speak to the question of electoral accountability as a means to screen out shirking politicians. When times are good, shirkers are less likely to be forced out of office while, conversely, even dedicated politicians are more likely to suffer the electorate's wrath in bad times. While not the most surprising finding, it brings up an important normative and positive question: is representative democracy inefficient? For example, the incentives are less present for more capable individuals to enter politics when times are bad– i.e., precisely when such individuals are most needed. Similarly, when times are good, the incentive to enter politics is high for all comers – individual ability and/or effort may not matter as much as it should. Since politicians' electoral fates may be essentially beyond their control – perhaps being determined only by exogenous shocks – why would anyone seek a career in politics? Aside from the material advantages that successful politicians may derive from their offices – which may not be large compared to the material rewards offered by management positions in the private sector of most modern democracies – what would drive an individual to surrender the success of his or her career to the vagaries of exogenous shocks? An interesting possibility is that politicians themselves – including policy-driven ones – fall victim to a self-referential form of the fundamental attribution error, known as the "illusion of control." [19] The illusion of control is characterized by people believing that they have more control over somewhat exogenously determined outcomes than they actually possess. Obviously, individuals with an illusion of control with respect to public policy should be more likely to consider careers in politics than other individuals, *ceteris paribus*.[20]

[19] Indeed, an argument could be made that policy-motivated individuals are predisposed towards possessing the illusion of control.

[20] Since the voter bias examined in this paper is similar to the fundamental attribution error – which has been demonstrated in economic settings such as the evaluation of management effectiveness in firms – perhaps the voter bias is simply a fact of life that must be faced anytime an

In sum, any bias in citizens' evaluations of politicians is relevant to democratic theory as a whole. As Dahl writes, "[A]t a minimum ... democratic theory is concerned with processes by which ordinary citizens exert a relatively high degree of control over leaders...." (Dahl, 1956, p. 3) The bias examined and demonstrated in this paper predicts that the citizens' collective ability to control their leaders may be undermined due to the voters' mistaken belief that their leaders have more of an effect than they actually do. That is, the voter bias implies that the citizenry's control of elected leaders may be reduced or even negated precisely because the electorate believes the need to control politicians is greater than it truly is.

5 Conclusions

There are several implications of this research. Specifically, the cognitive bias(es) discussed in this paper pose serious normative and positive questions for the study of representative democracy. Normatively, do representative institutions serve the public interest? For example, is the economic performance of representative democracy robust to voters' cognitive biases? An important positive question posed by this research is with regard to the selection effects such biases might induce in the candidates for public office, are the candidates for office during good times generally of a different quality from those during bad times? Similarly, can such selection effects lead to vicious cycles in representative democracies?

The work that remains to be done in this area is vast. In addition to the philosophical and theoretical questions raised by the possibility of cognitively biased electorates, there are some serious empirical questions that need to be answered. First of all, how do voters carry out the attribution process when considering which incumbents are responsible for past policy outcomes? Secondly, how do voters construct mental estimates of the potential effect of proposed policies? In addition, joining these two questions, do voters attribute blame and credit for past performance in a manner that is consistent with how they consider the possible effects of proposed policy changes?

The model and experiments presented in this paper are purposely very sparse – there is no "context"; that is, there is little or no real-world political content presented to subjects or accounted for by the model. We proceeded in this manner in order to establish the baseline bias. Subjects do overcondition their inferences regarding the effort exerted by the leader on the observed outcomes. However, the

individual is in a principal-agent situation.

next step for this particular research agenda must be to include the possibility of additional signals by the leader (and possibly other potential leaders). How can leaders profitably frame unusually good or bad outcomes? How do individuals' beliefs react to sudden shocks (positive or negative) when dealing repeatedly with the same leader or set of leaders?

Finally, the previous section mentioned the analogy between our experimental setting and pocketbook voting. There are, of course, several dimensions to pocketbook, or economic, voting. One key dimension which our experiments do not examine is the distinction between voting based on aggregate and personal economic outcomes, since all voters receive identical payoffs. An interesting extension of our experiments would be to subject individual payoffs to unobserved idiosyncratic perturbations in addition to the systematic shock, inform voters of their own payoffs and the average payoff received, and observe how voters' inferences are driven by the two pieces of related information. It is obvious that a "rational" voter in this setting will conditional her beliefs only on the average outcomes. We suspect, however, that subjects will overly infer the expended effort level from their individual draws. We look forward to pursuing these topics in future work.

References

Alvarez, R. Michael. 1997. *Information and Elections*. Ann Arbor: University of Michigan Press.

Chappell, Jr., Henry W. and William R. Keech. 1985. "A New View of Political Accountability for Economic Performance." *The American Political Science Review* 79(1): 10-27.

Dahl, Robert A. 1956. *A Preface to Democratic Theory*. Chicago: University of Chicago Press.

El-Gamal, Mahmoud and David Grether. 1995. "Are People Bayesian? Uncovering Behavioral Strategies." *Journal of the American Statistical Association* 90(432): 1137-45.

Fiorina, Morris P. 1978. "Economic Retrospective Voting in American National Elections: A Micro-Analysis." *American Journal of Political Science* 22(2): 426-443.

Gant, Michael M. and Dwight F. Davis. 1984. “Mental Economy and Voter Rationality: The Informed Citizen Problem in Voting Research.” *The Journal of Politics* 46(1): 132-153.

Hibbing, John R. and John R. Alford. 1981. “The Electoral Impact of Economic Conditions: Who is Held Responsible?” *American Journal of Political Science* 25(3): 423-439.

Kanthak, Kristin, and Rebecca Morton. 2000. “Minding Payoffs and Cues: Measuring Voter Accuracy in Evaluating Elected Officials.” Mimeo, University of Arizona.

Kinder, Donald R. and D. Roderick Kiewiet. 1979. “Economic Discontent and Political Behavior: The Role of Personal Grievances and Collective Economic Judgments in Congressional Voting.” *American Journal of Political Science* 23(3): 495-527.

Kramer, Gerald. H. 1971. “Short-Term Fluctuations in U.S. Voting Behavior, 1896-1964.” *The American Political Science Review* 65(1): 131-143.

Lord, R. G., J. F. Binning, M. C. Rush and J. C. Thomas. 1978. “The effect of performance cues and leader behavior on questionnaire ratings of leadership behavior.” *Organizational Behavior and Human Performance* 21: 27-39.

Lodge, Milton, Marco R. Steenbergen and Shawn Brau. 1995. “The Responsive Voter: Campaign Information and the Dynamics of Candidate Evaluation.” *The American Political Science Review* 89(2): 309-326.

MacKuen, Michael B., Robert S. Erikson and James Stimpson. 1992. “Peasants or Bankers? The American Electorate and the U.S. Economy.” *The American Political Science Review* 86(3): 597-611.

Mitchell, T. R., J. R. Larson, Jr. and S. G. Green. 1977. “Leader behavior, situational moderators, and groups performance: An attributional analysis.” *Organizational Behavior and Human Performance* 18: 254-268.

Norpoth, Helmut. 1996. “Presidents and the Prospective Voter.” *The Journal of Politics* 58(3): 776-792.

Powell, Jr., G. Bingham and Guy D. Whitten. 1993. “A Cross-National Analysis of Economic Voting: Taking Account of the Political Context.” *American Journal of Political Science* 37(2): 391-414.

Suzuki, Motoshi and Henry W. Chappell, Jr. 1996. “The Rationality of Economic Voting Revisited.” *The Journal of Politics* 58(1): 224-236.

Raporport, Rondald B., Kelly L. Metcalf and Jon A. Hartman. 1989. “Candidate Traits and Voter Inferences: An Experimental Study.” *The Journal of Politics* 51(4): 917-932.

Salancik, G. R. and J. Pfeffer. 1977. “Constraints on administrative discretion: The limited influence of mayors on city budgets.” *Urban Affairs Quarterly* 12: 475-498.

Staw, B. M. and J. Ross. 1980. “Commitment in an experimenting society: A study of the attribution of leadership from administrative scenarios.” *Journal of Applied Psychology* 65(3): 249-260.

Stigler, George J. 1973. “General Economic Conditions and National Elections.” *The American Economic Review* 63(2): 160-167.

Weber, Roberto A., Colin Camerer, Yuval Rottenstreich and Marc Knez. 1999. “The Illusion of Leadership: Misattribution of Cause in Coordination Games”, forthcoming *Organization Science*.

Wolfers, Justin. 2002. “Are Voters Rational? Evidence from Gubernatorial Elections.” Stanford Graduate School of Business Research Paper No. 1730.

A A Derivation of the Cutpoint Model

This section derives the estimation of the cutpoint model examined in Section 4.3. A very nice example of a similar model and estimation technique is El-Gamal and Grether (1995), from which most of this section is derived.

We assume that each subject i is using a cutpoint inference rule, c_i, such that subject i infers that the leader exerted effort if and only if the observed outcome is (strictly) greater than c_i. However, each subject may “tremble” with probability $\varepsilon > 0$, in which case the subject makes a random inference, inferring that the leader exerted effort with probability 1/2 and inferring otherwise with probability 1/2.

In our experimental setting, there are essentially 45 such rules, in particular, $c_i \in C = \{1, 2, \ldots, 45\}$ for all subjects i. Note that $c_i = 1$ and $c_i = 45$ correspond to always and never inferring that the leader exerted effort, respectively.

We allow for heterogeneity of the rules across subjects but impose homogeneity of the error rate. Denoting the number of rounds in which subject i's behavior was observed and given a sequence of observed inferences $z^i = (y_1^i, \ldots, y_{T_i}^i)$, where $z_t^i = (x_t^i, a_t^i)$, where $x_t^i \in \{2, 3, \ldots, 45\}$ is the outcome observed by subject i in period t and $a_t^i \in \{0, 1\}$ is subject i's inference (with $a_t^i = 1$ corresponding to an inference by voter i that the leader did exert effort), and a cutpoint $c \in C$, define the variable

$$\begin{aligned} y_{c,t}^i &= 1 \ \begin{array}{l} \text{if } (a_t^i = 1 \text{ and } y_t^i > c) \\ \text{or } (a_t^i = 0 \text{ and } y_t^i \leq c) \end{array} \\ &= 0 \ \text{otherwise.} \end{aligned}$$

By this definition, $y_{c,t}^i$ is simply a tally for whether c correctly predicted i's behavior in period t. We now define the sufficient statistic $Y_c^i = \sum_{t=1}^{T_i} y_{c,t}^i$. Given an error rate ε and a rule c, the likelihood of a vector $y_c^i = (y_{c,1}^i, y_{c,2}^i, \ldots, y_{c,T_i}^i)$ is given by

$$f^{c,i}(y_c^i) = \left(1 - \frac{\varepsilon}{2}\right)^{Y_c^i} \times \left(\frac{\varepsilon}{2}\right)^{T_i - Y_c^i}.$$

Given observation of data on n subjects, with each subject i generating T_i observations, the maximum likelihood estimate of the homogeneous cutpoint model, which assumes that each subject is using the same rule c (i.e., $c_i = c_j$ for all subjects i and j) and possesses the same error rate ε, is given by

$$(\hat{c}, \hat{\varepsilon}) = \arg \max_{(c,\varepsilon) \in C \times [0,1]} \prod_{i=1}^{n} f^{c,\varepsilon}\left(y_c^i\right).$$

Maximum-likelihood estimation of the heterogeneous cutpoint model is similarly straightforward. Supposing that k cutpoint rules (not necessarily distinct) are allowed, let δ_{ij} equal 1 if subject i is using the jth rule and 0 otherwise. The maximum likelihood estimate of $\left(c^1, \ldots, c^k, \varepsilon, \{\delta_{ij}\}_{1 \leq i \leq n; 1 \leq j \leq k}\right)$ is given by

$$\left(\hat{c}^1, \ldots, \hat{c}^k, \hat{\varepsilon}, \left\{\hat{\delta}_{ij}\right\}_{1 \leq i \leq n; 1 \leq j \leq k}\right) = \arg \max_{c^1, \ldots, c^k, \varepsilon, \{\delta_{ij}\}_{1 \leq i \leq n; 1 \leq j \leq k}} \prod_{s=1}^{n} \prod_{h=1}^{k} \left(f^{c^h,\varepsilon}\left(y_c^i\right)\right)^{\delta_{sh}},$$

where $\delta_{ij} \in \{0, 1\}$ and $\sum_{j=1}^{k} \delta_{ij} = 1$ for all $1 \leq i \leq n$. It turns out that this estimate is quite straightforward to calculate in our framework, given the finite nature of the problem (i.e., the fact that the data is binary and there are finitely many rules to check for each individual). In particular, the maximum likelihood estimate of subject i's decision rule, given a k-tuple of possible rules $\left(c^1, \ldots, c^k\right)$ is simply $\hat{c}_i = \arg\max_{c \in \left\{c^1, \ldots, c^k\right\}} Y_c^i$. The maximum likelihood estimate of the ε is then simply twice the proportion of the deviations in the aggregate data set, calculated with each subject's corresponding maximum likelihood rule.

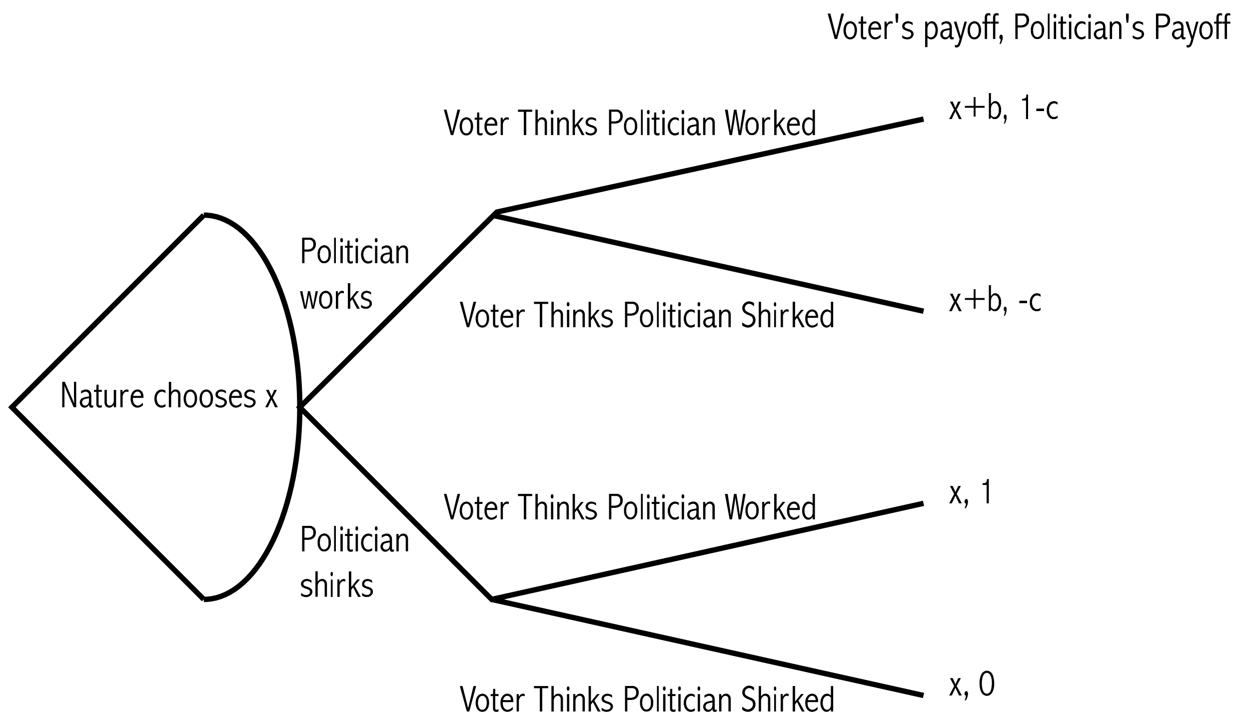

Figure 1: A Simple Principal-Agent Situation

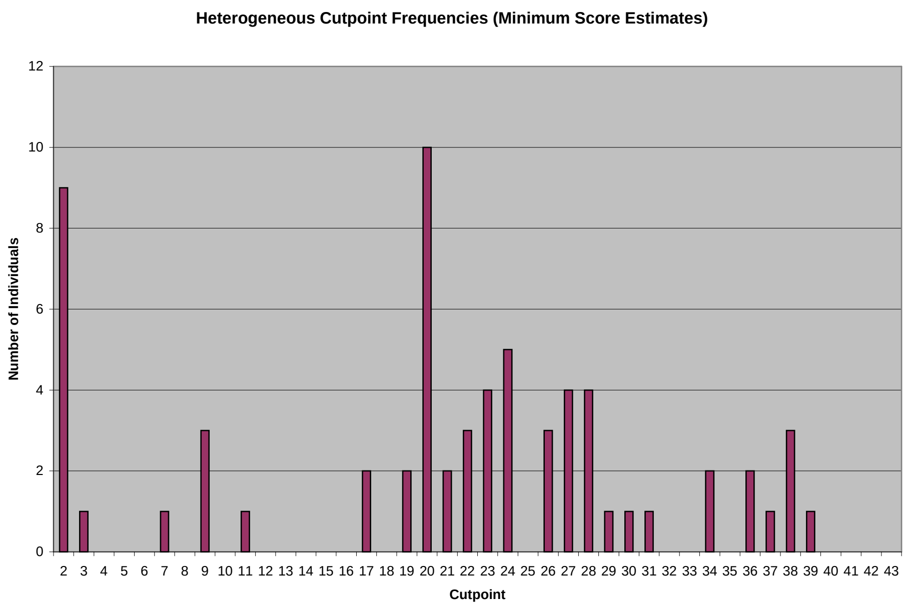

Figure 2: Histogram of Minimum Score Cutpoint Estimates

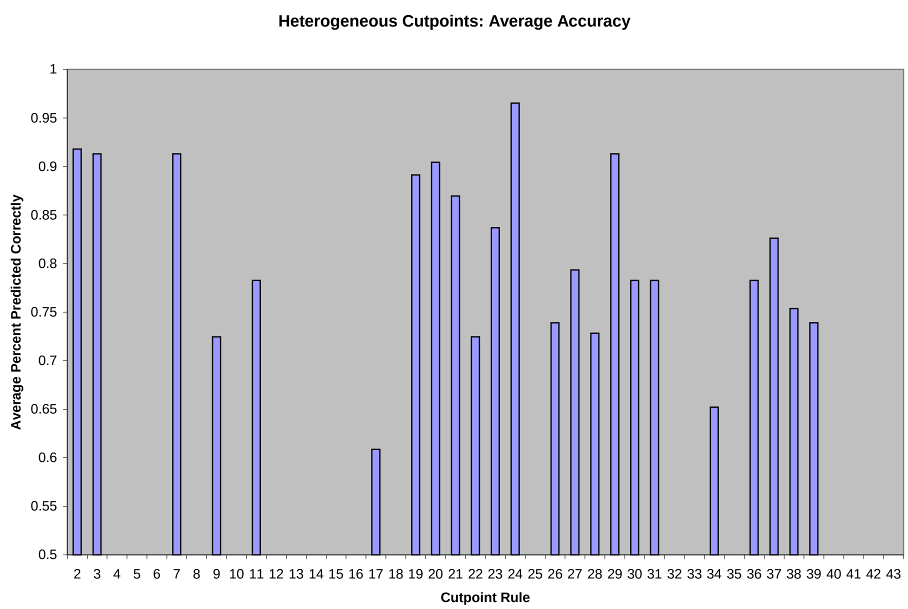

Figure 3: Average Accuracies of Cutpoints

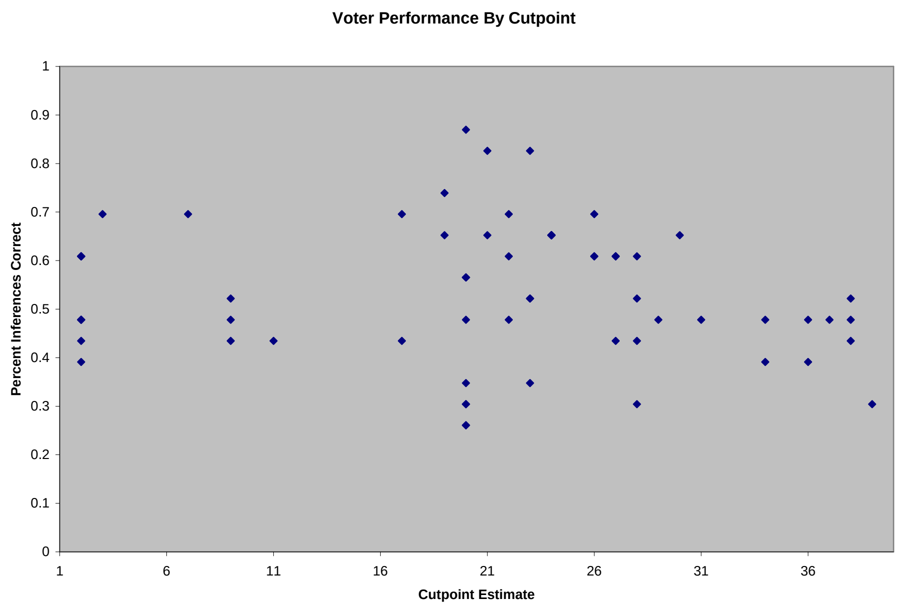

Figure 4: Average Performance of Estimated Cutpoints

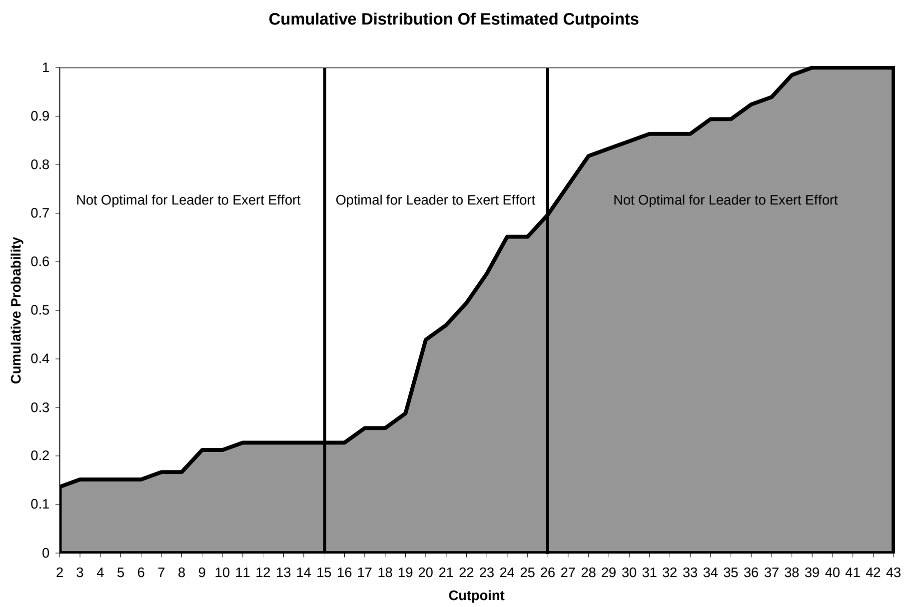

Figure 5: Cumulative Frequency of Cutpoints and Leader's Best Response Region