

Correct Voting in the 2008 U.S. Presidential Nominating Elections

Richard R. Lau

Published online: 7 March 2012
© Springer Science+Business Media, LLC 2012

Abstract Criticisms of the system by which the American political parties select their candidates focus on issues of representativeness—how choices are dominated by relatively small numbers of ideologically extreme primary voters, or how residents of small states voting early in the process have disproportionate influence. This paper adds a different concern, albeit one that still addresses representativeness. How well do primary and caucus voters represent their *own* values and interests with their vote choices? Lau and Redlawsk’s notion of “correct voting” is applied to the 2008 U.S. nominating contests. Four reasons to expect *levels* of correct voting to be lower in caucus and primary elections than in general election campaigns are discussed. Results suggest that voters in U.S. nominating contests do much worse than voters in general election campaigns, often barely doing better than chance in selecting the candidate who best represents their own values and priorities. Discussion focuses on institutional reforms that should improve citizens’ ability to make correct voting choices in caucuses and primaries.

Keywords Correct voting · Primary elections · Voting behavior · Institutional effects · Political cognition · Cognitive limitations

Historically political parties in the U.S. have determined how they select candidates to run for elective office. Through 1968 the candidate selection process was always dominated by party elites. That year, however, the growing protests against the

An earlier version of this paper was presented at the 69th annual meeting of the Midwest Political Science Association, Chicago, Illinois, March 31–April 3, 2011.

R. R. Lau (✉)
Department of Political Science, Rutgers University, New Brunswick, NJ, USA
e-mail: ricklau@rci.rutgers.edu

Vietnam War, the insurgent candidacies of Eugene McCarthy and later Robert Kennedy in the Democratic party, the withdrawal of incumbent President Lyndon Johnson from the race, the assassination of Robert Kennedy, and the resulting controversies surrounding the Democrat's ultimate selection of Vice President Hubert Humphrey as their party's nominee when he had not competed in a single primary, led the 1968 Democratic convention to appoint a reform commission to study the nomination process. The Democratic party accepted the reforms proposed by the McGovern–Fraser Commission in time for the 1972 election, and those reforms quickly led many state legislatures to establish statewide presidential primary elections (or their equivalent in party caucuses) for *both* parties, which essentially took the nomination process out of the hands of the party elites and gave it to mass elections (see Crotty and Jackson 1985; Kamarck 2009; Mayer 2000; Polsby and Wildavsky 1971; Polsby et al. 2008; Shafer 1983; but see Cohen et al. 2008, for a dissenting view).

This part of the story is well known, and the consequences of the growth of primaries has been the focus of most political science interest. One line of research has asked whether the electorate in primaries is representative of party members more generally. Turnout is usually so low in primaries (much less caucuses) compared to a presidential election that in theory it is much easier for a relatively small group of committed and more ideologically extreme activists to dominate the process, resulting in the nomination of more ideologically extreme, less representative of the party mainstream, and often less electable, candidates (although obviously this need not happen every election: compare Geer 1988; Kaufman et al. 2003; Norrander 1986, 1989; Wright 2009). A second line of inquiry, somewhat at odds with the first, explores how success in early primaries and caucuses is defined by the media and contributes to the subsequent “success” of those early victors (Aldrich 1980; Bartels 1988; Geer 1989; Kenney and Rice 1994; Norrander 2006), as momentum and bandwagon processes influence elite donors and the mass electorate alike. This line of argument views primary voters as driven by much less meaty concerns than ideology, but in the end also raises representativeness issues by asking why voters in small states such as Iowa and New Hampshire should have such disproportionate influence over candidate selection (Redlawsk et al. 2011).

This paper raises a different issue, albeit one that also reflects on representation. How often do voters in nominating (primary and caucus) elections “get it right”—by which I mean, vote for the candidate who best represents their interests, whatever they might be? Lau and Redlawsk (1997, 2006; Lau et al. 2008) call this *voting correctly*—choosing, inevitably under conditions of incomplete information, the candidate they would have chosen had they been able to gather and process complete information about all candidates in an election. If we adopt the widely accepted view of humans as “cognitively limited information processors” (Fiske and Taylor 1991; Lau 2003; Nisbett and Ross 1980), once information processing demands become great (as they do in most high-level political campaigns), it becomes almost inevitable that many people will adopt decision strategies that are far from optimal. Indeed, once we accept the idea that a choice can be mistaken, the vast literature on public opinion and political behavior would almost immediately lead us to worry that American voters—notoriously uninformed about political

matters, holding “nonattitudes” toward most political issues, with unconstrained ideologies organizing opinions toward the few issues they actually care about—would rarely get it right (Bartels 1996; Converse 1964, 1975; Delli Carpini and Keeter 1996; Kinder 1998). But in fact, what we know about correct voting in prior U.S. election campaigns suggests that voters do a surprisingly good job. Figure 1 updates findings originally reported in Lau and Redlawsk (1997; Lau et al. 2008), adding data from the 2008 ANES.¹ These data suggest that, in the ten U.S. presidential elections from 1972 through 2008, a little over 76% of voters reported choosing the candidate who, as objectively as we could determine, best represented their own values and concerns. The trend has been up, with over 85% voting correctly in the past two presidential elections.

But these data all come from the general election campaign when the Democrat’s nominee faces off against the Republican’s nominee (and the occasional moderately popular third party candidate). To date no one has explored the process by which actual politicians are nominated by their parties to become the candidates in those general election campaigns. There are at least four very good reasons to fear that voters would have a much harder time choosing correctly in primary elections. To begin with, many of the candidates are relatively new and unfamiliar to the American public. The first goal of any election campaign is to introduce the candidate to the voters, and primary elections and caucuses occur relatively early in that process. Knowledge matters to correct voting because it allows voters to better (and accurately) align their own values and preferences with the alternative candidates, on more dimensions of judgment, thus creating greater differentiation among them.² All voters ought to have much more accurate perceptions of the candidates during a general election campaign than they do earlier in the process when the parties are still selecting their candidates.

Second, there are often more candidates on the ballot in primary elections (at least early in the process) than there are in the typical 2-party dominated general election campaign in the U.S. Joe Biden, Hillary Clinton, Chris Dodd, John Edwards, Mike Gravel, Dennis Kucinich, Barack Obama, and Bill Richardson all campaigned for the Democratic nomination in Iowa, while on the Republican side Rudy Giuliani, Mike Huckabee, Duncan Hunter, John McCain, Ron Paul, Mitt Romney, Tom Tancredo, and Fred Thompson all won votes (if not delegates) in Iowa. Even if we limit ourselves to the “serious” candidates who raised significant

¹ This 2008 figure is not the “final” estimate, as the ANES has still not released any data with the open-ended questions coded. This includes answers to several of the political knowledge questions, which are used to identify experts whose mean responses provide “objective” estimates of where the candidates actually stand on the issues, how well the different traits describe them, and how good of a job President Bush did as president. The 2008 estimates will probably change slightly once these data have been coded and released.

² Lau et al. (2008) reported that, all else equal, the most knowledgeable voters were 20% more likely to vote correctly, compared to voters with very low political knowledge. Similarly, Lau and Redlawsk (2006) report from experimental data that a large number (15) of *accurate* memories about mock presidential election candidates increases the probability of a correct vote from .55 to .82. In practice knowledge will typically also be correlated with stronger and less likely to change preferences, which should also lead to higher levels of correct voting, although these influences are conceptually distinct from knowledge per se.

campaign war chests, in 2008 the Democrats were choosing among Clinton, Edwards, and Obama, while the Republicans had to choose from Giuliani, Huckabee, McCain, Romney, and Thompson. The more alternatives there are from which to choose, the more difficult the choice; and the more difficult the choice, the more likely people are to choose poorly.

Lau and Redlawsk (2006) report clear support for this hypothesis from their experimental data, where they randomly manipulated the number of candidates running in mock presidential primary campaigns. When there were two candidates competing for their party's nomination, subjects choose correctly 69% of the time, but when there were four candidates on the ballot, they only chose correctly only 31% of the time. A quick glance back at Fig. 1 reveals that the three elections with the lowest levels of correct voting over the past 40 years in the U.S. were the three elections with "serious" third party candidates (1980, 1992, and 1996). A recent study of voting decisions in 33 democracies around the world found correct voting declined from over 79%, when only two parties were on the ballot, to under 35% when nine parties were on the ballot (Lau et al. 2012). The number of alternatives matters, and it matters a lot.

Third, when people are overwhelmed by information, they rely on *heuristics*, cognitive shortcuts, to help them make decisions (Fiske and Taylor 1991; Nisbett and Ross 1980). In politics, by far the most important heuristic is party affiliation (Kam 2005; Lau 2003; Lau and Redlawsk 2001; Lodge and Hamill 1986). Strength of party identification was the most important factor in predicting correct voting in U.S. general elections (Lau et al. 2008). But primary elections and caucuses take party affiliation off the table. All candidates in the choice set are members of the




Fig. 1 Estimated levels of correct voting in recent U.S. presidential elections. *Note* The 1972–2004 figures were reported by Lau et al. (2008). All data come from the American National Election Studies. Yearly estimates are based on the mean of four closely related “normative naive” measures of correct voting

same political party. The presence or strength of party identification provides no basis for choosing among candidates, and thus cannot improve levels of correct voting.

A final reason to expect primary and caucus voters to have a more difficult task choosing correctly than voters in general elections follows from the third. When the choice set is restricted to candidates from a single party, they tend to be more similar to each other on policy issues than a broader, multi-party set of alternatives. As a result, primary voters tend to base their decisions more on “candidate” factors than policy issues (Gopoian 1982; Marshall 1984; Norrander 1986; Paulson 2009; Williams et al. 1976), but even in this domain voters in presidential primaries have a harder job than voters for lower offices, in that the serious contenders tend to be more strongly “vetted” (at least by the time the voting starts—insert your favorite joke about the current Republican field here) and thus of more uniformly high quality than candidates for lower offices. Candidate distinctiveness is another factor that past research has found to be associated with correct voting (Lau and Redlawsk 1997, 2006). When the alternatives are less distinct the choice is more difficult, and the probability of choosing correctly should decline.

All of these factors would predict lower levels of correct voting in nominating contests than has previously been reported in general election campaigns, particularly early in the process when voters are less familiar with the candidates and there are more of them to choose from. There is one factor that works in the opposite direction, however: political interest or motivation. Generally speaking, turnout in primaries is very low (averaging about 25%; see Dwyer 2011, and Patterson 2009, for recent reviews), and turnout in caucuses much lower still. The only people who bother to vote in primaries are the most interested and committed partisans—exactly the type of people who are generally the most knowledgeable about politics. Political interest, knowledge, and motivation should mitigate against the pernicious effects of having to choose among many relatively unfamiliar candidates sharing a party affiliation. Thus it is possible that despite the more difficult choice situation, the generally more sophisticated decision makers who are typically making the choices in nominating contests collectively “get it right” a high proportion of the time.

Data and Methods

One of the reasons we know relatively little about voting in primary elections is that we lack good survey data systematically exploring those elections. But in 2008 the Cooperative Campaign Analysis Project (CCAP), a large 6-wave internet-based survey, collected its first round of data in December, 2007, before any of the nomination contests in 2008, and subsequent early panel waves occur between clusters of nominating contests. Ultimately 5,707 respondents reported voting in a Democratic primary or caucus, and 3,114 respondent helped the Republicans select their nominee.³ But this is a very different project than the American National

³ See Jackman and Vavreck 2011, for a more complete description of this sample and the sampling procedures.

Election Studies, whose data and procedures provide all of the previous estimates of correct voting in U.S. elections. Before we do anything else we must convince ourselves that we can do a reasonably good job of estimating correct voting in the 2008 *general* election with the data available from the CCAP.

Operationalizing Correct Voting with the 2008 CCAP Data

Lau and Redlawsk (1997) developed a survey-based *normative naive* method of estimating correct voting in actual elections, normative in that it relies on expert judgments of candidate attributes and policy stands, and the rule that the same criteria of judgment must be applied to all candidates in the choice set; but naive in that it is based on respondent's own policy preferences, values, and priorities. Briefly, the procedure involves calculating candidate "utility scores" for how much each respondent should like each candidate, based on their own reported policy preferences and determination of which attributes of judgment they care about, and expert judgments about where the different candidates actually stand on those same policies and other attributes of judgment, rather than the highly subjective and inevitably biased perceptions of individual respondents.

Although past explorations of correct voting in U.S. presidential elections have relied on ANES data (Lau et al. 2008; Lau and Redlawsk 1997, 2006), there is no reason that other good surveys could not also be used to estimate the same concept. But just how close will those different estimates be, when the sampling frame, interview mode, question wording, and in many cases the topic of the questions themselves, differ as much as they do between the ANES and the CCAP? We have little prior experience addressing that question, and we simply do not know how robust the measure is against different operationalizations of the concept. Surely estimated *levels* of correct voting would vary with sampling error, just as any other survey-based estimate must. I would expect the correlates or predictors of correct voting to be much more stable across samples and different operationalizations, but this is an empirical question we are only beginning to address (see Lau et al. 2012; Patel 2010).

I followed the procedures of Lau et al. (2008) as closely as possible to estimate correct voting with the CCAP data. The details are reserved for the methodological appendix. Six sets of criteria go into calculating candidate utility scores in a general election: party identification (1 item); retrospective evaluations of incumbent (i.e., President Bush's) job performance (1 item) and the economy (1 item), relevant just to the incumbent party candidate John McCain; candidate-group connections (1 item per candidate), beliefs about the most important problem facing the country (1 item); trait ascriptions to the candidates (6 items); and agreement with the candidates on policy preferences (11 items) and ideology (2 items—both the presidential and vice presidential candidates). In each case, the perceptions of expert judges are used to estimate as objectively as possible where the different candidates actually "stand" on these different criteria of judgment. For most of these 24 items it is also possible to compute implicit "importance" weights tracking how much respondents cared about each of these different criteria of judgment. I then calculate

four slightly different estimates of how much “utility” each respondent should derive from each candidate being elected, by either (1) adding or (2) averaging all of the raw items described above, or by multiplying each item by its importance weight before calculating the sum (3) or mean (4) utility. It is easy to imagine situations where these mathematically different combination rules could come up with very different rankings of the candidates in an election, but in practice they almost never do. Whichever of these four slightly different combination rules are employed, the “correct” candidate is the one with the highest utility score. If a respondent voted for that candidate, they voted correctly. If they voted for a different candidate (about whom we have data), they voted incorrectly. Nonvoters, and people who voted for minor party candidates about whom the survey researchers did not ask any questions, are counted as missing.

Using these four different combination rules, I estimate anywhere between 85.3 and 86.8% (mean 86.1%) of McCain and Obama voters in the 2008 U.S. presidential election voted correctly. Compared to the right-most column of Fig. 1, this is just a tenth of a point higher than one gets from the 2008 ANES survey. Ha and Lau (2010) provide a more detailed analysis of correct voting in the 2008 general election from this CCAP data, and find that most of the predictors of correct voting in the ANES surveys similarly predict correct voting in the CCAP. These comparisons give us a great deal of confidence that, whether we are using ANES or CCAP data to estimate correct voting, we are talking about pretty much the same concept.⁴

As near as possible I follow the same general procedures to operationalize correct voting in the various Democratic and Republican primary and caucus elections in 2008, but the details become much more complicated (and tedious) to keep track of because of the changing pool of candidates competing for the nomination. The earliest contests in 2008 were the Iowa caucuses on January 3rd; the last was the South Dakota primary on June 3rd, with all the remaining contests spread out over the 5 months between these two endpoints. The two parties usually, but not always, hold their contests on the same day. From the outset, the principle investigators of the CCAP survey chose to restrict their attention to Clinton, Edwards, and Obama on the Democratic side, and Giuliani, Huckabee, McCain, Romney, and Thompson on the Republican side, and thus anyone who voted for any of the remaining candidates is, by necessity, treated as missing (only 3% of Democratic voters, but 10.5% of Republicans). But then many of these “major” candidates dropped out of the race at different stages throughout.⁵ Thus the choice facing voters is an ever-changing mix, depending on the date, state, and party of the primary or caucus in question.

⁴ The syntax commands used to operationalize correct voting in the 2008 CCAP data, along with replication data for the analyses reported below, are available from the author web page <http://fas-polisci.rutgers.edu/lau>.

⁵ Thompson withdrew after the South Carolina Republican primary on January 19th. Edwards and Giuliani withdrew after the Florida primary on January 29th. Romney withdrew after the “Super Tuesday” contests on February 5th. Huckabee withdrew and McCain claimed the Republican nomination after the four primaries on March 4th. In contrast, Clinton continued fighting through the last primaries on June 3rd.

The CCAP panel gathered data in late December, 2007, late January 2008, early March 2008, and then again in September after the two parties officially selected their nominees. I restrict the measurement of voter's values and beliefs that are used to determine their "correct" choice to surveys administered before a state's primary election or caucus. Thus for the earliest contests I can only use data from the initial baseline survey, while for later contests I can refine the measurement to include data available from the January and sometimes March waves of the survey. The vote choice itself is always measured from the survey wave occurring as soon as possible after the election—whenever it was—so from the January, March, or September waves of the CCAP panel. The Republican contest ended much sooner than the Democrats, so I ignore votes in all of the Republican primaries and caucuses occurring after March 3, when McCain clinched his party's nomination.

There are fewer items available for estimating the candidate "utility" scores in the primaries than there are for the general election contest between Barack Obama and John McCain, but still a goodly number: nine policy views, one reading of candidate ideology, three trait ascriptions, and one candidate-group connection for several of the candidates. It does no good to use strength of party identification or retrospective evaluations of President Bush's job performance even though they are available to us, as within-party these considerations apply equally to all candidates on the ballot.

In addition, it is quite rational for primary and caucus voters to factor the relative "electability" of the competing candidates—that is, their probability of winning the November election—into their vote calculus. Indeed, if primary voters in the extreme case figure they would prefer every one of their own party's candidates to any one of the opposition's candidates, electability would be the *only* criterion rational voters should consider. Although the CCAP survey asked electability questions in their December, January, and March waves, one can get more dynamic (and probably more accurate) estimates by averaging the findings from all available polls (as reported by Real Clear Politics.com) asking "horse-race" questions about potential match-ups in the general election throughout the primary season, the procedure followed here. Details are reserved for the "[Appendix](#)" section, but briefly I broke the primary season into ten periods, and estimated the relative electability of every remaining candidate in each party's nominating contest (potentially running against all remaining candidates from the opposite party) from all polls asking that question over the preceding 2–4 weeks. I then followed past practice and counted every available "consideration" equally in the determination of unweighted candidate utility scores, or employed available implicit importance weights to estimate weighted versions of those utility scores.

Results

Correct Voting in the 2008 Primary Elections

So, how well do American voters do in selecting the candidates who best represent their own values and priorities? The quick answer is, "Not very." The top panel of

Fig. 2 shows the data for Democrats. In January, when Clinton, Edwards, and Obama were all competing for the nomination, barely 30% of Democratic voters managed to select the candidate who, as far as I could tell, best represented their own interests. They would have done better if they had closed their eyes and picked randomly. Once Edwards dropped out and the two remaining candidates became increasingly familiar to voters, Democrats did noticeably better, getting it right 56% of the time, at least now better than chance—but not by much.

As shown in the bottom half of Fig. 2, Republican voters did a little better. Very early on when Guiliani, Huckabee, McCain, Romney, and Thompson were all in the mix, 31% of Republicans chose correctly—about the same as Democrats, but with more alternatives in the mix. But after Thompson drops out they jump up to over

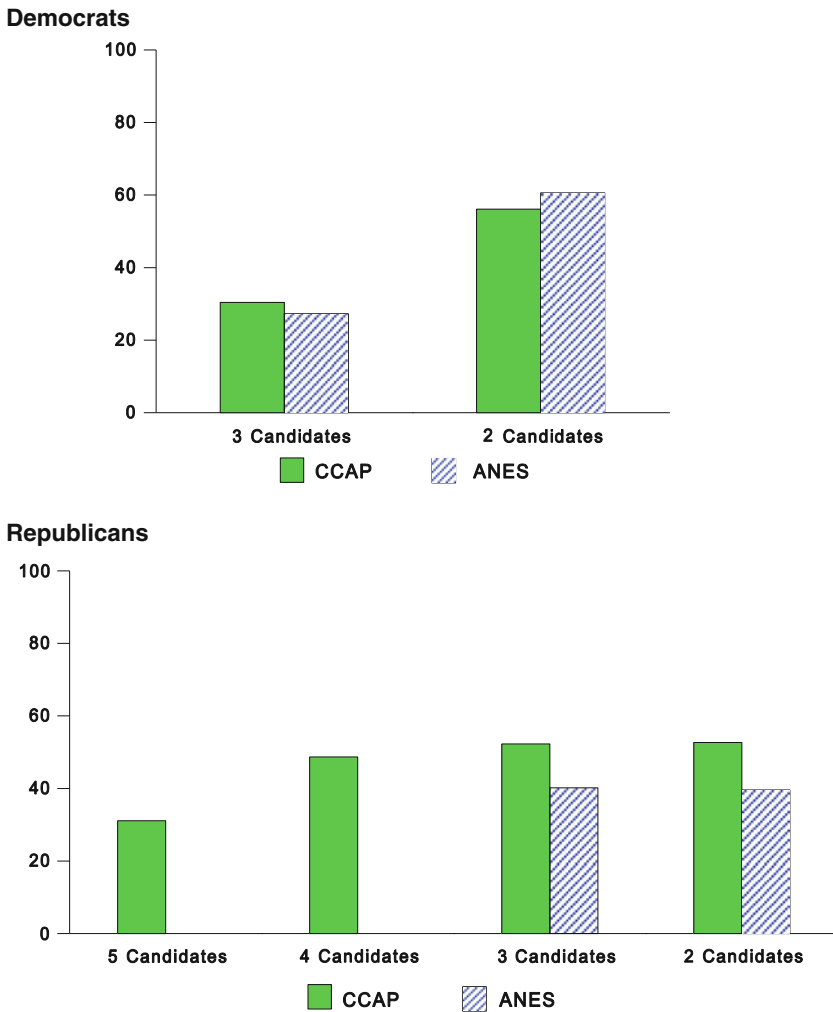


Fig. 2 Correct voting in the 2008 primaries and caucuses

48% correct (much better than chance), but never do much better than that as the field narrows to just Huckabee and McCain. I calculate that overall only 50% of Republican primary and caucus voters voted correctly. Compared to figures of over 85% in recent general elections, these numbers strike me as shockingly low—so low I was not sure whether to believe them.

A Partial Replication

In 2008 the ANES conducted their standard pre-post panel survey, contacting respondents initially in September after the two party conventions have officially nominated their candidates—and well after the primaries and caucuses were over. But fortunately they did ask respondents if they had voted in those nominating contests, and if so, who they supported. If we accept these answers as accurate, with a little simple statistical footwork to adjust the “objective” estimates of where the primary candidates stand on the issues to the question format employed by the ANES, we can also estimate correct voting in the primary elections with the ANES sample. These data are also shown in Fig. 2. The story for Democrats looks almost identical in the ANES sample as in the CCAP, with slightly higher figures for the drawn-out slugfest between Clinton and Obama. Republicans in the ANES sample look about 10–12% *worse* than they do in the CCAP. Thus there is nothing in the ANES figures to suggest that the CCAP sample (or my operationalization of correct voting using it) is somehow getting the story wrong.

Counting “Electability” More

Although each of the considerations that go into computing the candidate utility scores have been normalized to have a similar range, in practice policy considerations for most voters have a much greater influence over the final candidate utility score than electability for the simple reason that we have ten different items addressing policy (including liberal-conservative self-placement), but only a single reading of a candidate’s electability. This arbitrarily sets general limits on the relative importance of these two types of considerations that, in the case of a primary election, may not be warranted. Strategic voters might downplay the relatively minor differences among their own party’s contenders and focus on the potential electability of the different candidates in the upcoming general election. Would the results look dramatically different if we imposed an alternative weighting scheme, counting electability more heavily in the determination of the correct choice for each voter?

Figure 3 provides the answer to that question, where the results observed when electability is counted from 0–10 times in calculating candidate utility scores are reported. Among Democrats (top half of the figure), electability slightly decreases estimates of correct voting based on the summative measures, and slightly increases estimates of correct voting based on means, but the net effect is at most $\pm 1\%$. By the time electability reaches a weight of 5 (roughly half the potential weight of issues), it comes to completely dominate estimates of correct voting. Among Republicans, on the other hand, the more electability is counted, the lower the estimates of correct

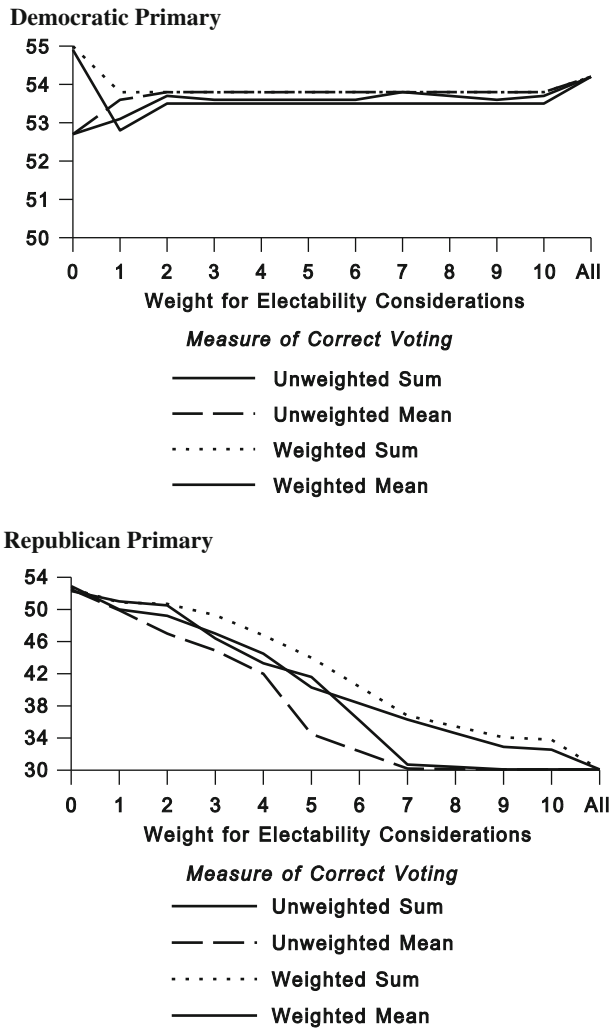


Fig. 3 Effect of electability considerations on estimated levels of correct voting

voting, and the total effect is over 20 points. Either Republicans are much worse than Democrats, or they just do not care about electability in their actual vote decisions. This would fit the stereotype of the ideologically driven, purity seeking, anti-Republican in name only (anti-RINO) extremist Republican primary voter, except that past research does not really support such a stereotype. It should be noted that John McCain was *always* perceived as the most electable Republican candidate, and after the first month Obama was always perceived as the most electable Democratic candidate (although their lead over their opponents varied a bit across the campaign), which may well explain why they won their party’s nominations, but clearly cannot explain the relative low overall levels of correct voting across the entire nomination process. The

only conclusion these data will sustain is that the American voter, who does a pretty good job of voting in general elections (at the presidential level at least), does a pretty lousy job of selecting the candidate who maximizes their own values and preferences among the contenders seeking their party's nomination to begin with. Any way you look at it, the picture is not pretty.

Explaining Correct Voting in Primary Elections

Even if primary voters are not producing particularly high levels of correct voting, we may still be able to predict who is more likely to be able to vote correctly. Doing so can add to a growing body of literature on correct voting, but more importantly can provide hints as to how we might improve the situation. What factors, then, *should* predict differential levels of correct voting? Beginning with the individual level, based on past research I hypothesize that political knowledge, political interest, and the policy-based distinctiveness of the candidates will all be positively related to the probability of a correct vote (Lau and Redlawsk 2006; Lau et al. 2008), although in practice there may be relatively little variance in political interest among caucus and primary voters (compared to general election voters), who almost by definition must be very interested in politics; and relatively less distinctiveness among the candidates running in each party's nominating contests, compared to a typical general election contest. In addition, I hypothesize that specific social "orientations" or prejudices will discourage people from supporting specific candidates, and if those candidates are the "correct" choice, be negatively related to correct voting. In particular, Democrats high in racial resentment (Feldman and Huddy 2005; Sears and Henry 2005) should be unlikely to support Barack Obama, and thus more likely to vote incorrectly, if Obama is the correct choice for them.

But these individual-level predictors are all based on the limitations of human cognition, or on early-learned prejudice, and they are not likely to change much any time soon. It is much easier to contemplate changing institutional rules, if they can be shown to influence the level of correct voting. Turning then to the aggregate level, I hypothesize that all else equal, the more candidates in the choice set, the more difficult the choice and thus the lower the probability of correct voting. Figure 2 provides strong confirmation of this hypothesis at the bivariate level, and I am confident those differences will hold up under multivariate controls.

I also hypothesize that the availability of relevant political information should positively impact the probability of voters making a correct choice. Two possible measures of information availability come immediately to mind. The first, somewhat indirect, is time—or more precisely, the number of days into the year 2008 that a state's primary or caucus occurred. The longer the campaign season continued, the more familiar citizens everywhere should be with the candidates running for president. This would give an advantage to citizens voting in later nominating contests. Unfortunately, time is very strongly (and negatively) correlated with the number of candidates in the choice set. Candidates drop out as time goes on, and empirically the effect of the number of candidates on the ballot is much stronger than the effect of time. I therefore dropped time from my initial model, but will return to it at the end of the results section.

A more direct measure of information availability is the local *intensity* of the competing candidates' campaigns. Lau et al. (2008) found support for this hypothesis in general election campaigns, but it should apply just as well to primary elections. I operationalized the intensity of each candidate's campaign in a respondent's local environment using a series of items asking respondents if they could remember seeing a candidate's television ad, hearing an ad from the candidate on the radio, seeing a yard sign from the candidate, receiving mail from the candidate, or hearing about the candidate in church. I took a count of these nonvolitional indicators of exposure to a candidate's campaign as a decent measure of the intensity of the candidate's campaign in the respondents local environment. Ignoring the counts from supporters of a particular candidate (which I assumed would be biased upward), I took the mean value of each of these counts for every candidate running in a state's primary or caucus as a measure of the intensity of local campaigning by that candidate. The sum of all candidates within each party (normalized to have a 1-point range) is thus a measure of the overall availability of campaign information for each party in each state.

Finally, I hypothesized that voters in caucus states would be better able to select a correct candidate than voters in states with primaries. My reasoning is two-fold. First, it is typically "harder" to vote in a caucus than it is to vote in a primary. It usually requires a much greater time commitment, and often the caucus is located further from a voter's home than the local polling place, which results in much lower turnout for caucus elections than primaries. Thus citizens who actually show up to a caucus and vote are typically more interested in and committed to politics, and these knowledge and motivational factors should contribute to higher overall levels of correct voting in caucus states. These are exactly the reasons many people use to argue against the fairness of caucuses as a means of selecting party candidates (Mann 2009; Sabato 2009), but simply in terms of encouraging correct votes from the people who do show up, restricting turnout to the most strongly motivated partisans is a good thing. My second reason is directly related to the nature of caucuses themselves, and one of the strongest arguments in their favor: citizens usually hear arguments from their fellow citizens about why they should support one candidate over the others (Redlawsk et al. 2011). These arguments should provide exactly the type of information that people need to make wise choices, and voters in caucus states are particularly likely to be exposed to them.

These arguments beg the question of *which* measure of correct voting we should be examining. Figure 3 presents the results of almost 50 different estimates. It does not appear to matter much, among Democratic voters, but the more electability is counted, the worse Republican voters appear to do. I would argue that electability concerns *should* matter in the determination of a correct vote choice. Apart from the notoriety gained by a few people closely associated with the campaign, what does anyone get if their favorite candidate wins their party's nomination, but that candidate is subsequently eviscerated in the general election? But how *much* should electability count? Looking at Fig. 3, a weight of 3 seems to be a happy median for Democratic voters, and is a plausible choice for Republicans as well, before estimated levels of correct voting fall too low.

Table 1 reports the results of a multilevel nonlinear analysis of correct voting in each party's nominating contests, where the dependent variable is the "weighted

Table 1 Explaining correct voting in the 2008 U.S. nominating elections

	Democrats			Republicans		
	Coeff.	(S.E.)	Prob. ^a	Coeff.	(S.E.)	Prob. ^a
<i>Level-2 predictors</i>						
Intercept	-.46***	(.09)		.14 [@]	(.07)	
Number of candidates	-1.07***	(.20)	-.22	-.64***	(.11)	-.44
Overall campaign intensity	.02	(.25)	.00	.33	(.29)	.08
Caucus state	.44**	(.17)	.10	.03	(.18)	.01
<i>Level-1 predictors</i>						
Political knowledge	.30**	(.09)	.07	.30*	(.15)	.07
Political interest	.30*	(.14)	.07	.06	(.25)	.01
Candidate distinctiveness	.27	(.39)	.03	2.46***	(.53)	.52
Age	-1.44***	(.24)	-.32	.93***	(.20)	.22
Male	.44***	(.07)	.10	.01	(.09)	.00
Black	1.15***	(.15)	.27			
Hispanic	-.37**	(.11)	-.09			
Obama correct	.18**	(.07)				
Racial resentment	-1.90***	(.23)				
Obama × racial resentment	-.61***	(.26)				
Born again Christian				-.60***	(.15)	-.15
Mormon				.58*	(.24)	.14

Data come from the 2008 CCAP

[@] $p < .10$; * $p < .05$; ** $p < .01$; *** $p < .001$

^a Figures in the “Prob.” column estimate the change in the probability of a correct vote due to a “full dose” of each predictor variable holding all of the other predictors at their mean or modal value—so, 50 year old (white non-Hispanic, among Democrats; not born-again, not Mormon, among Republicans) females with average political knowledge, political interest, and candidate distinctiveness, living in primary states experiencing mean levels of campaign intensity, when there are only two candidates still competing for the party’s nomination

mean” measure of correct voting with electability concerns counted three times.⁶ Starting with the aggregate level, the number of candidates on the ballot is very important in each equation, increasing the probability of a correct vote by as much as 44% as the field is winnowed down from five to only two remaining candidates. The intensity or visibility of the campaigning by candidates has its hypothesized positive effect and is associated with an 8% increase in the probability of a correct vote, but only in the Republican party. On the other hand, caucus voters do significantly better among Democrats, where it increases the probability of a correct vote by about 10%, while the effect for Republicans is not different from zero.

Turning to the individual level, as hypothesized political knowledge increase the probability of a correct vote by about 7%. I was not sure if there would be sufficient

⁶ All analyses were repeated with the weighted sums measure, and with electability counted only once. The results vary little from those reported in Table 1.

variance in political interest in the primaries for it to help predict correct voting; there was, among the Democrats, but not in Republican nominating contests. In contrast, candidate policy distinctiveness was very important among Republican voters but not Democrats. While I have no easy explanation for why political interest should be so much more important among Democrats than Republicans, *ex post* the differences in candidate distinctiveness is easy to explain. The three major Democratic candidates were all senators (or former senators) and took very similar stands on the issues, making it extremely difficult to discriminate among them on any objective policy basis; whereas there was considerably more variance in candidate background and issue stands to help guide Republican voters.

Two fairly obvious candidate-group connections were established by voters in the 2008 nominating contests and therefore became part of the operationalization of correct voting—the link between African Americans and Barack Obama on the Democratic side, and the link between Mormons and Mitt Romney on the Republican side. Although the effects are essentially “built in” by the standard method of determining a correct vote, these two variables are included as controls in the analysis reported in Table 1. The effects are quite strong and positive, for the relatively small number of voters who were members of these two social groups, and surely represent the heuristic value of such social group memberships (Dawson 1994). But note two other equally obvious candidate-group connections that, empirically, voters did not follow—a link between gender and Hillary Clinton, on the Democratic side, and between being a born-again Christian (or a Baptist) and Mike Huckabee, on the Republican side. Female Democrats showed only a slight preference for Clinton over her two primary male rivals, and born-again Christians (and Baptists) split their votes fairly evenly between Huckabee and Romney (with a small preference for the latter). Why these dogs did not bark is a topic for future research.

But group membership also has its dark side, and it can be a reason for prejudice from non-group members. There is one prominent instance of such prejudice in Table 1, the effect of racial resentment on correct voting among Democrats for whom Obama was otherwise the correct choice. As shown in Fig. 4, racial resentment reduced the probability of voting correctly among all Democrats, but its effects were particularly pervasive among voters for whom Obama was the correct choice. Such voters have a 67% probability of voting correctly, if they were very low in racial resentment, but not even a 14% probability of voting correctly, if they were high in racial resentment. Here is a clear instance where a reason for a vote choice, and a criterion for defining a correct choice, sharply diverge. We cannot take such social orientations or world views into account in defining a correct choice, as they are neither universally shared nor universally valued, but clearly many of our voters hold tight to these social orientations, and it influences their likelihood of voting correctly for candidates who are targeted by these perspectives.⁷

⁷ There was some evidence that cosmopolitanism (Jackman and Vavreck 2011) played the same role for Republicans that racial prejudice played for Democrats, decreasing the probability that Republicans who otherwise should have preferred Huckabee, from voting for him. But this finding was not consistent across all of the different variants of the dependent variable, and I therefore dropped it from my reported analysis.

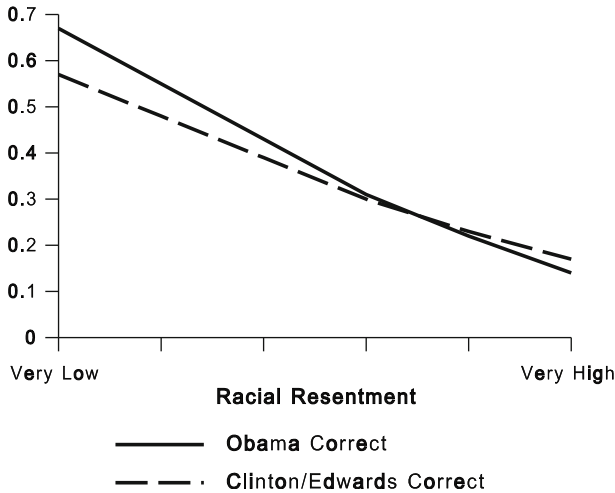


Fig. 4 Interaction of racial resentment and Obama being the correct choice on probability of correct vote, Democratic nominating contests

I want to return briefly to the indirect measure of information availability—how early or late during the campaign an election occurred. Time is positively correlated with information availability, but negatively corrected with the number of alternatives remaining in the choice set. In this data, in each party the number of alternatives was easily the more important of the two, but in certain ways this broader summary mis-characterizes the true nature of the Democratic primary. The Democrats started out with eight candidates, but once people actually started voting, the field quickly narrowed to only two—Hillary Clinton and Barack Obama. But those two candidates slugged it out for a solid 5 months. I therefore conducted a second analysis of the Democrat’s nomination, limiting attention to only those respondents who reported voting for either Clinton or Obama during that 5 month period. Time now replaces the number of alternatives in the equation, and it proves to have a positive sign and be highly significant.⁸ All else equal, this analysis suggests that the probability of a primary voter voting correctly more than doubles (from .30 to .64) across this 5-month period. This result must be taken with several very large grains of salt, as it is based on an artificially created “contest,” and the implicit assumption that states have been randomly assigned to election date, which is clearly not the case. But it is consistent with the idea that voters do learn about the candidates over time, and they can translate that increased knowledge into a greater probability of choosing correctly.

⁸ Time is operationalized as the number of days in the year 2008. It starts at 3, when the Iowa caucuses were held on January 3rd, and ends at 155, when the Montana and South Dakota primaries were held on June 3rd.

Discussion

Political scientists who have studied U.S. primary elections have long worried that the process encourages ideologically extreme and nonrepresentative candidates to be selected, and moreover gives undue influence to residents of the early primary and caucus states. These concerns are real, but I have raised a more basic, and fundamentally troubling, concern: How well do voters do in choosing their party's standard bearers? Judging by the criterion of selecting the candidate who represents their own values and concerns, the answer is not very well at all—in fact in many circumstances, no better than chance.

This finding raises several immediate questions, some methodological, others normative. At the methodological level, we can worry that the standard method of determining correct voting that I have utilized here is just not working very well when we apply it to a within-party nominating election. That is possible, but then we have to explain why a procedure that seems to work so well in a general election campaign seems to fail in a primary election campaign. It is not clear to me what additional questions I would want to ask, if I were running the CCAP the next time around. Rickershauser and Aldrich (2007) suggest that because primary candidates differ little in their actual policy stands, primary voters choose among them more on the basis of perceived policy *priorities*, and we could certainly ask questions aimed at measuring those perceptions. But the “weighted” versions of correct voting operationalized here take the voter's priorities (if not the perceived candidates' priorities) into account, but differ very little from the unweighted (equal weights) versions of those same measures. I might also recommend re-thinking how we measure “candidate” factors, which have such a strong partisan flavor (everyone thinks their own candidate is the most expert, hardworking, caring, trustworthy, strong leader) that it is difficult to discern “objective” differences among the candidates on such criteria.

Furthermore, it is not at all clear that we should judge the method as “failing” in the present circumstances. I am finding much lower levels of correct voting in primary elections than in general election contests, but is that the measure's fault? There are very good theoretical reasons that this should happen. Voting in primary elections is downright hard. There are often a lot of alternatives, who may not be all that different from each other. To the extent the measure provides predictable, theoretically consistent results—and the data in Table 1 generally suggests that it does—then we must conclude the measure has high *construct validity*, and as such it is doing everything we can ask of our measures in the social sciences.

But if the measure is valid, then who is to blame for these discouraging results, the voters or the party nominating system itself? It seems wrong to point to voters, who somehow manage to do a pretty good job in general election campaigns. True, most people do not “raise their game” sufficiently to do as well in the more difficult primary election, but then all of the informal signs in the larger political environment indicate that this choice just is not as important as the general election vote, so why should most people bother to invest the extra effort that would be required? I can therefore only point to the institutions, the rules and regulations the parties have developed over the past 40 years for selecting their nominees. Elections

provide a number of benefits, such as transparency and greater commitment to the democratic process itself. But the incentives built into the system favoring the “frontloading” of primaries and caucuses earlier and earlier in the season seem horribly misguided (Ridout and Rottinghaus 2008). Moving a state’s primary election up so that one’s citizens have more “say” in the selection of candidates also means they generally face a more difficult choice and, as a result, essentially choose randomly. This is not the type of public policy that many people would consciously want to encourage, but in 2008 and their schedules for 2012, that is exactly what most states did. The Republicans’ many winner-take-all elections quickly led to their selection of a candidate, but one who, according to my data, was the “correct” choice of only 37% of the Republicans who helped choose him. This was more than any other candidate, but is hardly a ringing endorsement.

We need a procedure that somehow weeds down the field to a manageable number (as we seem to have already), but one that simultaneously makes it difficult for a single candidate to lock up the nomination too quickly before voters have had the opportunity to get to know them better. The Democrats in fact did that in 2008, and by my estimation, it made a big difference by the end of the process in the proportion choosing correctly. But most observers were complaining about the length of the process rather than commenting on its virtues. The longer we have to become familiar with a set of candidates, the more we ought to know about them, and the greater should be the probability of choosing correctly among them. The “rush to judgment” that is encouraged by the Republican party’s frequent use of winner-take-all rules in their nominating contests is in many ways a very real detriment to correct voting.⁹ In any case, I now add my voice—and some new evidence—to the growing chorus calling for serious reform of the procedures we have developed to pick the candidates who compete in our most important elections.

Furthermore, if primary elections are generally going to provide a more challenging decision context than general election campaigns—and in the U.S., with its deeply entrenched two party system, this seems almost inevitable—then I would also argue that we would be better off *limiting*, rather than encouraging, turnout, assuming that the people who would still turn out would be the most knowledgeable and committed party members. One of the expressed goals of the McGovern–Fraser commission was to make the selection of the party’s nominee more democratic, and one visible means of achieving that goal was increasing turnout in the party’s nomination contests. I agree with the goal but not the means. All party members should have the *opportunity* to become involved in selecting candidates, but we should ask more of those who take advantage of that opportunity than simply casting a random ballot. If more states held caucuses rather than primaries, for example, turnout overall would be much lower but those actually voting would be, on average, much more informed about the candidates. I estimate that voters in the current caucus states are about 5–6% more likely to vote correctly than are voters in primaries. Political knowledge and political interest together add another 6–14% to

⁹ See also Carey and Hix (2011). Interestingly, the rules adopted by the Republican party for the 2012 election make proportional allocation of delegates a *punishment* for any state (like Florida and Michigan in 2008) that jumps out of line and holds its nominating contest too early.

the probability of casting a correct vote. Shouldn't we be encouraging these people to choose their party's contenders? There are, of course, drawbacks to this plan,¹⁰ but if a party's goal in selecting candidates to run in general election campaigns is *representation*, then somehow we need to get primary voters to cast more informed ballots.

The data generally suggest that intense, high level political campaigns contribute to higher levels of correct voting, presumably by providing easy access (in the guise of television ads, local campaign headquarters, and the like) to highly relevant information. This finding argues that any attempt to limit campaign spending would be misguided, at least from the perspective of helping voters choose the candidate who best represents their own values and priorities.¹¹ But no matter how full a candidate's campaign war chest, they cannot simultaneously sustain high level campaign efforts in more than a few states at any one time. A candidate and his or her top advisers can only be in one place at a time, after all. Indeed, if a large number of states all decide to hold their primary on the same day (as has become the practice on "Super Tuesday" over the past few election cycles), the citizens of those states are poorly served because the candidates must spread their finite campaign resources too widely. If we simply compare the average campaign intensity score of the 24 states contesting their primary or caucus on "Super Tuesday" (Feb 5) with the campaign intensity of all remaining states holding their election on every other day, we see a 14–20% lower mean effort on that 1 day ($p < .08$ for the Democrats, $p < .03$ for Republicans) with, concomitantly, a 2–3% lower probability of voting correctly among citizens choosing on that 1 day. Some sort of more coordinated effort where, say, at most seven or eight states held their nominating contest every 2 weeks, would serve voters (and the country) far better by allowing all citizens a better opportunity to hear the campaign messages the candidates would like us to hear. But that would require more control over the scheduling of the state nominating process than either national party seems willing to exert.

Acknowledgments I want to thank Scott McKee and other panelists for their comments on an earlier version of this paper, and David Redlawsk for a very careful reading and several excellent ideas.

Appendix: Details of Constructing Correct Voting Measures

Four conceptually distinct sets of items go into the calculation of candidate "utility scores" that determine which candidate best represents a respondent's own political values and priorities, and therefore which candidate a respondent "should" support in the primary election—policy stands, candidate-group connections, "electability"

¹⁰ Including the recent finding that turnout in primary and caucus elections spills over into a higher probability of voting in the subsequent general election campaign (Jones-Correa and Walker 2011), which I *would* like to encourage. Gerken and Rand (2010), building off of deliberative polls, offer an interesting but more radical idea that might accomplish my goal, a sort of "citizen assembly" that would be randomly chosen and charged with learning about and "vetting" potential party nominees.

¹¹ I do favor some sort of public financing of election campaigns that would help equal out imbalances in campaign finances, however.

beliefs, and personality traits.¹² In every case we accept survey respondent's reports of their own values and priorities, but try to find some semi-objective expert judgment about how closely each candidate actually "fits" those same values and priorities.

First, the survey asked respondent's to place themselves and the major presidential candidates on five policy scales: a question about illegal immigrants, a question about government health care, a question about increasing taxes on the rich, a question about the war in Iraq, and an overall liberal-conservative scale. I relied on the mean ratings of expert respondents (those scoring in the top quarter of the distribution of political knowledge) to determine where the candidate's stood on these issues. Respondents were also asked their opinion about the circumstances under which abortion should be legal, whether we should take action to slow climate change or protect jobs instead, how strongly they favored capital punishment, whether gay couples should be allowed to have legalized civil unions, and whether gays and lesbians should be allowed to adopt children; but they were never asked to place the major presidential candidates on those same issues. I therefore recruited a separate panel of experts (seven graduate students) to place these candidates on these five additional policy scales after reading an extensive file of information about the candidates' actual stands on these issues. I used the mean responses of these seven experts as an objective reading of where the candidate's stood on these five additional issues so that I could calculate agreement with the candidates on these five additional issues. Table 2 reports the "objective" candidate placement scores calculated for each candidate on each issue.

I then calculated each respondent's proximity to each candidate on these issues.¹³ The policy proximity scores were reversed (so that policy agreement is scored high) and rescaled to vary between -1 and $+1$, and then added to each candidate's utility score. *These policy agreement scores only factor into the candidate utility ratings for respondents who cared enough about an issue to report a position on it.*

Respondents were also asked to indicate how well three positive traits (strong leader, trustworthy, has the right experience) described each candidate. I assume that everyone universally prefers strong, trustworthy, and experienced leaders. Responses to these questions were recoded to range between -1 "Not Well at All" to $+1$ "Extremely Well." I again relied on the mean ratings of our expert respondents, and added these three ratings to each candidate's utility score, but again only if respondents showed that they cared enough about these attributes that they answered the questions. These "objective" trait scores are reported at the bottom of Table 2.

To estimate candidate-group linkages, I considered whether a majority of (expert) members of 11 easily-identifiable social groups—men, women, blacks, whites, Latinos, the working class, rich people, Southerners, born-again Christians, Mormons, Muslims—voted for any particular candidate in either party's primary.

¹² Party identification and retrospective evaluations of the incumbent's job performance, which are important parts of the calculation of candidate utility scores in a *general* election campaign, are irrelevant in a primary election because those considerations apply equally to all candidates within a given party (when no incumbent is seeking re-election).

¹³ All analyses were repeated with a directional algorithm (Rabinowitz and MacDonald 1989), which produces essentially identical results as those reported herein.

Table 2 Objective measures of candidates actual policy stands and trait qualities

	Democrats			Republicans				
	Clinton	Edwards	Obama	Giuliani	Huckabee	McCain	Romney	Thompson
Liberal–conservative identification ^c	2.5	2.1	2.2	3.0	3.5	3.2	3.7	4.1
Immigration policy ^a	.1	.1	.1	.7	.8	.2	1.5	1.5
Health Care Policy ^a	.5	.4	.7	1.5	1.3	1.5	1.5	1.5
Tax Policy ^b	1.7	1.3	1.4	3.3	2.8	3.3	3.5	3.6
Iraq War Policy ^d	1.7	.9	1.3	2.8	2.4	2.9	2.7	2.9
Abortion Policy ^b	1.4	1.3	1.7	1.9	3.6	2.9	2.9	2.7
Global Warming Policy ^c	2.3	1.7	2.3	3.7	3.1	3.0	4.0	4.7
Policy toward Civil Unions ^b	1.0	1.6	1.1	1.9	3.9	3.6	3.4	2.9
Policy toward Gay Adoptions ^c	1.1	1.1	1.1	1.2	1.9	1.9	1.5	1.6
Policy on Capital Punishment ^b	1.6	1.7	2.4	1.3	1.4	1.0	1.7	1.3
Strong Leader ^d	1.8	1.8	2.1	2.1	1.5	2.1	2.1	1.8
Trustworthy ^d	1.2	2.1	2.1	1.5	1.8	1.8	1.8	2.1
Has the Right Experience ^d	1.8	1.8	1.8	1.8	1.2	2.1	2.1	1.5

^a 0–2 scale; ^b 1–4 scale; ^c 1–5 scale; ^d 0–3 scale; ^e 1–2 scale

All policy scales are coded so the liberal views are scored low and conservative views high

This procedure assumes that politically expert members of these social groups are able to ascertain important (linked-fate?) candidate-group linkages that are independent of policy stands. Empirically, there were only two such linkages, between Barack Obama and black Democrats, and between Mitt Romney and Mormon Republicans. I therefore created dummy variables representing membership in these two demographic categories, and added them to the utility scores associated with Obama or Romney, respectively.

An important consideration for voters in their party’s nomination contests is how “electable” the different candidates are—that is, how likely they are to win November’s general election campaign. CCAP respondents were asked this question about both parties’ candidates in the baseline and January waves of interviews, and just about Clinton and Obama in the March interview. No consensus had developed among experts in either party about which candidate was the most electable until the March wave, when a large majority of experts believed Obama

Table 3 Candidate electability estimates across the nominating campaign

	Democrats			Republicans				
	Clinton	Edwards	Obama	Giuliani	Huckabee	McCain	Romney	Thompson
December 2007	2.5	10.8	7.2	-3.5	-9.0	3.0	-11.1	-12.7
Jan 1–Jan 19	6.5	10.0	10.2	-13.0	-12.1	4.5	-14.8	
Jan 20–Jan 31	8.3	*	12.2	-17.0	-13.5	2.0	-12.5	
Feb 1–Feb 5	5.4		7.8		^a	.4	-13.5	
Feb 6–Mar 4	-1.8		.4		^a	.7		
Mar 5–Mar 11	.7		3.0					
Mar 12–Apr 10	-.3		3.4					
Apr 11–May 6	-4.0		4.8					
May 7–May 20	-4.0		3.0					
May 21–June 3	-4.0		3.0					

Data report the average results of all “horserace” poll questions asked during different periods of the campaign (as reported by Real Clear Politics.com) pitting each candidate against all possible opponents from the opposite party. So for example in December of 2007, just before the nominating contests began, Edwards would have beaten all five Republican candidates “if the election were today,” most of them handily

^a No new poll data reported for this period, so the last available estimate is repeated

was much more likely to win in November than Clinton. This did not give me much to work with.¹⁴ Instead, I relied on the mean results of all survey questions (as reported by Real Clear Politics.com) pitting each viable Democrat against each viable Republican. I broke the nomination campaign into ten distinct periods, and calculated the average margin (across all surveys conducted during that period) by which each candidate would win or lose (“if the election were today”) against each remaining candidate from the other party. The mean of these average margins (averaging across all possible opponents) is the “electability” of that candidate during that time period. The data for all 8 candidates are reported in Table 3. I again normalized these electability scores within party so that they had a maximum range of -1 to +1 across the entire nomination process.

An *unweighted sum* measure of candidate utility is computed by simply adding together these 15 different criteria of judgement for each of the major candidates competing for their party’s nomination—Hillary Clinton, John Edwards, and Barack

¹⁴ The surveys also asked respondents which candidate they thought was the most *viable*—that is, which candidate they thought was most likely to win their party’s nomination. Despite Abramson et al.’s (1992) use of such viability beliefs in their study of “sophisticated” voting in the 1988 presidential primaries, and their claim (in fn 1) that it does not matter much empirically whether they use viability or electability beliefs in their analysis, I see obvious reasons why voters should consider electability in their vote choices, but can think of no logic for similarly including viability beliefs into a determination of correct voting, *after* accounting for electability concerns. Viability clearly does matter to the big donors who are providing the money to the different candidates who are seeking the nomination, and to the campaign consultants who will continue to be employed during the general election campaign if the pick a winner in the primaries. But I cannot see any rationale for why the average citizen would be any “better off” if they voted for the winning candidate in a primary election. Abramowitz (1989) might disagree.

Obama on the Democratic side, Rudy Guiliani, Mike Huckabee, John McCain, Mitt Romney, and Fred Thompson, on the Republican side. Similarly, an *unweighted mean* measure of candidate utility is computed by averaging together these different criteria of judgement for each of the major candidates. Many respondents do not provide valid responses on all of these criteria of judgment, of course, but we use as many as we can in calculating the candidate utility scores. With the exception of the candidate-group linkages, however, if a criterion of judgment is available for one candidate it is available for all candidates.

The above procedure treats all 15 criteria of judgment as if they are equally important—a good first approximation, but one my intuitions suggest does not match reality for most people. I also devised simple measures of how important each criterion was to each respondent. People were asked to express an opinion on numerous occasions for almost all of these different criteria of judgment, either their own opinion on an issue across multiple waves of the study, or their beliefs about where different candidates stand on an issue or how a particular trait would apply to different candidates, or both. Thus it was almost always possible to compute implicit “importance” weights for each of the criteria of judgment—simply the proportion of relevant questions on which a respondent provided a valid response. A *weighted sum* measure of candidate utility is computed for each of the eight candidates by multiplying each criterion of judgment by its importance weight and then summing across these 15 products. An *weighted mean* measure candidate utility is computed for each candidate by averaging together the 15 different (criterion \times importance weight) products.

The method described above treats each distinct consideration equally, but as described above I also considered the possibility that electability considerations should be counted more than once (in fact, up to ten times). The analysis in Table 1 utilizes the weighted-sum measure with electability counted three times, but the results are pretty much the same if we use any of the other three operationalizations of correct voting, and count electability less.

References

- Abramowitz, A. I. (1989). Viability, electability, and candidate choice in a presidential primary election: A test of competing models. *Journal of Politics*, 51(4), 977–992.
- Abramson, P. R., Aldrich, J. H., Paolino, P., & Rohde, D. (1992). “Sophisticated” voting in the 1988 presidential primaries. *American Political Science Review*, 86(1), 55–69.
- Aldrich, J. H. (1980). *Before the convention: Strategies and choices in presidential nomination campaigns*. Chicago: University of Chicago Press.
- Bartels, L. M. (1988). *Presidential primaries and the dynamics of public choice*. Princeton, NJ: Princeton University Press.
- Bartels, L. M. (1996). Uninformed votes: Information effects in presidential elections. *American Journal of Political Science*, 40(1), 194–230.
- Carey, J. M., & Hix, S. (2011). The electoral sweet spot: Low-magnitude proportional electoral systems. *American Journal of Political Science*, 55(2), 383–397.
- Cohen, M., Karol, D., Noel, H., & Zaller, J. (2008). *The party decides: Presidential nominations before and after reform*. Chicago: Chicago University Press.

- Converse, P. E. (1964). Attitudes and non-attitudes: Continuation of a dialogue. In E. Tufte (Ed.), *The quantitative analysis of social problems*. Reading, MA: Addison-Wesley.
- Converse, P. E. (1975). Public opinion and voting behavior. In F. Greenstein & N. Polsby (Eds.), *Handbook of political science* (Vol. 4, pp. 75–170). Reading, MA: Addison-Wesley.
- Crotty, W. J., & Jackson, J. S. (1985). *Presidential primaries and nominations*. Washington, DC: CQ Press.
- Dawson, M. C. (1994). *Behind the mule: Race and class in African-American politics*. Princeton, NJ: Princeton University Press.
- Delli Carpini, M., & Keeter, S. (1996). *What Americans know about politics and why it matters*. New Haven, CT: Yale University Press.
- Dwyer, C. E. (2011). *Packed primaries and empty caucuses: Voter turnout in presidential nomination contests*. Paper presented at the 69th Annual Meeting of the Midwest Political Science Association, Chicago, IL, 31 Mar–3 Apr 2011.
- Feldman, S., & Huddy, L. (2005). Racial resentment and White opposition to race-conscious programs: Principles or prejudice? *American Journal of Political Science*, 49(1), 168–183.
- Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd ed.). New York: McGraw-Hill.
- Geer, J. G. (1988). Assessing the representativeness of electorates in presidential primaries. *American Journal of Political Science*, 32(4), 929–945.
- Geer, J. G. (1989). *Nominating presidents: An evaluation of voters and primaries*. Westport, CT: Greenwood Press.
- Gerken, H. K., & Rand, D. B. (2010). Creating better heuristics for the presidential primary: The citizen assembly. *Political Science Quarterly*, 125(2), 233–253.
- Gopoian, J. D. (1982). Issue preferences and candidate choice in presidential primaries. *American Journal of Political Science*, 26(3), 523–546.
- Ha, S. E., & Lau, R. R. (2010). Personality traits and correct voting. (under review).
- Jackman, S., & Vavreck, L. (2009). *Constructing the 2007–2008 Cooperative Campaign Analysis Project Sample*. Unpublished manuscript, Stanford University.
- Jackman, S., & Vavreck, L. (2011). *Cosmopolitanism*. Unpublished manuscript, Stanford University.
- Jones-Correa, M., & Walker, A. (2011). *Primary voting and general election turnout*. In Paper presented at the 69th Annual Meeting of the Midwest Political Science Association, Chicago, IL, 31 Mar–3 Apr 2011.
- Kam, C. D. (2005). Who toes the party line? Cues, values, and individual differences. *Political Behavior*, 27(2), 163–182.
- Kamarck, E. C. (2009). *Primary politics: How presidential candidates have shaped the modern nominating system*. Washington, DC: Brookings Institution Press.
- Kaufman, K. M., Gimpel, J. G., & Hoffman, A. H. (2003). A promise fulfilled? Open primaries and representation. *Journal of Politics*, 65(2), 457–476.
- Kenney, P. J., & Rice, T. W. (1994). The psychology of political momentum. *Political Research Quarterly*, 47(4), 923–938.
- Kinder, D. R. (1998). Communication and opinion. In *Annual review of political science* (Vol. 1, pp. 167–197). Stanford, CA: Annual Reviews.
- Lau, R. R. (2003). Models of decision making. In D. O. Sears, L. Huddy, & R. Jervis (Eds.), *Oxford handbook of political psychology* (pp. 19–59). New York: Oxford University Press.
- Lau, R. R., Andersen, D. J., & Redlawsk, D. P. (2008). An exploration of correct voting in recent U.S. presidential elections. *American Journal of Political Science*, 52(2), 395–411.
- Lau, R.R., Patel, P., Fahmy, D., & Kaufman, R. (2012). Correct voting across 33 Democracies (under review).
- Lau, R. R., & Redlawsk, D. P. (1997). Voting correctly. *American Political Science Review*, 91(3), 585–599.
- Lau, R. R., & Redlawsk, D. P. (2001). Advantages and disadvantages of cognitive heuristics in political decision making. *American Journal of Political Science*, 45(4), 951–971.
- Lau, R. R., & Redlawsk, D. P. (2006). *How voters decide: Information processing during election campaigns*. New York: Cambridge University Press.
- Lodge, M., & Hamill, R. (1986). A partisan schema for political information processing. *American Political Science Review*, 80(2), 505–519.
- Mann, T. E. (2009). Is this any way to pick a president? Lessons from 2008. In S. S. Smith & M. J. Springer (Eds.), *Reforming the presidential nomination process*. Washington, DC: Brookings Institution Press.

- Marshall, T. R. (1984). Issues, personalities, and presidential primary voters. *Social Science Quarterly*, 65(3), 750–760.
- Mayer, W. G. (2000). *In pursuit of the White House 2000: How we choose our presidential nominees*. Chatham, NJ: Chatham House.
- Nisbett, R. E., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.
- Norrander, B. (1986). Selective participation: Presidential primary voters as a subset of general election voters. *American Politics Quarterly*, 14(1), 35–53.
- Norrander, B. (1989). Ideological representativeness of presidential primary voters. *American Journal of Political Science*, 33(3), 570–587.
- Norrander, B. (2006). The attrition game: Initial resources, initial contests and the exit of candidates during the US presidential primary season. *British Journal of Political Science*, 36(3), 487–508.
- Patel, P. (2010). *The effect of institutions on correct voting*. PhD Dissertation, Rutgers University.
- Patterson, T. E. (2009). Voter participation: Records galore this time, but what about next time? In S. S. Smith & M. J. Springer (Eds.), *Reforming the presidential nomination process*. Washington, DC: Brookings Institution Press.
- Paulson, A. (2009). Party change and the shifting dynamics in presidential nominations: The Lessons of 2008. *Polity*, 41(3), 312–330.
- Polsby, N., & Wildavsky, A. (1971). *Presidential elections: Strategies of American electoral politics* (3rd ed.). New York: Charles Scribner's Sons.
- Polsby, N., Wildavsky, A., & Hopkins, D. A. (2008). 1971. *Presidential elections: Strategies and structures of American politics* (12th ed.). Lanham, MD: Rowman & Littlefield.
- Rabinowitz, G., & MacDonald, S. E. (1989). A directional theory of issue voting. *American Political Science Review*, 83(1), 93–121.
- Redlawsk, D. P., Tolbert, C. J., & Donovan, T. (2011). *Why Iowa? How caucuses and sequenced primaries improve the presidential nomination process*. Chicago: University of Chicago Press.
- Rickershauser, J., & Aldrich, J. H. (2007). It's the electability, stupid'—or maybe not? electability, substance, and strategic voting in presidential primaries. *Electoral Studies*, 26(2), 371–380.
- Ridout, T. N., & Rottinghaus, B. (2008). The importance of being early: Presidential primary front-loading and the impact of the proposed western regional primary. *PS: Political Science & Politics*, 41(1), 123–128.
- Sabato, L. J. (2009). Picking presidential nominees: Time for a new regime. In S. S. Smith & M. J. Springer (Eds.), *Reforming the presidential nomination process*. Washington, DC: Brookings Institution Press.
- Sears, D. O., & Henry, P. J. (2005). Over thirty years later: A contemporary look at symbolic racism. *Advances in Experimental Social Psychology*, 37, 95–150.
- Shafer, B. E. (1983). *Quiet revolution: The struggle for the Democratic Party and the shaping of post-reform politics*. New York: Russell Sage Foundation.
- Williams, D. C., Weber, S. J., Haaland, G. A., Mueller, R. H., & Craig, R. E. (1976). Voter decisionmaking in a primary election: An evaluation of three models of choice. *American Journal of Political Science*, 20(1), 37–49.
- Wright, G. C. (2009). Rules and the ideological character of primary electorates. In S. S. Smith & M. J. Springer (Eds.), *Reforming the presidential nomination process*. Washington, DC: Brookings Institution Press.