

Voting Correctly

RICHARD R. LAU *Rutgers University*

DAVID P. REDLAWSK *Educational Testing Service*

The average voter falls far short of the prescriptions of classic democratic theory in terms of interest, knowledge, and participation in politics. We suggest a more realistic standard: Citizens fulfill their democratic duties if, most of the time, they vote "correctly." Relying on an operationalization of correct voting based on fully informed interests, we present experimental data showing that, most of the time, people do indeed manage to vote correctly. We also show that voters' determinations of their correct vote choices can be predicted reasonably well with widely available survey data. We illustrate how this measure can be used to determine the proportion of the electorate voting correctly, which we calculate at about 75% for the five American presidential elections between 1972 and 1988. With a standard for correct vote decisions, political science can turn to exploring the factors that make it more likely that people will vote correctly.

The classic texts of democratic theory (such as J. S. Mill and Rousseau) assume that for a democracy to function properly the average citizen should be interested in, pay attention to, discuss, and actively participate in politics. The attention and discussion provide information about political affairs, which allows citizens to make political decisions (e.g., a vote) based on rationally considered principles reflecting their own self-interest and the common good. All citizens may not be able to live up to these standards—some may be too disinterested, or lack sufficient information or the skills to understand politics, and as a consequence they vote by habit or narrow prejudices, or do not vote at all. But as long as a clear majority lives up to these standards, the collective wisdom of the people will prevail.

Five decades of behavioral research in political science have left no doubt, however, that only a tiny minority of the citizens in any democracy actually live up to these ideals. Interest in politics is generally weak, discussion is rare, political knowledge on the average is pitifully low, and few people actively participate in politics beyond voting (e.g., Berelson, Lazarsfeld, and McPhee 1954; Campbell, Converse, Miller, and Stokes 1960; Converse 1964; Delli Carpini and Keeter 1996). And what good is even voting if for so many it is based on so little information?

The wide divergence between classic normative theory and political reality has led to two widely divergent

responses. On the one hand are those who accept both the normative theory and the empirical data and who conclude as a consequence that governments calling themselves "democracies" are not truly democratic. An apathetic public cannot possibly constrain government officials, this line of argument goes; instead, some capitalist power elite, military-industrial complex, and/or giant media conglomerate uses democratic institutions and a complacent citizenry to manipulate government policy toward its own ends (e.g., Bennett 1988, 1992; Burnham 1965, 1974; Fishman 1980; Gans 1979; Mills 1971).

On the other hand are those who accept the empirical evidence but revise downward the requirements of normative theory so that modern governments can still be considered mostly "democratic." For instance, Page and Shapiro (1992) show how aggregate public opinion can be fairly stable and "rational"—and even, perhaps, guide public policy—while based on mostly ill-formed "nonattitudes" among individual citizens. Fiorina (1981) shows the advantages of basing vote decisions on retrospective judgments of the party's past performance rather than spending the time to learn about the candidates' future policy proposals. Others argue that the widespread ignorance of and indifference toward politics typically seen in Western democracies is in fact a good thing, for it reduces social conflict and contributes to greater system stability (Berelson, Lazarsfeld, and McPhee 1954; McClosky 1964; Mueller 1992; Prothro and Grigg 1960). In other words, democracy still "works," and in fact may even work better, if only some *minority* of the population is attentive to politics, ideological in its thinking, and so on (see also Dahl 1961, 1989; Huntington 1968; Lindblom 1965; Schattschneider 1960).

Although we are somewhat in sympathy with each of these divergent responses, we take issue with the very point on which they both agree: Do the empirical data in fact require so drastic a revision of classic democratic theory? True, if modern citizens paid the type of attention to public affairs that Rousseau prescribed several centuries ago, they would do nothing else but follow politics. Such standards are unrealistically high and, we argue, not necessary for the average citizen.

Richard R. Lau is Professor of Political Science, Rutgers University, New Brunswick, NJ 08903. David P. Redlawsk recently received his Ph.D. in political science from Rutgers and is now at the Educational Testing Service, Rosedale Road, Princeton, NJ 08541.

This research was funded by National Science Foundation grants SBR 93-21236 to the first author and SBR-9411162 to the second author. The authors would like to thank Eric Johnson, who provided invaluable advice at an early stage in this research; Racheal Ankrah, Jennifer Holt, Jill Locke, John Manyo, and Jeff Schnug, who worked as experimenters; Rachele Brooks and Aletia Morgan, who served as expert judges; Licia DeVivo, who coded much of the open-ended data; and especially Paul Babbitt and Liz Felter, who served in the multiple roles of expert, project manager, experimenter, and reader of a preliminary version of this manuscript. We also acknowledge the very helpful comments of Jane Junn, Ben Peck, Gerry Pomper, and Richard Smith on earlier versions. Any remaining errors are, of course, our own.

Classic democratic theory prescribes active attention to and close scrutiny of government policy because, logically, it seems the only way that citizens can make "correct" decisions. If A (an active, attentive public) is necessary for B (democracy), and A is not true, then logically B cannot be true.

This syllogism holds only if we accept the premise that close attention to politics and the actions of government officials and the promises of competing candidates are *necessary* for correct voting decisions. But what if they are not? What if people can make reasonably good decisions, most of the time, without all the motivation and attention and knowledge required by classic theory?

Modern cognitive psychology teaches us that humans are limited information processors (Fiske and Taylor 1991, Lau and Sears 1986). People process only a small fraction of the information to which they are exposed. The human mind, although effectively unlimited in its long-term storage capacity, is severely limited in how much information can be kept in short-term or active memory at any given time (Simon 1979). As a consequence, both perception and storage of incoming information, as well as subsequent recall of that information from memory, are structured (and thus biased) by prior expectations or cognitive "schemata" that help determine what information is noticed, where it is stored, and how likely it is to be retrieved from memory. Thus, citizens do not have all the information about politics that is required of them by classic democratic theory, and they do not process that information in as logical a way as those theorists hoped, in large part because of strict cognitive limitations. It is not so much that we do a particularly bad job of processing political information; rather, we do an equally bad job of processing any type of complex information. If the same standards that classic democratic theory holds up for citizens were to be applied to any other area of human life, such as finding mates or buying cars or choosing colleges, then people would be found to be just as inept in those areas as well.

Most people nonetheless seem to make an adequate marriage, get a decent education, and make reasonably good automobile purchases. How is this possible, if people behave and make decisions in ways that are so far from the dictates of normative theory? Quite simply, human beings have adaptively developed a large series of cognitive heuristics or shortcuts that allow them to make "pretty good" judgments most of the time (Kahneman, Slovic, and Tversky 1982; Nisbett and Ross 1980; Tversky and Kahneman 1974). These heuristics "do sometimes lead people astray when they are overextended or misapplied . . . [but] people's intuitive inferential strategies are probably used appropriately and effectively in the great majority of cases" (Nisbett and Ross 1980, 255).

As a consequence, if we are going to make judgments about the "democratic" nature of different forms of government, we should do so at least initially on the basis of the quality or "correctness" of the political decisions citizens make within that system of government rather than on the basis of the ways in which

those decisions are reached. Democracy is not a simple form of government, and judgments about the nature of different governments that claim to be democratic should not be made in a simplistic, either-or manner. Certainly, "degrees" of democracy are possible; and we argue that a crucial criterion is the proportion of citizens voting correctly at any particular time, rather than the manner in which those vote decisions are reached. That is, if most people, most of the time, vote "correctly," then we should not be too concerned if those vote decisions are reached on the basis of something less than full information about the different policies espoused by different candidates, much less information about how government actually carries out policy decisions or who the important players are.

WHAT IS A "CORRECT" VOTE DECISION?

Determining the "correctness" of a vote decision is not an easy task. Who is to decide what is "correct"? We are reluctant to define what is "good" for everyone; even if we were not, we doubt that many people would be willing to accept our judgments. Instead, we begin by defining "correctness" *based on the values and beliefs of the individual voter*, not on any particular ideology that presumes the values and preferences which *ought* to be held by members of different social classes, for instance, and not on any larger social goods or universal values. Given the limitations of human cognitive abilities discussed above, however, we are equally reluctant to accept as "correct" any individual vote just because it is freely chosen by that individual, as Downs (1957) and his followers might. Instead, we adopt a theoretical middle ground by defining "correctness" based on the "fully informed" interests of individual voters. As Dahl (1989, 180-1) writes, "a person's interest or good is whatever that person would choose with fullest attainable understanding of the experiences resulting from that choice and its most relevant alternatives." (See also Bartels 1990, Connolly 1972, Delli Carpini and Keeter 1995, Lippmann 1955, Mansbridge 1983). Thus, we define a "correct" vote decision as one that is the same as the choice which would have been made under conditions of full information. Ideally, this determination can best be made subjectively, by the voter, on an individual basis.

This paper has several purposes. First, we describe experimental work that attempts to operationalize this ideal of fully informed voters determining for themselves what is the correct vote decision. Second, we use this operationalization as a means of validating an easily obtainable measure of correct voting. In so doing, we move from a completely subjective, individually determined definition to a more objective, expert-determined judgment of which candidate best matches the voter's own stated preferences. The experimental data are crucial for justifying and validating an objective, externally determined measure. Third, having obtained it, we briefly illustrate the use of such measures with the 1972-88 American National Election Studies (ANES). These latter results provide important predictive validity for this measure. Finally, we return

to the question of what ought to be required of citizens by democratic theory, once we have a better idea of just how well our disinterested, apathetic, uninformed citizens actually do in making their vote decisions.

A DYNAMIC, PROCESS TRACING METHOD FOR STUDYING THE VOTE CHOICE

We have designed an interactive experimental paradigm to study voter decision making that captures the crucial features of modern political campaigns: They are media-based; they provide an overwhelming amount of relevant information, some of which voters choose to expose themselves to, some of which comes to voters without any conscious decision to learn it, and much of which is simply missed; and they are dynamic, in the sense that information available today may be gone tomorrow (see Lau 1995 and Lau and Redlawsk 1992 and n.d. for more details). Our experiments involve a primary election with from two to four candidates in the Democratic and Republican primaries, followed by a general election campaign involving one candidate from each party.

In our experiment, brief descriptions of information available about candidates in the election (e.g., "Fisher's Stand on Welfare," "Martin's Family Background") scroll down a computer screen, and subjects (voters) must actively "access" the information in which they are interested by clicking on it with a mouse. This conscious search reveals the type of information which might be learned about a candidate from an opening paragraph in a newspaper or magazine article. Six of these descriptive labels appear on the screen at any given time, but only one item may be accessed at a time. Moreover, when an item is accessed the scrolling continues, hidden in the background, while subjects read the information they have chosen. Thus, an entirely new set of options faces a subject when s/he has finished reading the particulars of the item s/he chose and returns to the campaign. The information available at any point in a campaign is determined by a stochastic process within certain constraints meant largely to replicate the type of information available during American presidential campaigns.¹ At certain points, the computer screen is "taken over" by 20-second political advertisements for one of the candidates running in the election. Voters can close their eyes and ears when the ads come on, but they cannot do anything else relevant to the campaign until the commercial is over.

Our goal in designing this mock election simulation was to overwhelm people with information to force them to be selective in what they learn. The experiment that provided the data used here involved 212 distinct items which could be accessed during the primary campaign, including 26 individual policy stands, back-

ground information, and personality descriptions about each of six different candidates in the two parties' primary elections; endorsements by 11 different groups for candidates in each primary; the results of 24 different surveys conducted at various points during the campaign; and the results of five earlier primaries in each party. Eight to ten "televised" political advertisements also appeared during the course of the primary campaign. Thus, the simulation provided far more information than could possibly be handled by anyone, even under ideal conditions.

METHOD

Subjects

Subjects were recruited from the central New Jersey area during fall 1994. The only two provisos were that they be American citizens at least 18 years old (i.e., eligible voters) and not currently going to college. Subjects were recruited by ads in local newspapers and by approaching charitable organizations (such as churches, PTA chapters, the American Legion) that could provide experimental subjects in return for the subject payment fee of \$20. Subjects who responded directly to our advertisements were paid the \$20 upon completion of the experiment.

We recruited 303 subjects for the experiment; the analyses reported below rely on data from 293 of them.² We make no claims that this is a representative sample of central New Jersey voters, much less the nation as a whole. We were successful in recruiting a broad range of people, but our subjects were slightly more likely to be female, college educated, and retired than is true of the area from which we recruited (according to the 1990 Census).

Procedure

All subjects participated in the experiment individually and in a reasonably private setting where there would be no interruptions. The complete experimental procedure is sketched in some detail in Figure 1. For current purposes, the crucial features are these. (1) Subjects initially completed a fairly standard political attitudes questionnaire in which their political preferences, knowledge, and so on, were determined. (2) After the nature of a mock 1996 presidential election was described and subjects were given practice using

¹ To make these probabilities realistic, we first conducted an elaborate study of the prevalence of different types of information in newspapers during the 1988 presidential campaign (Lau 1992). We then modeled the probabilities after the actual prevalence of those types of information during the 1988 campaign.

² Of the 303 recruits, three were eliminated when they could not complete the study for time reasons (some subjects took more than 2½ hours to complete the experiment), two were eliminated because of computer errors, and another five were dropped because the experimenter, after running a subject, believed s/he had not taken the experiment at all seriously. Of the remaining 293 subjects, 60% were female; 84% were non-Hispanic whites; 22% had a high school education or less, while 52% were college educated; and 22% had family income under \$25,000, while 25% had family income over \$75,000. The median age was 52; half were currently employed, and one-third were retired; 30% were Catholic, 29% Protestant, 18% Jewish, with the remainder expressing some other or "no" religious preference. Forty percent were Democrats, 25% Republicans, and the remainder independents.

FIGURE 1. Outline of Experimental Procedure

1. Political Attitudes Questionnaire
Questions to measure subjects'
 - political preferences;
 - political interest, participation, knowledge, and media usage;
 - importance of different types of political information for 1992 vote choice;
 - background/demographic information(about 30–40 minutes)
2. Mock Election Campaign
 - a. Practice session accessing information with the mouse about 1988 presidential election (about 8 minutes)
 - b. Explicit instructions and 1996 campaign scenario; random assignment to different experimental conditions (hidden from subjects) (about 5 minutes)
 - c. Primary election campaign involving six candidates (about 22 minutes)
 - d. Vote in party's primary election; evaluate all six candidates; manipulation check on difficulty of choice (about 3 minutes)*
 - e. Break for party conventions; general election candidates selected (about 2 minutes)
 - f. General election campaign involving two candidates (about 12 minutes)
 - g. Vote in party's primary election; evaluate all six candidates; manipulation check on difficulty of choice (about 3 minutes)
3. Unexpected Memory Task
Subjects were asked to remember as much as they could about two general election candidates. (about 10 minutes)
4. Correct Voting Determination
Subjects were presented with complete information about two candidates from the primary (the one for whom they voted and, of the remaining candidates in that same party, the candidate closest to the subject on the issues) and asked to decide the one for whom they would have voted if they had obtained all this information when they actually had to make their choice (about 10–15 minutes)
5. Debriefing
Subjects' general impressions of the experiment were gathered; any remaining questions were answered; and so forth. (about 5 minutes)

*Experiment II subjects completed the unexpected memory task (step 3) at this point for candidates in their party's primary. They then completed a detailed guided protocol analysis, to explain why they had selected the items they chose to examine during the primary, before determining their correct vote choice (step 4).

the computer, (3) subjects registered to vote in either the Democratic or Republican primary. They then experienced a primary campaign involving six candidates (two from one party and four from the other) and much more information than can possibly be processed. (4) After the primary campaign, subjects voted for one of the candidates in their party's primary. Finally, after several activities of no concern here,³ (5) subjects determined if their vote was "correct."

Determining "Correct" Vote Choices

After the mock election campaign was over, subjects were questioned about their impressions of the experiment and were asked to complete one final task. The experimenter commiserated with each subject about how difficult the election had been but explained it was very important to know whether the subject thought

s/he had voted "correctly." The pictures of two candidates from the primary election were shown on the computer screen, the one for whom the subject voted and another from that same party. If only two candidates from the party ran in the primary, then those two were shown. For the four-candidate primary there was some discretion, however, and to make this choice as difficult as possible, the computer was programmed to select the candidate (from the three available) who was closest to the subject on the issues.

The experimenter brought out a notebook in which all the information about these two candidates was laid out side-by-side, so that it was very easy to compare. Explaining that it was not possible to learn everything about any candidate during the experiment, the experimenter asked the subject to examine the material carefully and decide if s/he still would have voted the same way if s/he had known all the available information about these two candidates. These instructions were designed to get subjects to take this final task seriously and to set up a context in which they would feel free to change their mind about their initial vote choice in the primary without feeling defensive or foolish about that choice.

After the subjects carefully considered all the information about the two candidates, if they decided they still would have voted as they did, then we classified

³ We actually ran two different experiments that diverged after the vote in the primary election, but in this paper we only use information from the primary election campaign, which was essentially identical across the two experiments. The experiments also involved several other random manipulations, but only the two- versus four-candidate manipulation affected the probability of a correct vote, and for present purposes the other manipulations are ignored. For simplicity we have collapsed together subjects from the two experiments.

them as voting "correctly." If they would have *changed* their initial vote, then we classified them as voting "incorrectly." This is our first, and probably most defensible, measure of correct voting.⁴

A Second Measure of Correct Voting

None of us can learn everything there is to know about the candidates in a real election, and few people are motivated to learn even what is readily available. Nonetheless, almost everyone *tries* to vote correctly, given what they have learned about the candidates (and parties) by election day. We now attempt to model how people could "naively" or "intuitively" go about making these decisions.

From the preelection questionnaire, we knew subjects' political preferences and policy stands. We also knew the candidates' stands, and those can be expressed on the same scales employed by subjects through the use of expert judges.⁵ Most important, we knew exactly the information to which a voter was exposed about the different candidates. We did not know precisely how that information was evaluated by each individual, but we could make very informed guesses based on our knowledge of the voter's preferences. Thus, we were in an excellent position to try to model the vote decisions that people actually made in our experimental elections.

In fact, assuming that people try to vote correctly, this is essentially a second measure of correct voting—a *normative* measure of naive vote preferences. It is based on the voter's own values, to be sure, but ultimately is determined by the authors rather than the voter. We consciously juxtapose the terms "normative" and "naive" in the description of what we are trying to model. This variable is "naive" or "intuitive" in that it is based on the voter's actual information-gathering strategies rather than any ideal, logical, or expert-determined process; but it is "normative" in that it is based on an objective evaluation of that information. Thus, an objective determination can be made of the candidate for whom the subjects *should* have voted,

⁴ Besides random error, there are at least three plausible reasons subjects may not be completely accurate in their own assessment of the correctness of their initial vote choice. One is self-presentation: Despite our efforts to make it acceptable to revise one's decision—and we were careful to avoid such words as "mistake" and "incorrect"—some subjects still may have been reluctant to do so. A closely related reason is avoiding postdecision dissonance, or any unpleasant internal state resulting from learning that one has made a bad decision. A third reason is fatigue. This final task was presented after an average 126 minutes of prior effort (the range was 93 to 160, with a standard deviation of about 13 minutes), and the material about the two candidates was almost twelve pages long, single spaced. Nonetheless we stressed the importance of this "final task," and most subjects made a very serious effort.

⁵ Seven "experts" (the two authors, four graduate students in political science at Rutgers, and one local elected official) read all issue positions taken by eight different candidates and rated them on seven-point scales, in keeping with the items and scales on the political attitudes questionnaire given to subjects. We used the mean rating of these seven judges, rounded to the nearest whole number, as the objective stand of each candidate on that issue.

given their own political preferences and the differential candidate information to which they were exposed.

Unfortunately, there is an almost infinite number of ways in which voters can combine information about the candidates with their own values and policy preferences to form candidate evaluations. After examining in some detail various methods that span the range of possibilities (but provide results very similar to those presented below), we provide data from one of the simplest procedures. We assume that voters average together the favorableness of all information they have learned about a candidate, weighting some parts of it more heavily than others. After comparing the summary evaluation of each candidate, voters should choose the candidate with the highest evaluation. This is essentially the method assumed by many popular models of the vote decision (e.g., Fiorina 1981; Kelley and Mirer 1974; Lau 1986, 1989; Lodge, McGraw, and Stroh 1989; Lodge, Steenbergen, and Brau 1995; Markus and Converse 1979).

Three types of information went into each candidate evaluation in our experiment. First, the favorableness of a candidate's issue stands was calculated by the directional method (Rabinowitz and MacDonald 1989), with the mean rating of our seven experts providing the stands (refer to footnote 5). Whenever a voter learned a candidate's position on an issue, and if the voter had expressed an opinion about it in our initial questionnaire, agreement or disagreement on that issue (rescaled to range from -1 to $+1$) was added into the candidate's total evaluation. Second, group endorsements learned by a subject were scored $+1$ if the subject liked the group (i.e., rated it above the mean of all groups evaluated and above the midpoint of the scale), -1 if the subject disliked the group, and 0 if the subject was indifferent toward the group. Third, the favorableness of personality descriptions learned about candidates and the attractiveness of their picture used in campaign ads were based on the ratings of 67 independent pretest subjects, again rescaled such that $+1$ was the highest possible evaluation and -1 was the lowest possible.⁶

Each bit of information was weighted according to subjects' judgments (from the preexperiment questionnaire) about how important different broad types of information were to their vote in the last presidential election. The weights ranged between 0 ("Not at All Important") to 1 ("Most Important"). Because every individual evaluation was scaled to range between -1 and $+1$, and the total evaluation was divided by the number of items upon which it was based, the summary evaluation itself ranges between -1 and $+1$. A detailed

⁶ Thus, most information learned about a candidate figured into the evaluation. We have no way of knowing how subjects evaluated some of the available information, however, particularly about candidates' personal background. Although we have no reason to suspect systematic biases, it is quite possible that some subjects preferred senators over governors, people from Florida over people from California, lawyers over former newsmen, 45-year-olds rather than 65-year-olds, and so on. It is important to note that our knowledge of subjects' preferences is incomplete.

example of our method of calculating candidate preference is given in the Appendix.

RESULTS

Candidate Choice

The first analytic task is to see how well this simple method of determining differential candidate preference predicts the actual vote choice of our subjects. This question is not as interesting in the context of a mock election campaign as it would be in a real election, but it is the question in which most studies of voting behavior are primarily interested, and it will demonstrate that our simple normative method of determining naive candidate preference is reasonable. We specified a logistic regression in which vote choice in the primary election is regressed on the difference between the candidate evaluation measures calculated for the two candidates offered to subjects for closer inspection after completing the main experiment. This single variable (which, of course, incorporates issue voting, group endorsements, and candidate appearance and personality) was highly significant ($p < .001$), correctly predicting more than 60% of the actual vote choices against a baseline of 50%—about as good as can be expected in a primary election campaign in which party cannot be used as a voting cue.

Predicting "Correct" Voting

Survey designs of actual elections are far better vehicles for learning why one particular candidate wins, of course, or for building methods to predict the direction of the vote choice. In contrast, we are primarily interested in understanding whether people voted correctly according to their fully informed interests. Since 70% of the subjects (206 out of 293) would not have changed their vote after learning everything there was to learn about the two candidates, by our definition they voted correctly. The remainder, by their own determination, voted "incorrectly."⁷ This is our first important finding, although it is impossible to say how generalizable this proportion is, given that it is based on a mock election, albeit one designed to simulate the crucial aspects of real campaigns.

The first major question we pose is how well our normative measure of naive or intuitive candidate preference predicts fully informed correct voting. If the prediction is good, then we are justified in referring to it as a "normative" measure, and we are justified in using it as a second measure of correct voting, a surrogate for the more complete, but much more difficult to obtain, fully informed correct vote determination.

Thus, we specified a second logistic regression in which the subject's fully informed determination of the

correct candidate was regressed on the normative candidate differential variable described above. This single variable is again highly significant ($p < .001$) and correctly predicts almost 66% of the correct vote choices—better than this same variable predicts the actual vote choice! This is strong validation of our normative candidate differential variable as another measure of correct voting. Just as important, *our normative method of determining candidate preferences does almost as well in determining "correct" voting decisions as did voters themselves* (who voted for the "correct" candidate 70% of the time).

This finding is the crux of our argument, and we want to put it in clear perspective. We have good, but certainly incomplete, knowledge of voters' preferences; based on this knowledge, we can make reasonable (but again far from perfect) inferences about how subjects evaluated the information they learned about the candidates. Knowing nothing about how voters actually combine these evaluations into a vote choice but by modeling a plausible alternative, *we can do almost as well in determining correct vote decisions as do voters themselves*, who have perfect knowledge of their own preferences and perfect knowledge of how favorably they responded to the candidate information to which they were exposed.

The reason our normative measure of naive candidate preference predicts a fully informed vote choice better than an actual vote choice is that voters, under normal information processing circumstances, cannot possibly achieve the care and objectivity that would be possible if they were given more time and the opportunity to become fully informed about two opposing candidates. With the aid of a powerful computer to help keep track of what has been learned about the multiple candidates, we can reasonably approximate that care and objectivity. In other words, under normal circumstances the vote decision is an intuitive, global judgment, and people with limited cognitive resources have a very hard time combining complex sets of information to make such judgments. Only when given the time and presented with the information in a very focused (only two candidates, not all six) and easily comparable format (similar information about the two candidates presented side by side) can people approach the objectivity of our simple averaging algorithm.⁸

Our results suggest that (1) voters in our experiment, confronted with the same type of time constraints and information overload faced by voters in actual elections, nevertheless do a pretty good job of selecting the candidate for whom they would have voted had there

⁷ Since fully 30% of the subjects were willing to say they would change their vote, implicitly admitting that they had voted "incorrectly," self-presentation concerns were probably not a major issue to most subjects (refer to footnote 4).

⁸ The finding that fairly simple but objective algorithms for combining multiple criteria for judgment outperform naive (or expert) decision makers who rely on a global judgment is fairly common in the decision-making literature. Perhaps the classic example is Meehl's (1954) summary of 20 different studies comparing what he called "clinical" judgment to a statistical summary of objective information available to the decision maker. In no case was the global judgment found to be superior to the statistical summary. Dawes (1988) reviews many subsequent studies, all of which reach the same conclusion.

TABLE 1. Willingness to Change Original Vote as a Function of the "Quality" of the Original Choice

	Would Not Change	Would Change	Total	(N)
Worst candidate chosen	56%	44%	100%	(82)
Intermediate candidate	62%	38%	100%	(58)
Best candidate chosen	83%	17%	100%	(148)
Tau c = .26, p < .001				

been no constraints on their information-gathering capabilities (i.e., if they have full information); but (2) they could clearly do better than they do under current circumstances, given more time or presented with information in a more easily "digestible" manner; and (3) we as researchers or external observers can determine fairly accurately who individual voters, given full information, would want to pick as their best choice.

Further Validation of the Normative Measure

We can provide additional evidence that our normative candidate differential measure is a reasonable approximation of a fully informed correct vote decision and is, in fact, what voters were trying to achieve during the correct vote determination stage of our experiment. First, we can ask if the *additional* information subjects learn about the candidates during the final correct voting determination task can predict whether they will want to change their initial choice. If we cannot understand how additional information may *change* decisions, then we probably do not have a very good grasp of how those decisions are reached in the first place.

To explore this question, we put subjects' actual vote choices into three categories according to our normative measure of candidate preference: Had subjects voted for the best possible alternative, the worst possible alternative, or (in the four-candidate condition) a candidate who fell between these two extremes? If our measure reasonably captures voters' fully informed preferences, then we should observe a much higher percentage of voters who want to change their vote among those who (according to our measure) originally voted for the worst possible candidate than among those who voted for the best possible candidate. The data, shown in Table 1, reflect just this pattern: Almost 44% of those who we thought picked the worst alternative were willing to admit a mistake after examining more information about the candidates, while less than 17% of those who we believed picked the best possible candidate wanted to change their mind.

As a final check that our normative measure of differential candidate preference is a good approximation of how voters would try to process information and decide who is their best choice if they had the opportunity, we can use the same procedures to model the *additional* information subjects gained about the can-

didates during the final task—that is, information not available to them when they made their actual vote choice. If our method is a good one, then this new variable, which incorporates all the new information presumably gained only after the campaign was over, should predict which of the two candidates was selected by subjects as their correct choice, after controlling for the actual vote choice.

To address this issue we (1) recoded the dependent variable to represent whether Candidate A or Candidate B (an arbitrary distinction) was determined by the subject to be the correct choice, (2) used as predictors whether subjects had originally voted for Candidate A or Candidate B, and (3) added a new variable reflecting an evaluation of the information learned about the two candidates after the original vote. The original vote reflects the intuitive or naive candidate choice. Whatever information was learned about the candidates during the primary election (plus whatever inferences people were willing to make) is reflected in this vote choice. We already know this will be a highly significant predictor; after being presented with more information, 70% of subjects reported they had voted correctly—a substantial improvement over a chance level of 50%.

The more interesting variable is the second in the equation, that is, the effect of the additional information gained from the fully informed correct voting task. If we as outside observers, knowing only the voters' preferences and the stands of the candidates, can predict fully informed choices better than can voters with their own intuitive methods (that is, if we do indeed have a good idea of what voters will believe is their best choice), then our information-gain variable should add significantly to the predictive power of the equation. As shown in Table 2, that is the case. Therefore, we are now even more confident that we understand how voters are determining their "correct" vote choice because we can predict how additional information about the alternatives influences that decision. Thus, we feel reasonably confident in offering our "normative" method for determining candidate preference as an alternative criterion for correct voting.

TABLE 2. Effect of New Information on Decision to Change Vote

	B	Wald	Significance
Direction of actual vote	1.76 (.27)	41.86	.000
Effect of new information	.28 (.12)	5.35	.010
Constant	-1.32 (.20)	42.12	.000

Note: The B coefficients come from a logistic regression analysis; standard errors are in parentheses. For this analysis, -2 times the log likelihood was 323.93, while the model chi-square with 2 degrees of freedom was 5.03 (p < .000). Classification was correct for 71.2% of the cases. N = 288.

AN APPLICATION TO AMERICAN PRESIDENTIAL ELECTIONS

The results from our experiment are quite impressive. Confronted by an information environment in some ways even more difficult than an actual campaign, our subjects nevertheless voted correctly 70% of the time. Some skepticism may remain, however, about generalizing to an actual election. Any experimental study of the vote choice, no matter how realistic it attempts to be, is a far cry from a real campaign. In this section we illustrate how our method can be applied to ANES data.⁹

It is certainly possible to construct a measure analogous to our normative method of determining naive or intuitive candidate preference using the information in the typical ANES survey. There is no analog to our importance weights, but in analyses not reported here, an equal-weights version of our normative measure performed almost as well as the weighted version. Likewise, we have no direct measure of the candidate information to which a voter has been exposed, but we can use willingness to answer ANES questions about the candidates (e.g., to attribute an issue stand to a candidate) as a reasonable indirect measure. Using this analog, we can determine (with a great many qualifications, of course) the percentage of "correct" voting in different American presidential elections.

We pose three simple macro-level hypotheses about differences across elections in the percentage of voters who should have voted "correctly." To the extent these hypotheses are supported, they provide important predictive validity for our surrogate measure of correct voting as it can be operationalized from survey data.

HYPOTHESIS 1. *Given limited cognitive resources, voters are more likely to make correct decision when there are fewer candidates than when there are more alternatives in the choice set. This suggests that voters were more likely to make correct choices in 1972, 1976, 1984, and 1988 than in 1980, when there was a reasonably successful third-party candidate running in the general election.*¹⁰

HYPOTHESIS 2. *If the number of candidates is held constant, then voters will be more likely to make correct decisions when the candidates are easy to distinguish than when they are difficult to distinguish. This suggests rates of correct voting were higher in 1972, 1984, and 1988, when more ideologically extreme candidates*

captured the two parties' nominations, than in 1976, when more centrist candidates opposed each other.

HYPOTHESIS 3. *All else equal, voters will be more likely to make correct decisions when campaign resources are reasonably balanced, giving all candidates an equal opportunity to get their case across, than when resources are imbalanced. Phrased more cynically, this suggests that candidates whose campaign resources are greater than those of opponents can attract more support than should have been the case if everyone had voted "correctly."*

Table 3 presents the results of a preliminary test of these hypotheses. Without going into all the details (which are left for the Appendix), we have replicated as closely as possible with ANES data an equal-weights version of our normative measure of naive candidate preference. In addition to party identification, every policy issue and candidate-group linkage that could be objectively estimated was included in the analysis. Taking this measure as our criterion of "correct" voting, we can determine the proportion of correct voters in recent American presidential elections.

Table 3 is broken into five sections, one for each presidential election between 1972 and 1988. The top row of each section displays the proportion of voters for each candidate (and the overall proportion) who, by our determination, voted correctly. These numbers range between a low of 51% for John Anderson in 1980 to a high of 89% for Walter Mondale in 1984. Overall accuracy of voting across these five elections ranges from just under 68% up to 79%, with a mean of 75%.

Consistent with hypothesis 1, the mean number of correct votes in years with two major candidates, 77.1%, is significantly higher than the mean of 67.8% ($z = 5.81, p < .001$) for 1980, when there was an important third-party candidate.¹¹ Consistent with hypothesis 2, the mean number of correct votes in the three two-candidate elections with ideologically distinct candidates, 77.7%, is greater than in the one year with two centrist candidates, 75.5% ($z = 1.83, p < .04$, one-tailed).

The data necessary for testing hypothesis 3 are less straightforward. Each section of Table 3 provides ANES data on the percentage of voters who reported voting for each of the major presidential candidates

¹¹ If we consider only Carter and Reagan voters as the most appropriate comparison group (and ignore Anderson voters), then the 1980 proportion increases to 69.4% voting correctly, but the hypothesis is still strongly supported ($z = 4.53, p < .001$).

One reviewer suggested this hypothesis could be confounded by "strategic" vote considerations (e.g., Abramson, Aldrich, Paolino, and Rohde, 1992), for example, voters realizing that Anderson was their "best" choice but that he could not win and thus voting for their "second best," who had a much better chance. Such considerations should not affect the relative rankings of the three candidates on feeling thermometer evaluations, however. We repeated this analysis, substituting the relative ranking measure for the actual vote choice. With this alternative criterion the overall proportion of correct votes (or rankings) actually falls to 66.9%. If we limit the analysis to voters who evaluated Carter or Reagan highest, the proportion correct remains at 69.4%. In neither case would it appear that strategic voting considerations provide an alternative explanation for these results.

⁹ The data used in the following analyses were made available by the Inter-university Consortium for Political and Social Research. The 1972 data were originally collected by Warren Miller, Arthur Miller, Richard Brody, Jack Dennis, David Kovenock, Merrill Shanks, and the Center for Political Studies (Miller et al. 1975). The 1976 data were originally collected by Warren Miller, Arthur Miller, and the Center for Political Studies (Miller et al. 1977). The 1980, 1984, and 1988 data were originally collected by Warren Miller and the National Election Studies (1982, 1986, 1989). Neither the collector of the original data nor the consortium bears any responsibility for the analyses or interpretations presented here.

¹⁰ The sporadic Ross Perot candidacy in 1992 appeared unlikely when the ANES staff made the final decisions about their pre-election survey. This resulted in a dearth of questions about Perot in the survey and precludes that election from this analysis.

TABLE 3. Correct Voting in American Presidential Elections, 1972–88

Presidential Candidate	% Supporters Voting Correctly	% Reported Vote	% Predicted Correct Vote	Deviation, Reported – Predicted	% General Election Spending	Deviation of Spending from Proportional Share	Overall Accuracy
1972							79%
McGovern	74.5	35.7	38.5	–2.8	33.7	–16.3	
Nixon	81.5	64.3	61.5	2.8	66.3	16.3	
1976							75.5%
Carter	72.2	51.1	47.1	4.0	52.4	2.4	
Ford	78.9	48.9	52.8	–3.9	47.6	–2.4	
1980							67.8%
Carter	64.0	40.0	34.6	5.4	40.0	6.7	
Reagan	73.7	51.6	46.5	5.1	47.7	14.3	
Anderson	50.6	8.5	19.0	–10.5	12.3	–21.0	
1984							76.8%
Mondale	89.0	41.8	55.7	–13.9	40.5	–9.5	
Reagan	68.0	58.2	44.2	14.0	59.5 ^a	9.5	
1988							77.1%
Dukakis	78.2	47.1	49.4	–2.3	53.0	3.0	
Bush	76.1	52.9	50.5	2.4	47.0	–3.0	

Note: Data are from the 1972, 1976, 1980, 1984, 1988 ANES.

^aIncludes spending during primaries because Reagan was essentially unchallenged, and his campaign manager is quoted as saying that all spending during the primaries was aimed at the general election campaign.

that year; the percentage each candidate would have received if all voters, by our calculations, had voted correctly; and the difference between the two (a positive difference indicates the candidate received more votes than he “should” have, while a negative difference means he received fewer votes than he should have); the percentage of all money spent by each candidate (or on his behalf) during the general election campaign¹²; and the difference between that percentage and a proportional share of all spending (i.e., 100% divided by the number of candidates).

The crucial data for hypothesis 3 are in the two rows of deviation or difference scores. If hypothesis 3 is correct, then candidates with a disproportionately larger share of campaign resources (i.e., a positive difference in the bottom row of each section of the table) should be able to win more votes than they correctly “should” have (i.e., a positive difference in the fourth row of each section of the table), while candidates with a disproportionately smaller share should, all else equal, receive fewer votes than they should have. The Spearman rank-order correlation between these two difference scores is .77, $p < .01$, providing strong confirmation of the third hypothesis.

DISCUSSION

Any political philosophy presumes a view, a psychological theory, of human nature (Lau 1985). Classic democratic theory sets unrealistic standards for ideal citizens at least in part because it holds unrealistic expectations about the very nature of human cognition. Beginning with a more circumscribed human psychol-

¹² These spending figures are reported in great detail in Alexander 1975, 1979, 1983; Alexander and Bauer 1991; and Alexander and Haggerty 1987.

ogy, we can set more realistic goals for democratic citizens—and still judge how readily those goals are met.

We offer a very simple standard: Irrespective of how the vote decision is reached, how frequently do voters vote correctly? To ask this question implies that one has an answer, or at least a method for obtaining the answer. Relying on notions of “fully informed” interests, we have suggested one such method. Our analyses of both experimental and survey data show that our method does a reasonably good job of measuring correct voting. Had we relied only on the experimental data and its rather artificial full-information determination of correct voting, our findings would have been of more limited value. But the second measure of correct voting, available in both experiments and surveys, should prove to be much more useful because it is much more widely applicable. Moreover, that we have found corroborating evidence from two very different research designs lends much credence to either set of findings alone.

We do not want to suggest that our method is the only way, or even the best way, of determining “correct” voting. Recently, Bartels (1996) and Delli Carpini and Keeter (1995) have presented research that, although aimed at a very different purpose, could be construed as alternative approaches to correct voting. Bartels (1996) seeks to find the political consequences of an electorate which, as a whole, is notoriously low in political knowledge. He addresses this question by estimating a logistic analysis which includes the vote decision regressed on 20 variables representing various demographic characteristics and the interaction between each of them and a measure of political knowledge. Based on significant interaction terms, Bartels concludes that the informed and uninformed portions

of the electorate do vote differently and, moreover, that the results of close elections could change if uninformed voters voted as if they had full information. Although Bartels does not use the term "correct" to describe the votes of the fully informed, it is easy to make that interpretation and to call the votes of the less well informed, to the extent they differ, incorrect. His approach is drastically different from ours in that it completely avoids the "political values and preferences" that form the basis of our determination of what is correct, presumably because if the vote decisions of the less informed could be wrong, then so, too, could the political preferences upon which those decisions are based.

Delli Carpini and Keeter (1995) have the same aim as Bartels, to demonstrate the political consequences of low political information. They begin with two plausible assumptions: (1) Material interests differ across various gender, class, race, and age divisions in society, and (2) the more politically informed will be better able to discern their interests. The latter assumption clearly allows us to label the opinions of the most politically informed as "correct" and the opinions of the less politically informed, if they differ, as "incorrect," although Delli Carpini and Keeter are not this explicit. Their finding that the well-informed hold political opinions different from those of the poorly informed shows that information matters and, in general, tends to increase group differences on political issues. Delli Carpini and Keeter go on to show that the amount of political information affects the extent of issue voting. They note that the most politically informed women had a strong vote preference for Bill Clinton in 1992, but women in the lowest quartile of information preferred George Bush. Delli Carpini and Keeter do not take the step that for our purposes would be most relevant—tracing the extent to which holding "incorrect" opinions led the least politically informed to vote "incorrectly"—although we could imagine how one could perform this analysis.

Thus, we can point to three quite different methods of determining a "correct" vote. Each has its merits and problems, and we are not suggesting that our method is superior to the others. But we do suggest that the issue of "correct voting" is a very important one, and we hope that we have illustrated at least one way in which it can be addressed empirically.

Taking our results at face value, we can return to the question with which we began: Is 75% of voters voting correctly in a typical presidential election "good enough" for a system of government to be considered truly democratic? We are pleasantly surprised by these results: This high a level of "correct" voting certainly validates the efficiency of heuristic-based information processing that underlies our view of human nature (but see Bartels 1996). Moreover, it challenges those critics who hold that democracies' problems stem primarily from people not having the motivation to gather sufficient information to figure out what is in their best interest. Most people, most of the time, can make this calculation, at least in presidential elections.

But is this level high enough for us to consider that (at least) the American version of democracy "works" at some minimal level? It is certainly too soon to draw any firm conclusions about this question in any case, and we should mention several very important caveats.

First, if 75% of voters are voting correctly in the typical presidential election, then 25% are voting incorrectly. If this group were distributed randomly, then it would not be much of a problem, but our test of hypothesis 3 demonstrates otherwise. Candidates with more money have an advantage. Here is yet one more argument for serious campaign reform of American elections.

Second, the analyses in Table 3 only consider voters, but barely half of eligible Americans vote. The interests of nonvoters are beyond the scope of this paper but certainly not beyond the scope of theoretical concern.

Third, we have only examined presidential elections, but our federal system ensures that much of what is important in politics happens at lower levels of government. At those levels, we would expect even less correct voting—or at least for elections with less media attention than presidential campaigns, which is to say all others. If 75% of voters were to vote correctly for president and 65% were to vote correctly for mayor, we would be quite happy. If 75% were to vote correctly for president and 25% were to vote correctly for mayor, we would not.

Political science as a field has only begun to map out the "correct voting" landscape. A great deal more research must be conducted before empirical political science can be of much help to normative theorists struggling with this question.

Whatever may ultimately prove to be the answer to the question of the extent to which any system of government can be considered truly democratic, no one would argue that things cannot be improved. Given a metric of "correct" voting, we can turn to the equally important issue of what leads people to make more or less optimal decisions. This can be addressed at both the individual and institutional level. For example, we could ask if particular information search tactics or different decision or choice strategies lead some individuals to make better decisions than others. Bartels's (1996) recent findings warn us not to assume that all voters use heuristics and other information shortcuts equally effectively. We could ask whether different practices the media have developed for covering campaigns encourages or discourages the more effective information-processing strategies. We could ask whether institutional arrangements that favor two-party systems or that separate the fates of executives and legislators change the probability of correct voting. And we could study whether certain campaign tactics are particularly effective in distorting the "correct" outcome of an election and, if so, consider means of discouraging those tactics. All these questions—which a "relevant" political science ought to be asking—can only be addressed with a defensible measure of correct voting.

APPENDIX

Calculating "Naive" Candidate Evaluations in the Experiments

We assume that voters consider four general types of information in forming evaluations of candidates: party affiliation, issue stands, group endorsements, and a candidate's personality/appearance. For the first three categories the voter's own values or preferences are important (i.e., their party identification, issue stands, and their evaluation of the groups doing the endorsing), along with "objective" standards of where the candidates stand on the issues and the facts about the party each candidate belongs to and which candidate a group endorses. Because we are looking only at choice during a primary election, party affiliation becomes irrelevant as a means of choosing between candidates and, for present purposes, was ignored.

Issue Stands. Every issue for which a subject (1) expressed a stand on the pre-experiment questionnaire, and (2) learned a candidate's position (generally by choosing to "access" that information during the campaign, but also by exposure to a campaign advertisement) contributed one item toward that candidate's summary evaluation. We employed Rabinowitz and MacDonald's (1989) "directional theory" of issue voting, which argues that being on the same "side" of an issue is the most important consideration in issue voting. For computation, Rabinowitz and MacDonald suggest subtracting the midpoint of the scale from the voter's own opinion and from some determination of the candidate's position before multiplying the two together. With 7-point scales, the resulting product can range in theory from -9 to $+9$, but in practice depends on the range of opinions expressed by voters and the actual stand taken by the candidate. Each item was rescaled to range between -1 and $+1$. The candidate's positions were determined objectively by a panel of experts who carefully read each position taken by a candidate and placed it on the same scale subjects were using to express their political opinions (see footnote 5).

Group Endorsements. Eleven groups endorsed a candidate in each party's primary. Subjects had previously evaluated those groups (e.g., the American Civil Liberties Union) on 100-point feeling thermometers. Voters were considered to "like" groups they evaluated above the midpoint of the scale (50) and (to control for individual differences in use of the feeling thermometer) above the mean of all groups evaluated; they were considered to "dislike" groups rated below 50 and below their personal mean of all groups evaluated. Then any time a voter learned (through accessing that item) that a group endorsed a particular candidate, the endorsement (scored $+1$ if the group was liked and -1 if the group was disliked) counted one additional item toward that candidate's evaluation.

Candidate Personality/Appearance. Four different brief personality descriptions were available about each candidate, and one picture. (The picture could be accessed as a separate item, but it also appeared for the last five seconds of every campaign commercial.) All of these items were rated for their "attractiveness" by 67 pretest subjects. The mean ratings of the 67 pretest subjects, rescaled to range between -1 and $+1$, were treated as the "objective" value of that information. Every personality description accessed by a subject contributed one additional item toward that candidate's evaluation. Because all subjects saw every candidate's picture at least once (during a campaign advertisement sponsored by the

candidate), the mean rating of the picture contributed one last item to each candidate's total.

Importance Weights. Early in the pre-election questionnaire subjects were asked to think back to the 1992 presidential election and to explain briefly, in their own words, why they had voted for (or preferred, if they had not voted in 1992) one of the candidates over the others. After this recollection, subjects were asked to look at ten reasons people sometime give for voting for one candidate over another, and to pick the one that was the most important factor they used in deciding how to vote in 1992. The ten factors listed for consideration were "Appearance/I Just Liked Him Best," "Candidates' Party Affiliations," "Groups/Persons Supporting Him," "Couldn't Stand the Other Two Candidates," "Just Time for a Change," "Most Trustworthy Candidate," "Foreign/Defense Policies," "Economic Policies," "Social Policies," "Most Competent/Qualified Person." After subjects had chosen the one most important factor, they were asked to consider the remaining ones and select those that were "very important" in their vote decision. Subjects were allowed to select any number of reasons as very important. Finally, subjects were asked to consider any remaining factors and select those that were "not at all important" in their vote decisions. Again, any number of factors could be chosen. The most important factor was given a weight of 1; "Very Important" reasons were given a weight of .67; "Not at All Important" reasons were given a weight of 0; and any remaining reasons were considered to be "Somewhat Important" and given a weight of .33.

The information available about the mock candidates was then matched up with one of these ten reasons (all but "Couldn't Stand the Other Two Candidates" and "Just Time for a Change"), and the weight assigned to that reason was used as the "importance weight" of that particular type of information.

Summary Evaluations. Each item was multiplied by its weight before the mean was calculated. Then "naive differential candidate evaluations" were computed by subtracting the mean rating of the candidate a subject did not vote for from the summary evaluation of the candidate a subject did vote for. This difference score is positive if the subject voted for the candidate we calculated they should like the most, and negative if they voted for the other candidate. The greater the difference in evaluation of these two candidates, the stronger the presumed preference.

As an example, suppose a subject voted for "Chris Rodgers" during the primary election and, at the conclusion of the experiment, was asked to consider carefully all information about Rodgers and "Pat Thomas," another candidate in the same primary. During the campaign this subject learned Rodgers's stand on affirmative action (objectively determined to be 6), abortion (7), and defense spending (5), and Thomas's stand on abortion (4) and crime (6). To make life easy, let us assume this subject took conservative stands (i.e., 6) on all of these issues herself. During the course of the campaign two group endorsements were learned: the National Taxpayers Association preferred Rodgers, as did the National Rifle Association. This subject had previously rated "groups that try to limit taxes" 85 on a feeling thermometer, well above her personal mean of all groups rated (63); but disliked the National Rifle Association (thermometer rating of only 25). Two personality descriptions were learned about Rodgers, a family friend's description (rated .75 by the 67 pretest subjects) and an anecdote from the candidate's mother (rated .20); and two personality descriptions were learned about Thomas, a staff member's description (.60) and a political opponent's description ($-.50$). Pictures of both Rodgers

(.10) and Thomas (.23) were seen during political advertisements. During the pre-election questionnaire this subject had selected "Social Policies" as the most important reason for her 1992 presidential vote choice (given a weight of 1), "Foreign/Defense Policies" and "Groups/Persons Supporting Him" as very important reasons (weights of .67), "Most Trustworthy Candidate" and "Most Competent/Qualified Person" as somewhat important (weights of .33), and "Appearance/I Just Liked Him Best" as not at all important (weight of 0).

Table A-1 summarizes this information, and the way each bit of it was presumed to be evaluated. Rabinowitz and MacDonald's "directional" calculations (shown in the table only for Rodgers) are straightforward. These values are the numerator for the "rescaling." The denominator varies across issues; it is based on the presumption that subjects will use the entire 1 through 7 range of the attitude scale (and hence the first multiplicand is -3 or +3 at its extreme. The second multiplicand is constant across subjects, however, always equaling the objective rating of the candidate's position minus 4. The difference in the potential range of values across Rodgers's stand on abortion, affirmative action, and defense spending is reflected in the denominator used for the rescaling of those three issues (9, 6, and 3, respectively). The National Taxpayers Association endorsement counts +1 for Rodgers (because that group is liked), but the NRA endorsement counts -1, as that group is disliked. Ratings of the personality descriptions and appearance of the candidates have already been rescaled in the figures listed in the table. The mean candidate evaluation is simply calculated by summing the evaluation of every individual bit of information learned about a candidate, multiplied by its importance weight, and dividing by the total number of such bits of information. We would predict that this subject, "naively" calculating differential candidate evaluations, should decide to vote for Rodgers, because in this example he was the more highly evaluated candidate.

Calculating "Correct Voting" in the ANES Surveys

Our goal in establishing criteria for determining the candidate a respondent should have voted for was to include as many of the criteria for candidate evaluation that have been

identified by prior research as possible, allowing respondents to determine their own preferences while relying primarily on "objective" criteria for rating the candidates. We included four categories of predictors: party affiliation, agreement with a candidate's policy stands, candidate-social group linkages, and (for incumbents running for reelection) performance evaluations. Emotional responses and trait ascriptions were not included because we could think of no defensible way to get an objective measure of them. The bulk of the items on which we rely were not introduced into the ANES surveys until 1972, and thus we go back only that far. We do not include 1992 in the analysis because there is a dramatic difference in the number of items available to evaluate Perot relative to Clinton and Bush. In general (although the survey analysis included a few more "types" of information), the "correct voting" determination in these surveys was calculated in exactly the same way as the "naive" candidate evaluations in the experiments, with one major exception: All weights were either 0 or 1. In all cases we rescaled the individual items to range between -1 and +1, with -1 representing the most negative evaluation and +1 representing the most positive evaluation. Summary evaluations of each candidate were determined by the simple mean of all nonmissing evaluative items; the "correct" candidate that should have been voted for is the one with the highest summary evaluation.

Party Identification. The 7 points of the standard party identification scale were recoded to equidistant values between -1 and +1, with "Strong Democrat" at -1 and "Strong Republican" at +1. The resulting item was included in this form in the summary evaluation of Republican candidates; it was multiplied by -1 before being included in the summary evaluation of Democratic candidates. As an example of how this variable worked, all weak Democrats received a +.67 toward their summary evaluation of McGovern in 1972, and a -.67 toward their summary evaluation of Nixon. Because Anderson ran in the general election as an independent in 1980, party identification was not relevant to his summary evaluation.

Agreement with Candidates' Policy Stands. All political attitudes on which respondents were asked for their own opinions and for their perceptions of the candidates' positions were considered. We established "objective" criteria for

TABLE A-1. Hypothetical Ratings of Two Candidates in Experiment

Issue stands	Hypothetical		Rabinowitz	Importance Rescaled	Thomas's	Importance	
	Voter's Stand	Rodgers's Stand	Directional Calculation				Weights
Abortion	6	7	$(6 - 4 \times (7 - 4)) = 6$	6/9 = .67	4	0	$\times 1.00$
Affirmative action	6	6	$(6 - 4) \times (6 - 4) = 4$	4/6 = .67			$\times 1.00$
Crime	6				6	4/6 = .67	$\times 1.00$
Defense spending	6	5	$(6 - 4 \times (5 - 4)) = 2$	2/3 = .67			$\times .67$
Group endorsements							
Nat'l Taxpayers Assoc.		(liked group)		1.00			$\times .67$
Nat'l Rifle Assoc.		(disliked group)		-1.00			$\times .67$
Personality descriptions							
from family friend		.75					$\times .33$
from staff member					.67		$\times .33$
from political opponent					-.50		$\times .33$
anecdote from mother		.20					$\times .33$
Candidate's appearance		.10					$\times 0$
			Weighted sum		.25		$\times 0$
			Mean candidate evaluation		$\Sigma = 2.08$		$\Sigma = .72$
					.30		.18

TABLE A-2. Correct Voting, 1972 ANES

Issue Stands	Respondent's Position	Political Experts' Mean Determination	
		Nixon's Stand	McGovern's Stand
Guaranteed jobs	V172	4.7	2.3
Tax rates	V178	4.5	2.8
Vietnam withdrawal	V184	4.8	1.6
Fighting inflation	V190	3.4	3.1
Legalizing marijuana	V196	5.5	3.3
School busing	V202	5.3	3.1
National health insurance	V208	4.8	2.4
Pollution	V214	3.3	2.5
Women's role	V232	3.5	2.7
Treatment of criminals	V621	4.3	2.8
Minority aid	V629	4.3	2.4
Liberalism-conservatism	V652	5.1	2.2
Solving urban unrest	V670	4.3	2.5
Student demonstrations	V678	5.1	2.8
Job performance evaluation	V46	3.9	
Groups with close links to Nixon: Businessmen, southerners, conservatives			
Groups with close links to McGovern: Liberals, poor people, blacks			
Party identification	V140		

Note: Data come from the 1972 ANES (Miller et al., 1975, Study #7010).

where each candidate stood on those issues by first creating a scale of political knowledge (from responses to all questions that clearly had a correct answer), and then taking the mean perceptions of where the candidates stand on the issues of all respondents above the median on the resulting political knowledge scale. Tables A-1 and A-2 list the particular items that were used to gather the respondent's opinion of these issues, and the resulting objective placements of each candidate on those same issues. As in the experiments, agreement scales were computed via the "directional" method suggested by Rabinowitz and MacDonald (1989). Thus, each political attitude on which a respondent expressed an opinion contributed one additional item to the evaluation of each candidate.

Candidate-Social Group Linkages. Beginning in 1972 the ANES surveys ask respondents which of a list of social groups they feel "close to." We then calculated a simple crosstab between each "closeness" item and the reported vote. Whenever this crosstab resulted in a significant relationship between feeling close to a group and preference for one of the candidates, a candidate-group linkage item was created, scored 0 for respondents who did not feel close to the group and 1 for respondents who did feel close to the group, and added to the summary evaluation of that candidate. As an example, in 1972 respondents who felt close to liberals, poor people, and blacks all showed a significant vote preference for McGovern; while respondents who felt close to businessmen, southerners, and conservatives showed a significant vote preference for Nixon.

Incumbent's Job Performance. In every election year where a sitting president was running for reelection, survey respon-

dents were asked for their perceptions of the incumbent's job performance, both an overall perception and more specific (e.g., economic policy; foreign policy) perceptions. We recoded each such performance evaluation to run from +1 (Strongly Approve) to -1 (Strongly Disapprove), and added them to the summary evaluation of the incumbent candidate.

Table A-2 lists the specific items that were used in the correct voting determination for the 1972 ANES, along with the objective criteria that were used in this determination. Similar tables for the 1976-88 ANES are available from the authors upon request.

REFERENCES

- Abramson, Paul R., John H. Aldrich, Phil Paolino, and David Rohde. 1992. "'Sophisticated' Voting in the 1988 Presidential Primaries." *American Political Science Review* 86(March):55-69.
- Alexander, Herbert E. 1975. *Financing the 1972 Election*. Lexington, MA: Lexington Books.
- Alexander, Herbert E. 1979. *Financing the 1976 Election*. Washington, DC: CQ Press.
- Alexander, Herbert E. 1983. *Financing the 1980 Election*. Lexington, MA: Lexington Books.
- Alexander, Herbert E., and Monica Bauer. 1991. *Financing the 1988 Election*. Boulder, CO: Westview Press.
- Alexander, Herbert E., and Brian A. Haggerty. 1987. *Financing the 1984 Election*. Lexington, MA: Lexington Books.
- Bartels, Larry M. 1990. "Public Opinion and Political Interests." Presented at the annual meeting of the Midwest Political Science Association, Chicago.
- Bartels, Larry M. 1996. "Uniformed Votes: Information Effects in Presidential Elections." *American Journal of Political Science* 40(February):194-230.
- Bennett, W. Lance. 1988. *News: The Politics of Illusion*. 2d ed. New York: Longman.
- Bennett, W. Lance. 1992. *The Governing Crisis: Media, Money and Marketing in American Elections*. New York: St. Martin's Press.
- Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Burnham, Walter Dean. 1965. "The Changing Shape of the American Political Universe." *American Political Science Review* 59(March):7-28.
- Burnham, Walter Dean. 1974. "Theory and Voting Research: Some Reflections on Converse's 'Change in the American Electorate.'" *American Political Science Review* 68(December): 1002-23.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1960. *The American Voter*. New York: Wiley.
- Connolly, William E. 1972. "On 'Interests' in Politics." *Politics and Society* 2(Summer):459-77.
- Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." In *Ideology and Discontent*, ed. David E. Apter. London: Collier-Macmillan.
- Dahl, Robert A. 1961. *Who Governs? Democracy and Power in an American City*. New Haven: Yale University Press.
- Dahl, Robert A. 1989. *Democracy and Its Critics*. New Haven: Yale University Press.
- Dawes, Robyn M. 1988. *Rational Choice in an Uncertain World*. New York: Harcourt Brace Jovanovich.
- Delli Carpini, Michael X., and Scott Keeter. 1995. "Political Knowledge and Enlightened Group Interests." Presented at the annual meeting of the American Political Science Association, Chicago.
- Delli Carpini, Michael X., and Scott Keeter. 1996. *What Americans Know about Politics and Why It Matters*. New Haven: Yale University Press.
- Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper & Row.
- Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven: Yale University Press.

- Fishman, Mark. 1980. *Manufacturing the News*. Austin: University of Texas Press.
- Fiske, Susan T., and Shelly E. Taylor. 1991. *Social Cognition*. 2d ed. New York: McGraw-Hill.
- Gans, Herbert. 1979. *Deciding What's News*. New York: Vintage Books.
- Huntington, Samuel P. 1968. *Political Order in Changing Societies*. New Haven: Yale University Press.
- Kahneman, Daniel, Paul Slovic, and Amos Tversky, eds. 1982. *Judgement under Uncertainty: Heuristics and Biases*. New York: Cambridge.
- Kelley, Stanley, Jr., and Thad W. Mirer. 1974. "The Simple Act of Voting." *American Political Science Review* 68(June):572-91.
- Lau, Richard R. 1985. "Changing Perspectives on Political Thinking." Invited address to the 81st annual meeting of the American Political Science Association, New Orleans.
- Lau, Richard R. 1986. "Political Schemata, Candidate Evaluations, and Voting Behavior." In *Political Cognition: The 19th Annual Carnegie Symposium on Cognition*, ed. Richard R. Lau and David O. Sears. Hillsdale, NJ: Lawrence Erlbaum.
- Lau, Richard R. 1989. "Construct Accessibility and Electoral Choice." *Political Behavior* 11(March):5-32.
- Lau, Richard R. 1992. "Searchable Information during an Election Campaign." Rutgers University. Typescript.
- Lau, Richard R. 1995. "Information Search during an Election Campaign: Introducing a Process Tracing Methodology for Political Scientists." In *Political Information Processing*, ed. Milton Lodge and Kathleen McGraw. Ann Arbor: University of Michigan Press.
- Lau, Richard R., and David P. Redlawsk. 1992. "How Voters Decide: A Process Tracing Study of Decision Making during Political Campaigns." Presented at the 88th annual meeting of the American Political Science Association, Chicago.
- Lau, Richard R., and David P. Redlawsk. N.d. "An Experimental Study of Information Search, Memory, and Decision Making during a Political Campaign." In *Political Psychology and Public Opinion*, ed. James Kuklinski. Cambridge: Cambridge University Press. Forthcoming.
- Lau, Richard R., and David O. Sears, eds. 1986. *Political Cognition: The 19th Annual Carnegie Symposium on Cognition*. Hillsdale, NJ: Lawrence Erlbaum.
- Lindblom, Charles E. 1965. *The Intelligence of Democracy*. New York: Free Press.
- Lippmann, Walter. 1955. *Essays in the Public Philosophy*. Boston: Little, Brown.
- Lodge, Milton, Kathleen McGraw, and Patrick Stroh. 1989. "An Impression-Driven Model of Candidate Evaluation." *American Political Science Review* 83(June):399-419.
- Lodge, Milton, and Marco R. Steenbergen, with Shawn Brau. 1995. "The Responsive Voter: Campaign Information and the Dynamics of Candidate Evaluation." *American Political Science Review* 89(June):309-27.
- Mansbridge, Jane J. 1983. *Beyond Adversary Democracy*. Chicago: University of Chicago Press.
- Markus, Gregory B., and Philip E. Converse. 1979. "A Dynamic Simultaneous Equation Model of Electoral Choice." *American Political Science Review* 73(December):1055-70.
- McClosky, Herbert. 1964. "Consensus and Ideology in American Politics." *American Political Science Review* 58(June):361-82.
- Meehl, Paul E. 1954. *Clinical versus Statistical Predictions: A Theoretical Analysis and Review of the Literature*. Minneapolis: University of Minnesota Press.
- Mills, Charles W. 1971. *The Power Elite*. London: Oxford University Press.
- Meuller, John. 1992. "Democracy and Ralph's Pretty Good Grocery: Elections, Equality, and the Minimal Human Being." *American Journal of Political Science* 36(November):983-1003.
- Miller, Warren E., Arthur H. Miller, Richard Brody, Jack Dennis, David Kovenock, Merrill Shanks, and the Center for Political Studies. 1975. *The CPS 1972 American National Election Study: Pre- and Post-Election Waves, Volumes I and II*. [computer file] (Study #7010). Conducted by the Center for Political Studies of the Institute for Social Research, the University of Michigan. 1st ICPSR ed. Ann Arbor: University of Michigan, Center for Political Studies/Inter-university Consortium for Political and Social Research [producers], 1977. Ann Arbor: Inter-university Consortium for Political and Social Research [distributor], 1977.
- Miller, Warren E., Arthur H. Miller, and the Center for Political Studies. 1977. *The CPS 1976 American National Election Study: Pre- and Post-Election Waves, Volumes I and II*. [computer file] (Study #7381). Conducted by the Center for Political Studies of the Institute for Social Research, the University of Michigan. 1st ICPSR ed. Ann Arbor: University of Michigan, Center for Political Studies/Inter-university Consortium for Political and Social Research [producers], 1977. Ann Arbor: Inter-university Consortium for Political and Social Research [distributor], 1977.
- Miller, Warren E., and the National Election Studies. 1982. *American National Election Study, 1980. Volume I: Pre- and Post-Election Surveys*. [computer file] (Study #7763). Conducted by the Center for Political Studies of the Institute for Social Research, the University of Michigan. 2nd ICPSR ed. Ann Arbor: University of Michigan, Center for Political Studies/Inter-university Consortium for Political and Social Research [producers], 1982. Ann Arbor: Inter-university Consortium for Political and Social Research [distributor], 1982.
- Miller, Warren E., and the National Election Studies. 1986. *American National Election Study, 1984. Pre- and Post-Election Survey File*. [computer file] (Study #8298). Conducted by the Center for Political Studies of the Institute for Social Research, the University of Michigan. 2nd ICPSR ed. Ann Arbor: University of Michigan, Center for Political Studies/Inter-university Consortium for Political and Social Research [producers], 1986. Ann Arbor: Inter-university Consortium for Political and Social Research [distributor], 1986.
- Miller, Warren E., and the National Election Studies. 1989. *American National Election Study, 1988. Pre- and Post-Election Survey*. [computer file] (Study #9196). Conducted by the Center for Political Studies of the Institute for Social Research, the University of Michigan. 2nd ICPSR ed. Ann Arbor: University of Michigan, Center for Political Studies/Inter-university Consortium for Political and Social Research [producers], 1989. Ann Arbor: Inter-university Consortium for Political and Social Research [distributor], 1989.
- Nisbett, Richard E., and Lee Ross. 1980. *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Page, Benjamin I., and Robert Y. Shapiro. 1992. *The Rational Public: Fifty Years of Trends in Americans' Policy Preferences*. Chicago: University of Chicago Press.
- Prothro, James W., and Charles W. Grigg. 1960. "Fundamental Principles of Democracy: Bases of Agreement and Disagreement." *Journal of Politics* 22(May):276-94.
- Rabinowitz, George, and Stuart Elaine MacDonald. 1989. "A Directional Theory of Issue Voting." *American Political Science Review* 83(March):93-121.
- Schattschneider, Elmer E. 1960. *The Semi-Sovereign People: A Realist's View of Democracy in America*. New York: Holt.
- Simon, Herbert A. 1979. "Information Processing Models of Cognition." *Annual Review of Psychology* 30:363-396.
- Tversky, Amos, and Daniel Kahneman. 1974. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185(September):1124-31.