Jumping on the Bandwagon or Jumping Ship?
Testing Alternative Theories of Vote Share Overestimation

Abstract

Post-election poll results typically overstate the proportion of people who voted for winning candidates at all levels of government. Using original survey data and the American National Election Study, we test several alternative explanations of this apparent post-election "bandwagon" effect: conventional ones include expectations that respondents misrepresent how they voted to save face, genuinely forget how they voted, or experience shifts in opinion just before an election. We develop and test an unexplored alternative hypothesis, that post-election surveys inflate the winner's vote because a greater proportion of people who voted for the winning side want to participate in a post-election survey than people who voted for the loser. We devise empirical tests to distinguish and test each of these hypotheses. We find evidence that, rather than misrepresenting their votes to poll-takers, people who voted for the losing side are less likely to participate in post-election surveys.

Introduction

Exit polls and other public opinion surveys conducted after an election often indicate winning candidates have a higher level of support than they did in the election itself. The long-standing interpretation of this exaggerated support for the winner is that survey respondents incorrectly indicate support for the winning side (Atkeson 1999, Wright 1990). On the face of it, this may be a welcome, if puzzling, insight. It is a welcome finding because it helps to underscore the value of elections in providing legitimacy to political leaders (Harrop and Miller; 1987: 244-269). The very fact of winning an election seems to confer on winners several additional points of support among voters.

As comforting as this pattern may be for assessing the value of elections, a puzzle remains of why we see this overestimation. Beginning with Thomsen's 1938 paper, a series of studies seek to explain this post-election bump in support (e.g., Atkeson 1999, Beasley and Joslyn 2001, Belli et al. 1999, Ginsberg and Weissberg 1978, Wright 1990). Grounded in the psychology of voters, this research infers that losers do not simply become reconciled to winners but, in effect, misrepresent or misremember having supported the winner all along. This is the case not simply for primary elections (Atkeson 1999) but for general elections, too. The most prominent explanation for vote share, as well as turnout, overestimation[1] is that survey respondents deliberately dissemble their voting behavior to conform to perceived norms of social desirability. That is, they "overreport," claiming to have voted for the winner when they really voted for someone else (and to have voted when they really did not). In Atkeson's turn of phrase, respondents jump on the "bandwagon" after the election (1999).

We offer a heretofore unexplored alternative explanation for the pattern: non-response bias. We anticipate that people who voted for the losing side in an electoral contest may actually decline to answer post-election surveys proportionally more often than those who voted for the winning side. Rather than jumping on the winner's bandwagon in post-election polls, they "jump ship" and refuse to take part in the political discourse represented by the polling enterprise. In this case, vote share overestimation owes to overrepresentation of electoral winners and under-representation of losers in the sample, rather than to willful overreporting or memory failure. Despite an intense preoccupation with non-response among survey methodologists, scholars of elections have paid surprisingly little attention to non-response bias as a cause of vote share overestimation.[2] We show this non-response bias to hold in both candidate and, in a new area for this literature, ballot proposition elections.

Our analytic strategy is two-pronged. First, we assess classification error produced by comparing model-predicted with self-reported vote preferences for both the pre- (t_1) and post-electoral (t_2) cross-sections. We distinguish empirically among explanations for vote share overestimation using data from an original survey designed to detect shifts in stated support for California ballot propositions before and after an election. If classification error—that is, the proportion of respondents who report a vote for the winning side when the model predicts a vote for the losing side—is higher at t_2 than at t_1, respondents on the whole exaggerated the extent to which they voted for the winning side, consistent with the social desirability hypothesis. More or less equal classification error rates at t_1 and t_2 constitute evidence of non-response bias.

The second prong of our analytical strategy explores the generalizability of our findings using American National Election Study (ANES) panel data. We compare re-interviewed respondents' pre-election vote preferences with those who dropped out of the panel. If non-response bias obtains, re-interview rates should be higher for t_1 respondents who intended to vote for the eventual winner than those who intended to vote for the losing candidate.

Our research has important methodological and substantive implications. Substantively, the frequent overestimation of winning candidates' vote shares may create perceptions that the winners received more votes than they actually did and, thus, inflate victorious candidates' putative mandate following an election (Wright 1990, Atkeson 1999). Methodologically, understanding the mechanism underlying vote overestimation should inform choices for *post hoc* statistical adjustments to survey data, such as appropriate weighting schemes and selection models.[3] We also hope to shed light on potential survey respondents' psychological motivations in choosing whether to answer survey questions truthfully and whether to participate in the first place. We begin by reviewing the puzzle itself, overreport of winner's vote share in post-election surveys.

Exaggerated Support for Winners

Studies that seek to explain why people vote have long noted that post-election polls routinely overestimate the percentage of people who report having voted (e.g., Belli, Traugott, Young and McGonagle 1999). Wolfinger and Rosenstone found that in the American National Election Study (ANES), taken after every presidential and mid-term election since 1948, the percentage of respondents who report having voted is always between 5% and 20% higher than official turnout figures provided by the Federal Electoral Commission (1980: 115, see also Deufel and Kedar 2000: 24). Validated vote studies comparing self-reported voting behavior on post-electoral surveys to voting records maintained by county registrars, also find large differences between self-reported and actual turnout (Silver, Anderson and Abramson 1986).

Research investigating why people vote as they do also finds that post-election polls overestimate the percentage of people who report having voted for the winning candidate (Wright 1990). Averaging over ANES studies since 1952, Wright calculated that the "pro-winner" bias was 4.0% in U.S. Senate races, 4.7% in gubernatorial contests, and (between 1978 and 1988) 7.0% in races for the U.S. House of Representatives (1993: 295). Eubank and Gao (1984) find even bigger effects, a disparity of 14.0% between the average survey-reported vote share for incumbents in House races and their average share on ballot returns. Atkeson (1999) shows systematic post-election survey vote overestimation for presidential primary winners 1972-1992.

Prevalent explanations of both turnout and winning vote share overestimation largely point to respondent misreports—and more specifically, to overreports—of their own behavior as the culprit. Survey respondents overreport when they inaccurately claim to have voted when they did not and to have voted for the winning candidate when they voted for someone else. Conventional thinking on this problem anticipates that overreporting occurs because respondents misrepresent or misremember whether and for whom they voted.[4]

Respondents may also "forward telescope a remote voting experience," transforming prior votes into a vote cast in the most recent election (Belli et al. 1999: 91)—especially if, after an election, media coverage of the winning candidate emphasizes positive attributes of the winner and his/her opponent's campaign falls silent. The consequent revision of (remembered) history may well be especially prevalent among less sophisticated voters. When such voters cannot reconstruct how they voted accurately, they substitute judgments they make at the time of the survey for those they made at the time of the election, and may claim to have voted for someone other than the candidate for whom they really voted (Wright 1993).[5]

We also explore another alternative, the possibility that some voters have an 11^{th}-hour change in voting intentions. If large numbers of people who had supported a disadvantaged candidate change their minds after the last pre-election interviews and vote for the winning candidate, this could account for at least some of the sharp discrepancies observed between the pre- and post-election polls, and between the pre-election poll and actual voting results. Of course, the longer the pre-election poll is taken before the election, the less accurately it will predict election results (Crespi 1988).

Here, however, we focus on a (surprisingly) over-looked explanation: *non-response bias*. People who abstained or voted for the losing side may disproportionately refuse to take the survey, compared to actual supporters of the winner. Higher non-response among electoral losers and abstainers would result in overrepresentation of voters and winners and, consequently, overestimation of the percentages of citizens who voted, or voted for winning candidates. Citizens who cast ballots for losing candidates, momentarily dispirited, may be disinclined to take a survey about an election whose outcome they find disagreeable. It is possible, in theory, for overestimation to occur even when all respondents report their voting behavior truthfully, although it is likelier that misreporting and non-response bias contribute to overestimation in tandem.

While this is a straightforward argument, in practice it involves quite complex issues. Research on sampling and survey methodology has long grappled with the potential biases produced by non-response (e.g., Berinsky 2004, Cochran 1977, Groves and Couper 1998, Groves et al. 2002, Kish 1965). Practitioners distinguish between item non-response, in which respondents neglect to answer some (but not all) questions on a survey, and unit non-response, in which some people selected into the sample fail to take the survey altogether. We are concerned here with unit non-response, which may bias estimation when the probability of taking a survey is different for different segments of a population and there are significant differences between segments: "For the bias to be important, a large nonresponse must coincide with large differences between the means of ... two segments" (Kish 1965: 535).

Some of the existing literature does mention non-response bias in passing but avoids a more thorough exploration of the possibility. For example, Atkeson shows considerable pro-winner bias in the ANES 1988 Super Tuesday post-election poll. Since African American voters were underrepresented in the sample, the survey results understated support for Jesse Jackson and overstated support for the eventual winner, Michael Dukakis. *Post hoc* weighting adjustments brought vote share estimates in line with actual results, which raises the possibility that non-response bias drove overestimation of Dukakis's vote share—and, conceivably, other results (Atkeson 1999: 207). Similarly, in his study of presidential and congressional races, Wright speculates that "hostile" respondents are not likely to misreport vote choice intentionally, but rather "would generally refuse to be interviewed in the first place" (1993: 293). Neither Atkeson nor Wright, however, develops these asides into a full-blown consideration of non-response bias as a cause of vote preference overestimation.

Distinguishing Explanations of Winner Vote Overestimation

We identify four potential explanations for overestimating the vote share of winners in post-election surveys:

1. **Social desirability**: Respondents recall how they voted but deliberately misreport their electoral preference, embarrassed to admit voting for the losing side.

2. **Non-response bias**: The survey sample overrepresents citizens who voted for the winning side because those who voted for the losing side or abstained are less likely to participate in a post-election poll.

3. **Memory lapse**: Survey respondents, unable to recall how they voted, misreport their electoral preference.

4. **Late opinion shift**: Large numbers of voters change their minds too late for pre-election polls do not capture the shift, which is registered in the post-election survey.

The two main hypotheses under consideration—socially desirable overreporting and non-response bias—both favor the winner. Neither is directly testable because it is impossible to know who survey respondents really voted for. The ANES "vote validation" studies compared individual, survey-reported voting to country registrar records of actual voting; thus, researchers know which respondents reported having voted accurately and which misreported. While we can know who voted, the secret ballot means we cannot know for whom. Detecting evidence of our hypotheses thus necessarily implies drawing inferences indirectly from patterns we observe in the data measured against patterns we would expect to observe under both hypotheses.

To adjudicate among the hypotheses, we devised a two-pronged inferential strategy. We apply one prong or other, or both, to three electoral contexts, depending on the possibilities afforded by each. [6] In the first prong, which we dub the "Classification Error Comparison Method," we model vote preferences for the pre-election (t_1) cross-section and obtain both coefficient estimates and individual-level predicted probabilities of voting for the winner. We assume that the t_1 model is the true model of vote preferences (or close to it). Then, we plug the post-election (t_2) cross-sectional data into the t_1 coefficients to predict individual probabilities that a post-election respondent will report having voted for the winning side. Next, we calculate t_1 and t_2 classification error rates by comparing model-predicted with survey-reported vote preferences at both time periods. Equal classification error rates constitute evidence of non-response bias. On the other hand, if the t_2 classification error rate exceeds the t_1 benchmark, overreporting has occurred.

Given the assumption that the t_1 model approximates the true model of vote preferences reasonably well, erroneous predictions should occur at about the same rate in t_1 and t_2. If the t_2 false negative rate exceeds the t_1 baseline—that is, if the model predicts a vote for the losing side, conditional on self-reported votes for the winning side, more often at t_2 than at t_1—respondents on the whole exaggerated the extent to which they voted for the winning side, consistent with the social desirability hypothesis. On the other hand, if false negative classification error rates are about the same at t_1 and t_2, no overreporting has occurred.

[Table 1a-c about here]

A simple example might help illustrate the way non-response bias leads to winning vote share overestimation. Imagine that in a two-candidate election, Candidate A bests Candidate B by 54% to 46%. A pollster conducts a pre-election survey (N=1,800). An equal percentage, around 70%, of Candidate A and Candidate B supporters take the survey. The survey estimates the eventual winner's vote share accurately (as shown in Table 1a's marginal row percentages). The same pollster then carries out a post-election survey on a different cross-section, drawing a sample of 1,800 that again mirrors population vote shares. This time, though, 80% of Candidate A supporters and 58% of Candidate B supporters take the survey. Everyone reports voting preferences accurately but the survey estimates a vote share of almost 62% for Candidate A, inflating the victor's support by nearly eight percentage points (as shown in Table 1b's marginal row percentages). Different response rates between supporters of Candidate A and Candidate B, not overreporting, account for the overestimation.

To see how equal classification error rates at t_1 and t_2 constitute evidence of non-response bias, consider a standard statistical model for the pre-election sample that predicts about 80% of intended votes accurately. The remaining 20%—classification error—is divided between false negatives and false positives proportionally to each candidate's intended vote share. Table 1a depicts this scenario, representing classification errors as percentages within each category of reported vote intention (shown by the row percentages), rather than as a percentage of the total sample; that is, classification error is conditional on reported vote. A vote choice model is then estimated on the post-election data. The t_2 model coefficients are the same as those at t_1 (by assumption, the true model). Again, the model predicts about 80% of vote preferences correctly, and allots the 20% classification between false negatives and false positives in proportion to reported vote shares. As in Table 1a, classification error in Table 1b is conditional on reported vote preferences. The percentage of respondents who report having voted for the winner at t_2 is greater than at t_1 (reflected in Table 1a's and 1b's marginal reported vote percentages), but conditional classification error is the same (shown in Table 1a's and 1b's cell row percentages). The post-election survey overestimates Candidate A's vote share, but the t_2 model predicts reported voting exactly as well as the t_1 model because no overreporting has occurred.

As Tables 1a-c indicate, classification error comprises "false positives" (the upper right-hand cell, in which, given a reported or intended vote for the loser, the model predicts a vote for the winner) and "false negatives" (lower left-hand cell, with predicted votes for the loser, conditional on reported votes for the winner). We assume that all classification error at t_1 consists simply of the random deviations from model predictions inherent in all statistical modeling: respondents are voting contrary to how they "should" vote (according to the model). On the other hand (as Tables 1b and 1c show), classification error at t_2 consists of random error plus misreporting. False positives result from random error plus underreporting (in which respondents state they voted for the loser when they voted for the winner), and t_2 false negatives, from random error plus *overreporting*. Since only false negatives bear directly on the social desirability hypothesis, and since (in turnout studies, at any rate) underreporting is both empirically negligible and randomly distributed (see, e.g., Silver et al. 1982, Presser and Traugott 1992), we ignore false positives and analyze only false negative classification error.

Technically, our interpreting equal (conditional) classification error rates at t_1 and t_2 as evidence of non-response bias rests on the *invariance property of odds ratios*. Odds ratios (of which coefficients in logistic regression models are natural logarithms) are "invariant to changes in marginal distributions, since such changes are translated to proportional increases or decreases across rows or columns" (Powers and Xie 2008: 76-77). Here, the marginal probability of voting for the winner changes from Table 1a to 1b, but the odds ratio is the same in both tables:

$$\theta = \frac{[\Pr(PV_{win} = 1|RV_{win} = 1)]/[\Pr(PV_{win} = 0|RV_{win} = 1)]}{[\Pr(PV_{win} = 1|RV_{win} = 0)]/[\Pr(PV_{win} = 0|RV_{win} = 0)]} = \frac{542/135}{116/466} = \frac{622/156}{96/358} \approx 16.01,$$

where PV_{win} is a model-predicted vote for the winner and RV_{win} is a reported vote for the winner. The post-election survey exaggerates support for the winner, but no overreporting has occurred and both surveys predict winning side votes equally well—as revealed by equal odds ratios and conditional classification error rates.[7]

To illustrate the effect of overreporting, imagine now that a significant percentage of reported votes for the winner at t_2, say 10% (or around 78 votes), are overreports, as shown in Table 1c. The odds ratio, 8.49, is much lower, and conditional false negative classification error, 32.1%, much higher than in Table 1b. These differences are attributable to overreporting.

Extending our example to include covariates, suppose that the standard statistical model includes a variable positively related to voting for Candidate A, say membership in Party A, and that a greater proportion of Party A supporters takes the post-election poll than Party B supporters; that is, the marginal distribution of Party A supporters changes between t_1 and t_2. The post-election survey will overestimate support for Candidate A, even when the coefficient describing the effect of membership in Party A on support for Candidate A—and, consequently, the conditional false negative error classification rate—remains the same at t_1 and t_2.

Using t_1 model coefficients to predict t_2 vote preferences is critical to our method of detecting overreporting. Vote predictions at t_2 (and, therefore, the odds ratio in a cross-classification of model-predicted by self-reported vote preferences) are conditional on data and model parameters:

$$f(\theta|\alpha, \beta, x) = logit[\Pr(PV_{win} = 1|x)_i] = \alpha + \sum_1^K \beta_k x'_{ki} \ ,$$

where x'_{ki} is a $1 \times K$ vector of covariate values for individual i; β_k, a $K \times 1$ vector of coefficients associated with x; α, a (conditional) intercept, and all other notation is as above. The intercept α, in turn, may be decomposed into:

$$\alpha = \gamma_0 + \delta \ ,$$

where γ_0 is the conditional intercept and δ, a parameter that captures the effect of overreporting. If nobody overreported voting for the winner, δ is 0; with overreporting, δ will be positive, raising the false negative classification error rate and lowering the odds ratio. Predicting t_2 vote preferences by plugging t_2 data into the t_1 coefficients and intercept, in effect, allows dependent variables to reflect marginal changes in the independent variables while forcing all classification error into the overreporting parameter δ. Overreporting is then detectable as a higher conditional false negative rate at t_2.

The second prong of our analytic strategy for distinguishing social desirability and non-response bias is more straightforward, but requires pre- and post-election panel data. We compare re-interviewed respondents' reported pre-election vote preferences with those of pre-election respondents who dropped out of the panel. Non-response bias would imply higher re-interview rates among t_1 respondents who intended to vote for the (ultimately) winning candidate than among those who intended to vote for the (ultimately) losing candidate. On the other hand, more or less equal re-interview rates among the two groups are consistent with social desirability-induced overreporting for the winning side. The difference between survey-reported support for the winner and that actually obtained at the polls is not attributable to non-response bias, but to the "post-election bandwagon effect."

Of course, panel attrition occurs for reasons other than losing-side voters' turning down the follow-up interview, including low interest in politics, belonging to disadvantaged social and ethnic groups, and other factors (Groves and Couper 1998). So, we also model the decision to participate in the follow-up interview as a function of these factors as well as intended voice choice. This controls for potentially confounding variables, yielding cleaner estimates of pre-election vote preferences' effect on t_2 survey response. Both prongs of our research strategy afford evidence of non-response bias rather than social desirability-induced overreporting.

In light of our analytical strategy, these are formal statements of our hypotheses:

Non-Response Bias Hypothesis I: In a cross-classification of predicted by reported vote for both the pre- and post-election cross-sections, classification error produced by predicting a vote for the losing side-given a reported vote for the winning side—will be the same in the pre- and post-election samples:

$$\Pr(PV_{win} = 0|RV_{win} = 1, T = 2) = \Pr(PV_{win} = 0|RV_{win} = 1, T = 1) \ ,$$

where PV_{win} is a predicted vote for the winner, RV_{win} is a reported vote for the winner, and T is the survey period (1 = pre-election, 2 = post-election).

Social Desirability (Overreporting) Hypothesis I: In a cross-classification of predicted by reported vote for both the pre- and post-election cross-sections, classification error produced by predicting a vote for the winning side—given a reported vote for the losing side—will be higher in the post-election than in the pre-election sample:

$$\Pr(PV_{win} = 0|RV_{win} = 1, T = 2) > \Pr(PV_{win} = 0|RV_{win} = 1, T = 1) \ ,$$

where the notation is as before.

Non-Response Bias Hypothesis II: Re-interview rates will be higher for t_1 survey respondents who intended to vote for the (ultimately) winning candidate than for t_1 respondents who intended to vote for the (ultimately) losing candidate, *ceteris paribus*:

$$\Pr(R_{post} = 1|RV_{win,t_1} = 1, x) \text{ ¿} Pr(\mathrm{R}_{post} = 1|RV_{win,t_1} = 0, x) \ ,$$

where R_{post} is survey response in the post-election wave, RV_{win,t_1} is intent to vote for the winner declared in the pre-election wave, and x is a vector of covariates related to panel attrition.

Social Desirability (Overreporting) Hypothesis II: Re-interview rates will be the same, within sample error, for pre-election respondents who intended to vote for the (ultimately) winning candidate as for pre-election respondents who intended to vote for the (ultimately) losing candidate, *ceteris paribus*:

$$\Pr(R_{post} = 1|RV_{win,t_1} = 1, x) = \Pr(R_{post} = 1|RV_{win,t_1} = 0, x),$$

where the notation is as before.

Detecting evidence of the two remaining hypotheses, memory lapse and genuine late shifts in voter preferences, involves investigating the course of electoral preferences over the duration of the survey. If people forget who they voted for and then systematically misremember voting for the winner when they did not, we should observe a trend toward greater self-reported voting for the winning side after Election Day. Similarly, if late shifts of opinion affect survey overestimation of winner support, we should see a trend toward greater support for the winner across days leading up to the election, controlling for other attributes of survey respondents.

Memory Lapse Hypothesis: Support for the winning side of an election will increase over time *after* Election Day, *ceteris paribus*.

$$\Pr(RV_{win} = 1 | t > T) > \Pr(RV_{win} = 1 | t \leq T) \forall t > 0$$

where RV_{win} is a reported vote for the winner, t is the day on which the survey was taken, T is an arbitrarily fixed reference day, and 0 is Election Day.

Late Opinion Shift Hypothesis: Support for the winning side of an election will increase over time *before* Election Day, *ceteris paribus*. So,

$$\Pr(RV_{win} = 1 | t > T) > \Pr(RV_{win} = 1 | t \leq T) \forall t < 0,$$

where all notation is as before.

We distinguish among these hypotheses using two sources of data, an original survey designed to explore shifts in support for sides of California ballot propositions and support for presidential candidates expressed by respondents to the ANES.

2009 California Special Election

On May 19, 2009, California conducted a special election on a package of six ballot propositions—all originated by the state legislature—aimed at addressing a $23 billion fiscal deficit. Voters roundly rejected five of the proposals (Propositions 1A through 1E, which contained, collectively, a mixture of spending cuts and fee hikes) and overwhelmingly approved the sixth ballot proposal, Proposition 1F, forbidding legislative pay raises in years of budget deficits. The margins separating the winning "No" from losing "Yes" votes on the first five propositions were formidable, ranging from 24 to 33 percentage points. The difference for the sixth measure was nearly 49 percentage points. Wide though these margins of victory were, an original post-election survey inflated them further, by as much as 17%.

The survey research center at a Western public university fielded its 2009 California Special Election Survey between May 11-24, 2009, conducting interviews before and after the May 19 election as part of an experiential learning exercise in a course on public opinion. Survey participants constituted a simple random sample (SRS) drawn from California voting registration records. Between May 11 and 18, in the pre-election portion of the survey, 169 registered voters participated in the survey; the post-election portion, interviewed between May 20 and 24, comprised 107 respondents. We suspended data collection the day of the election. The survey asked about intended votes or, in the post-election poll, reported votes for four of the propositions, 1A, 1B, 1D, and 1F.

[Table 2 about here]

Table 2 shows the vote shares for the losing "Yes" side of Propositions 1A, 1B, and 1D estimated by the 2009 California Special Election Survey pre- and post-election cross-sections (columns labeled "Pre" and "Post" followed by the number of respondents that answered the question) as well as actual election results (in the column "Actual") for Propositions 1A, 1B, and 1D.[8] We first note that the pre-election estimates are quite close to actual vote percentages, within sampling error, for two of the three propositions (1A and 1B, as shown in the columns "Actual-Pre" and "p^a"), though the estimate is off the mark for Proposition D. In all three cases, the post-election poll significantly overestimates support for the winning side, by 14.6% for Prop 1A, 12.2% for Prop 1B, and 17.3% for Prop 1D (columns "Actual-Post" and "p^b").

Tests and results

To test the first *Social Desirability* and *Non-Response Bias* hypotheses, we first run three separate logistic regressions of votes for (1 = "Yes") or against (0 = "No") Propositions 1A, 1B, and 1D on a common set of explanatory variables for all pre-election (T=1) respondents. The explanatory variables are education, age, income, identification with the Democratic party, approval of the governor (on a scale of 1 to 10), approval of the legislature (also 1 to 10), agreement with budgeting through the citizen initiative process (yes/no), the county-wide percentage of citizens voting for the proposition, and reported votes (or intended votes) on the other two propositions.

Summarizing, model fit is quite good for Props 1A and 1B (Pseudo-R^2 of .40 for both) and reasonable for Prop 1D (.19), notwithstanding the fact that we have complete data for only 89 respondents.[9] Votes on one or more of the other propositions proved predictive of vote preferences on Props 1A, 1B, and 1D, indicating that voters tended to vote the propositions up or down as a block. Those who approved of the legislature's performance were likelier to vote for Props 1A and 1D, but those who agreed that citizens should vote directly on budget matters were less likely to vote for Proposition 1A. Democrats tended to vote for Prop 1B more than other citizens-hardly surprising, given teachers' unions endorsement of provisions that made up for shortfalls in spending on schools-but the likelihood of voting for this proposition decreased with age. Finally, greater educational attainment redounded in higher support for Prop 1A. In sum, many of these coefficients' predictive power, the models' overall fit, and the general closeness of pre-election survey estimates to final vote tallies all militate in favor taking the t_1 models as reasonable points of departure for predicting t_2 vote preferences.

We next combine coefficients from these models with data from the post-election (t_2) respondents to predict linear scores for each respondent. Using the logit transformation, we recover individual-level probabilities of a "Yes" vote and parlay these into a predicted binary vote preference equal to "1" if $p > 0.5$ and "0" otherwise. We then generate for each of the propositions two cross-classifications of predicted versus reported votes, one for the t_1 sample and the second for the t_2 sample, shown in Tables 3a-f.

[Tables 3a-f about here]

The tables afford no evidence of socially desirable overreporting, instead lending support to the non-response bias hypothesis. Cell values are row percentages, reflecting the probability of predicted vote preferences conditional on reported votes. In each table, the lower left-hand cell is the "conditional false negative classification error" rate—that is, respondents for whom a (losing) "Yes" vote is predicted are represented as a percentage of respondents who reported a vote for the (winning) "No" side. In two cases, Props 1B and 1D, the false negative rate for the post-election sample exceeded that of the pre-election sample. False negative classification error for Prop 1A is slightly higher at t_2 than t_1, but the difference is razor thin—nine-tenths of a percentage point (8.7% - 7.8%). Even if this difference reflected overreporting rather than sampling error, overreporting would still account for only a fraction, around 6%, of the 14.6-point difference between post-election survey estimates and actual voting for Prop 1A.

We acknowledge that our sample size is small. To help allay concerns that our results may be contingent upon the peculiarities of a small sample, we combine multiple imputation (MI) and bootstrap procedures to conduct supplemental analyses. Item missingness is particularly high for income, which was unobserved for 51 of the 169 pre-election, and 27 of 107 post-election, respondents. For each missing income observation, we impute ten values drawn from a conditional Gamma distribution where the shape parameter is estimated from a (GLM)

Gamma regression of income on covariates (education, gender, age, political knowledge, and agreement with citizen initiatives on budget matters) for all observations with complete data. Individual-specific scale parameters are obtained by predicting conditional mean incomes from the Gamma regression coefficients and known values of covariates, and then dividing predicted incomes by the shape parameter. Multiple imputation yielded 143 respondents at t_1 and 70 at t_2 (the remainder had missing values on other variables).

For its part, bootstrap estimation entails 1) drawing 2,000 samples with replacement of both the pre- and post-election cross-sections, where each sample included the imputed income values; 2) iteratively estimating the same logit model for the bootstrapped pre-election samples; 3) using the resulting t_1 model coefficients at each iteration (averaged across the multiply imputed datasets) to predict vote preferences for each of the t_2 bootstrapped samples; and 4) calculating false negative rates (conditional on a reported vote for the winning side) at t_1 and t_2.[10] Figures 1a-c show the results of the bootstrap iterations for Props 1A, 1B, and 1D.

[Figures 1a-c about here]

In each figure, the dark gray bars show the frequencies of t_1 classification error rates; the white bars, frequencies of t_2 classification error rates (more variable than the t_1 rates since the t_2 bootstrap samples were smaller); and the light gray region represents the overlap between the two time periods. For all three propositions, most of the probability mass lies in the overlapping region, suggesting that classification error rates are indistinguishable between the pre- and post-election cross-sections. The numbers bear this intuition out. For each of the propositions at t_1 and t_2, the medians, followed by the 95% confidence interval (the outer two numbers) and inter-quartile range (inner two numbers) in parentheses, are:

Prop 1A (t_1) – 4.9% (1.3%, 3.5%, 6.5%, 9.8%)
Prop 1A (t_2) – 6.6% (0.2%, 4.0%, 9.7%, 18.5%)

Prop 1B (t_1) – 7.7% (3.3%, 5.9%, 9.6%, 13.5%)
Prop 1B (t_2) – 7.5% (0.5%, 4.6%, 10.9%, 21.5%)

Prop 1D (t_1) – 5.7% (1.2%, 3.9%, 7.6%, 11.9%)
Prop 1D (t_2) – 5.7% (0.0%, 3.0%, 9.3%, 19.3%).

The t_1 and t_2 classification error rates are, for all intents and purposes, the same for Props 1B and 1D. For Prop 1A, classification error may be higher at t_2 than t_1 (consistent with our single-sample analysis above): the t_2 median, 6.6%, is slightly above the 75^{th} percentile of the t_1 distribution, 6.5%. Again, though, taking the numbers at face value, the 1.7-point difference between the t_2 and t_1 medians accounts for under 12% of the 14.6-point vote share overestimation observed for the winning side of Prop 1A.

Turning to the *Memory Lapse* and *Late Opinion Shift* hypotheses, the essence of our analytical strategy lies in modeling support for Propositions 1A, 1B, and 1D over time as a function of a counter for the day of interview, starting with May 11, 2009 (-8, for eight days before Election Day), running through May 24, 2009 (five days after Election Day). We interact this day counter with a dichotomous indicator for post-election respondents, so that the day counter represents the effect of time for the pre-election respondents (i.e., when the coefficient for the $Day \times t_2$ interaction is constrained to 0). This allows us to examine a different slope for the relationship between time and referendum vote in the pre-election and post election periods. Given that each of these propositions lost, we should see negative relationships between day of interview and support for each proposition before the Election Day (indicating *Late Opinion Shift*) and after Election Day (indicating *Memory Lapse*).[11]

Figures 2a-c graph the predicted probability of support for each of these three ballot measures over time, computed using Clarify (Tomz, Wittenberg, and King 2001). The graph and the regression results (not presented here) reveal little evidence of *Late Opinion Shift* or

Memory Lapse. Negative slopes for the day counter and the $Day \times t_2$ interaction would indicate shifts toward the winning side ("No") before the election and biases in recalled votes for the winner, respectively. There is no meaningful trend in pre-election support for Proposition 1B or 1D, and the model for Proposition 1A actually identifies a shift toward the *losing* side of the ballot referendum, with increased support for Proposition 1A closer to the election (the only statistically significant, time-oriented pre-election finding across these three models). Similarly, votes in support of Propositions 1A and 1B are flat across the days following the election-though Figure 2c does show a suggestive pattern in post-election support for Proposition 1D. Post-election support is signed in the direction anticipated by our *Memory Lapse* hypothesis: over time, people are more likely to say they voted against Proposition 1D. This shift does not reach conventional levels of statistical significance, although that might also be due to lack of power (i.e., limited post-election observations).

[Figures 2a-c about here]

Summing up, this evidence from the 2009 California Special Election Survey lends support to *Non-Response Bias Hypothesis I* over *Social Desirability Hypothesis I*. Little, if any, overreporting appears to have taken place. Some may have occurred on Prop 1A, but this accounts for a small fraction of winning vote share overestimation. The data also affords little evidence of *Memory Lapse* and none for *Late Opinion Shift*. However, given the sample size of the survey and the novelty of these elections, we expand our consideration of these hypotheses to traditional candidate elections and the ANES in the next section.

U.S. Presidential Elections, 1952-2008

The American National Election Studies have been carried out every U.S. presidential election since 1948. As Table 4 (seventh column) shows, in nine of the 16 presidential contests since 1952[12] the ANES overestimated the vote share of the winning candidate. Overestimation averaged 1.44% points over all 16 elections, and 3.06% taking into account just the elections that overestimated the winners' vote tally-outside the 1.4% margin of error for the smallest ANES sample of 1,212 in 2004. In contrast, underestimation, which averages -1.00%, can likely be chalked up to sampling error.

[Table 4 about here]

Test and results

We test our *Social Desirability I* and *Non-Response Bias I* hypotheses on the nine elections in which winner bias obtained—1952, 1956, 1964, 1968, 1972, 1980, 1992, 1996, and 2008 (see Table 4)—excluding years that underestimate winning vote share as irrelevant to a study on *overestimation*, and because underestimation appears to reflect random fluctuation of samples around the true vote share. The pre-post panel design of the ANES presented a challenge absent from the cross-sectional pre- and post-election samples in the California data. Potential "consistency bias" (respondents' tendency to remember and give the same answers they gave in previous waves) could artificially deflate post-election vote overreporting for the winner, stacking the deck against the overreporting hypothesis. So, we emulated pre- and post-election cross-sections by randomly dividing ANES respondents into two halves, modeling t_1 vote intention on one half and using t_1 model coefficients to predict t_2 vote choice on the other half. We repeated this process 1,000 times to ensure that our results do not depend on which observations were selected into each sample half.

Developing a model of winning candidate support for all presidential elections since 1952 presents a challenge for at least two reasons. First, virtually no attitudinal or behavioral variables that might explain winning candidate support (such as performance ratings, economic evaluations, etc.) are available for the entire time series.[13] Second, for nearly all variables measured at both t_1 and t_2, the ANES cumulative data file reports only the t_1 measure. Given these limitations, our independent variables are a winning party identification dummy variable equal to 1 if the respondent identifies with the same party as the winning candidate and 0 otherwise; a Republican Party identification dummy variable (1 = Republican, 0 = Other); African American ethnicity; a residual, "Other" ethnicity category; age; sex (1 = Male, 0 = Female); education (four categories treated linearly); and family income (five categories representing percentile ranges). We anticipate these characteristics will be associated with support for Republicans (age, education, and income) or Democrats (minority ethnicity).

We interacted each demographic variable with the Republican Party ID dummy. Thus, the explanatory variables' main effects correspond to non-Republican respondents, and interactive effects-the two component variables' main effects plus the interaction coefficient-to Republicans. Since Republicans won five of the nine contests considered here, we expect Republican identification will increase the likelihood of support for winning candidates, both alone and in combination with the socio-demographic variables. That is, demographic variables' effects on winning candidate support should be positive for Republican respondents and greaterthan for non-Republicans. Finally, to control for election-specific circumstances we include dummy variables for each election year (with 1952 as the reference category).

Winning party ID accounts for the lion's share of the model's explanatory power, and Republican Party ID (that is, when all other variables are equal to 0) is also strong and highly significant. Other significant predictors are African American ethnicity (positive for Republican African Americans, negative for non-Republicans), age (which had a negative effect for Republican respondents), education (positive effect for non-Republican respondents, negative for Republicans), and all the year dummies except for 1956. Model fit was reasonable: Pseudo-R^2 averaged .54 over the 1,000 half samples (about 8,800 respondents each).[14]

[Figure 3 about here]

Figure 3 depicts their substantive results, superimposing two histograms of, respectively, classification error at t_1 (dark gray bars) and at t_2 (white bars), with the overlapping area in light gray. The t_1 median is 20.2%, with a 95% confidence interval of (17.5%, 22.9%) and the t_2 median, 22.9%, with a 95% confidence interval of (20.5%, 25.7%). The nominal difference of 2.7 percentage points is greater than in the California study. It may be that the time-invariant predictors measured at t_1—the only ones available in the ANES cumulative data file—do a worse job predicting votes at t_2 than t_1. Still, overlap between t_1 and t_2 classification errors is large (2.5 percentage points between the t_2 confidence interval lower bound and the t_1 upper bound), and the t_2 median is below the t_1 upper bound, taking into account rounding error.

We also subject the *Social Desirability* and *Non-Response Bias* hypotheses to the second prong of our analytic strategy: comparing re-interview rates. The last three columns of Table 4 report, respectively, the percentage of pre-election supporters of winning party candidates who took the post-election poll, the percentage of losing party candidates who did so, and the difference between the two. Since 1952, post-election survey response rates were, on average, 0.9 points higher for pre-election supporters of winning candidates than for supporters of losing candidates; electoral losers' response probability was actually higher than winners' on five occasions (four of which were Democratic victories). We compare re-interview rates for the same subset of election years used in the Classification Error Comparison Method, minus the 1956 contest. The 100% panel retention rate that year can shed no light on differences in re-interview rates between supporters of winning and losing candidates.

[Table 5 about here]

Table 5 presents the results of three logistic models of post-election survey response on pre-election support for winning candidates. In Model 1, the coefficient for winning candidate support, the only explanatory variable in the model, is ($\beta = .25, p = .000.$). The model-predicted probability that a respondent who intended to vote, in the pre-election survey, for the ultimate winner is .913. It is .891 for respondents who intended to vote for a losing candidate, for a difference of 2.2 percentage points. To ensure that this difference is not spurious or attributable to omitted variables, we progressively control for potentially confounding variables by adding election year dummy variables to Model 2 and the year dummies plus a full complement of explanatory variables (African American ethnicity, other, non-white ethnicity, age, education, income, sex, interest in politics, and party ID) to Model 3. The effect of winning candidate support on post-election survey response (Model 2 $\beta = .25, p = .000$, Model 3 $\beta = .23, p = .000$) is robust to the addition of control variables. Predicted probabilities of survey response are .897 for "losers" and .919 for "winners" (difference of 2.2 points) under Model 2, and .903 and .922, respectively (difference of 1.9 points) under Model 3, with all differences significant at $p < .001$. The results of these tests, then, suggest that non-response bias accounts for 60% to 70% of the 3.06% winning vote share overestimation.

Finally, we examine the *Late Opinion Shift* and *Memory Lapse* hypotheses, regressing pre- and post-election reports of winning candidate support on a day-of-interview counter and the explanatory variables used to estimate classification error. We find no evidence for either. The pre-election counter's coefficient is indistinguishable from 0 (β_{pre} = -.002, p = .170) and, on average, support shifts away from winning candidates after the election (β_{post} = -.005, p = .007).

Discussion

The problem of post-election polls' overestimating winning vote share pervades survey research. Existing explanations for overestimation center on psychological factors that cause survey respondents to misreport their votes. Respondents are seen either remembering their votes inaccurately or, because they wish to present themselves as having engaged in socially sanctioned behavior, deliberately misrepresent their vote. We propose—and find evidence for—an alternative hypothesis: voters for the losing side may not lie about how they voted, but rather choose not to participate in a post-election survey in the first place. Despite a plethora of research on survey non-response and the reasons for it, scholars have not taken it into account in explaining overestimation of winning vote share. We find no evidence (in the elections we consider, at any rate) of two other possible explanations, late opinion shift and memory lapse.

Our findings suggest that survey researchers need to revise our understanding of survey psychology and respondents' motivations for participating in surveys. We know from vote validation turnout studies that survey participants will prevaricate when responding truthfully is embarrassing. Scholars assumed that this explanation also accounted for survey overestimation of winning candidates' vote shares: respondents would say they voted for the winner because they are ashamed to say they voted for the loser. This study, however, raises the possibility that overestimation occurs not because voters find it awkward to admit they voted for the losing side, but because they are simply less interested in taking a survey in the first place.

That survey respondents would lie about having voted more than they would about having voted for the winner stands to reason. Voting is a civic duty, but voting for a winning candidate is not. Greater shame probably attaches not having voted than to having cast a ballot for a candidate who ultimately lost. Thus, incentives to overreport having voted may be stronger than the incentives to overreport voting for winning candidates. This may be part of the reason turnout overestimation rates are generally much higher than winning vote share overestimation rates. Simply put, there is greater reason to lie about having voted than about whom one voted for. If there is not much social pressure to fib about having voted for winning candidates, there would seem to be even less to do so over ballot initiatives, given the absence of personalities and the emotional reactions they engender.

Although we have called non-response bias and social desirability-induced overreporting alternative explanations of winning side vote share overestimation, they are perhaps better seen as complementary. The tests we carry out favor non-response bias over social desirability, suggesting that anywhere from half (in our re-interview comparison analysis of the ANES cumulative dataset analysis) to 90% (in our classification error analysis of Prop A of the California ballot initiatives) or more of overestimation may be attributable to non-response bias. Devising a means for assessing the relative contributions of social desirability and non-response to overestimation remains an area for further research. We intimate possibilities—including classification error as a percentage of total overestimation—but developing them is outside the scope of this paper.

Our study's main implication for political research is that we need to do a better job of representing losers (and subpopulations that supported losing candidates) in post-election survey samples—especially given the importance of survey research in shaping voters' and political elites' perceptions of the public will. The fault for winner vote overestimation may not lie with deceitful respondents, but with survey recruitment techniques. Rather than unjustly excoriating survey respondents for giving dishonest but socially desirable responses, then, survey researchers turn a critical eye to the problem of making potential respondents who have suffered recent political loss feel more welcome in political discourse.

Notes

[1]We prefer the term vote "overestimation" to the more widely used "overreporting." "Overreporting" implies that respondents supply inaccurate information when taking a survey. But overestimation can also result from overrepresentation of citizens who voted for the winner in the sample. We therefore use "overestimation" in the general case where survey estimates are higher than actual vote share, and "overreporting" in the specific case where overestimation results from respondents' inaccurately claiming to have voted for the winner.

[2]Though this study focuses on vote share overestimation, we refer frequently to studies on turnout overestimation (in which survey-reported turnout rates exceed actual turnout). The psychological underpinnings of overreporting and survey participation are largely the same in both cases.

[3]See Brehm (1999) and Heckman (1979) for sample selection models and Raghunathan (2004) for sample reweighting techniques.

[4]Some studies argue overreporting is partly an artifact of the survey instrument. Eubank and Gow (1983) point to question order effects. The ANES's placement of the question on vote choice after questions that identified incumbents slanted responses in favor of more recognized, incumbent candidates. Jacobson and Rivers (1993) argue that the phrasing of questions on vote choice induces respondents to report voting for winning candidates and incumbents. The switch in 1978 from an open-ended question about whom respondents voted for to one that provided name and party cues boosted reported vote shares for incumbents, likely due to the higher name recognition enjoyed by winners.

[5]Studies have also established that misreporting increases the later the post-election survey is taken after the election. Using validated voting data in an Oregon study, Belli and his coauthors found that an experimental question wording designed to prod respondents' memories increased reliability of self-reported voting data for surveys carried out later in the

data collection period. The authors inferred that misreporting increased the more time had elapsed between the election and the survey (1999: 99). Atkeson (1999) noted that memory failure was an especially important explanation for vote misreporting given the large number of days between the primary election and the administration of the ANES, taken after the general election.

[6]Because of space limitations, we omit our analysis of the 1996 U.S. presidential election here, instead making it available as an on-line appendix available via the Editorial Manager site for the *American Journal of Political Science*, *http://www.editorialmanager.com/ajps*. The on-line appendix reports findings from the 1996 election entirely consistent with the findings presented here.

[7]We use *conditional* classification error rates rather than *overall* classification error rates (or incorrectly predicted votes as a percentage of total sample size). Although conditional classification error is invariant to changes in marginal proportions, overall classification error is not.

[8]Since the survey *underestimated* support for the winning "yes" side on Proposition 1F (estimated post-election vote share of 64.5% compared to 74.3% actual vote), we omit this proposition as irrelevant to our analysis.

[9]We omit this table of results (and some others, as noted) in the interest of space, but are happy to furnish it to interested readers. Coefficients are flagged as significant at $p < .10$ (one-tailed tests) at a 95% confidence level.

[10]We discard iterations in which R's "glm" algorithm estimated probabilities of 0 or 1 or failed to converge (notwithstanding, the "glm" library provides parameter estimates in these cases anyway), so the effective number of iterations is smaller—as low as 1,372 in the case of Prop 1A, but much higher for the other two propositions.

[11]The models also control for other correlates of vote preferences on the propositions, including party identification, approval of Governor Schwarzenegger's and the California

Legislature's job performance, each voter's electoral context (the percentage of votes for each proposition in the respondent's county), political knowledge (using a measure of the respondent's self-reported knowledge of the party identification of her representative in the California Assembly, and respondent age.

[12]We omit the 1948 study, in which the pre-election poll contained only seven questions about the coming election.

[13]In the on-line appendix, we explore the influence of retrospective pocketbook voting using the 1996 as a case study. The results substantively echo the results from the ANES cumulative file.

[14]An example of these statistical models are presented in the on-line appendix.

References

Atkeson, Lonna Rae. 1999. "'Sure I Voted for the Winner!' Overreport of the Primary Vote for the Party Nominee in the National Elections," *Political Behavior*, 21(3):197-215.

Beasley, Ryan K., and Mark R. Joslyn. 2001. "Cognitive Dissonance and Post-Decision Attitude Change in Six Presidential Elections," *Political Psychology* 22(3):521-540.

Belli, Robert F., Michael W. Traugott, Margaret Young, and Katherine A. McGonagle. 1999. "Reducing Vote Overreporting in Surveys: Social Desirability, Memory Failure, and Source Monitoring," *Public Opinion Quarterly*, 63(1):90-108.

Berinsky, Adam. 2004. *Silent Voices: Public Opinion and Political Participation in America.* Princeton NJ: Princeton University Press.

Brehm, John. 1993. *The Phantom Respondents. Opinion Surveys and Political Representation.* Ann Arbor: University of Michigan Press.

Brehm, John. 1999. "Alternative Corrections for Sample Truncation: Applications to the 1988, 1990, and 1992 Senate Election Studies," *Political Analysis*, 8(2):183-199.

Cochran, William G. 1977. *Sampling Techniques.* New York: John Wiley Sons.

Crespi, Irving J. 1988. *Pre-election Polling. Sources of Accuracy and Error.* New York: Russell Sage Foundation.

Deufel, Benjamin J., and Orit Kedar. 2000. "But Did They Really Vote? A Method to Correct for Overreporting of Turnout." Presented at the 2000 Annual Meeting of the American Political Science Association, Washington, D.C.

Ginsberg, Benjamin, and Robert Weissberg. 1978. "Elections and the Mobilization of Popular Support.," *American Journal of Political Science*, 22(1):31-55.

Groves, Robert M., and Mick P. Couper. 1998. *Nonresponse in Household Interview Surveys.* New York: John Wiley Sons.

Groves, Robert M., Don A. Dillman, John L. Eltinge, and Roderick J.A. Little, eds. 2002. *Survey Nonresponse*. New York: John Wiley Sons.

Harrop, Martin, and William Miller. 1987. *Elections and Voters: A Comparative Introduction*. New Amsterdam Books.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error," *Econometrica*, 47(1):153-161.

Herbst, Susan. 1998. *Reading Public Opinion. How Political Actors View the Democratic Process*. Chicago: University of Chicago Press.

Johnson, Marcia K., Shahin Hashtroudi, and D. Stephen Lindsay. 1993. "Source Monitoring," *Psychological Bulletin*, 114:3-28.

Little, Roderick J.A., and Donald B. Rubin, 2002. *Statistical Analysis with Missing Data*, 2^{nd} ed. New York: John Wiley Sons.

Kish, Leslie. 1965. *Survey Sampling*. New York: John Wiley Sons.

Noelle-Neumann, Elisabeth. 1993. *The Spiral of Silence: Public Opinion-Our Social Skin*. Chicago, IL: University of Chicago Press.

Powers, Daniel, and Yu Xie. 2000. *Statistical Methods for Categorical Data Analysis*. San Diego, CA: Academic Press.

Presser, Stanley. 1990. "Can Context Changes Reduce Vote Overreporting?" *Public Opinion Quarterly*, 54:586-593.

Price, Vincent. 1999. *Public Opinion*. Thousand Oaks: Sage Publications.

Raghunathan, Trivillore E. 2004. "What Do We Do with Missing Data? Some Options for Analysis of Incomplete Data," *Annual Review of Public Health*, 25:99-117.

Silver, Brian D., Barbara A. Anderson, and Paul R. Abramson. 1986. "Who Overreports Voting?" *American Political Science Review*, 80(2):613-624.

Southwell, Priscilla L. 1986. "The Politics of Disgruntlement: Nonvoting and Defection among Supporters of Nomination Losers, 1964-1984," *Political Behavior*, 8:81-95.

Thomsen, Arnold. 1938. "What Voters Think of Candidates Before and After Election," *Public Opinion Quarterly*, 2(2):269-274.

Tomz, Michael, Jason Wittenberg, and Gary King. 2001 CLARIFY: Software for Interpreting and Presenting Statistical Results (Version 2.0). Cambridge, MA: Harvard University.

Wolfinger, Raymond, and Steven J. Rosenstone. 1980. *Who Votes?* New Haven, CT: Yale University Press.

Wright, Gerald C. 1990. "Misreports of Vote Choice in the 1988 NES Senate Election Study," *Legislative Studies Quarterly*, 15(4):543-563.

Table 1: Hypothetical Cross-Classifications.

a. Hypothetical t_1 respondents by self-reported vs. model-predicted vote intention

		Predicted Vote Intention Loser	 Winner	
Reported Vote Intention	Loser	Correctly Predicted Vote for Loser 466 (80.1%)	Random False Positive Classification Error 116 (19.9%)	582 (46%)
	Winner	Random False Negative Classification Error 136 (20.1%)	Correctly Predicted Vote for Winner 542 (79.9%)	678 (54%)
		602 (47.8%)	658 (52.2%)	$N = 1,260$

b. Hypothetical t_2 respondents by self-reported vs. model-predicted vote preference
(no overreporting)

		Predicted Vote Preference Loser	 Winner	
Reported Vote Preference	Loser	Correctly Predicted Vote for Loser 386 (80.1%)	False Positive Random + Underreporting 96 (19.9%)	482 (38.3%)
	Winner	False Negative Random + Overreporting 156 (20.1%)	Correctly Predicted Vote for Winner 622 (79.9%)	778 (61.7%)
		542 (43%)	718 (57%)	$N = 1,260$

c. Hypothetical t_2 respondents by self-reported vs. model-predicted
(with overreporting)

		Predicted Vote Preference Loser	 Winner	
Reported Vote Preference	Loser	Correctly Predicted Vote for Loser 386 (80.1%)	False Positive Random + Underreporting 96 (19.9%)	482 (38.3%)
	Winner	False Negative Random + Overreporting 250 (32.1%)	Correctly Predicted Vote for Winner 528 (67.9%)	778 (61.7%)
		636 (50.5%)	624 (49.5%)	$N = 1,260$

Table 2: Comparison of Pre- and Post-Election Surveys and Actual Voting for Three Ballot Measures in the May 2009 California Special Election (Cells are percentage that voted "Yes" on each proposition).

Measure	*Pre*	N_{Pre}	*Post*	N_{Post}	*Actual*	*Actual–Pre*	p^a	*Actual–Post*	p^b
Prop. 1A	31.2%	154	20.0%	80	34.6%	3.4%	0.181	14.6%	.001
Prop. 1B	34.4%	151	25.9%	81	38.1%	3.7%	0.169	12.2%	.008
Prop. 1D	23.2%	151	16.3%	80	34.0%	10.8%	0.001	17.7%	.000

[a] independent samples t-test for one-sided difference of means between pre-election and actual vote share.

[b] single-sample t-test for one-sided difference of means between post-election and actual vote share.

Table 3: Cross-Classification of Predicted and Reported Vote, Before and After 2009 California Special Election.

		Pre-Election Sample: Predicted Vote				Post-Election Sample: Predicted Vote	
		Yes (Lose)	No (Win)			Yes (Lose)	No (Win)
Prop. 1A							
Intended Vote	Yes (Lose)	68% (17)	32% (8)	**Reported Vote**	Yes (Lose)	87.5% (7)	12.5% (1)
	No (Win)	7.8% (5)	92.2% (59)		No (Win)	8.7% (4)	92.3% (42)
		$N = 89$				$N = 54$	
Prop. 1B							
Intended Vote	Yes (Lose)	67% (22)	33% (11)	**Reported Vote**	Yes (Lose)	67% (8)	33% (4)
	No (Win)	12.5% (7)	87.5% (49)		No (Win)	7.1% (3)	92.9% (39)
		$N = 89$				$N = 54$	
Prop. 1D							
Intended Vote	Yes (Lose)	67% (14)	33% (7)	**Reported Vote**	Yes (Lose)	75% (6)	25% (2)
	No (Win)	7.4% (5)	92.6% (63)		No (Win)	0% (0)	100% (46)
		$N = 89$				$N = 54$	

Table 4: Winning Vote Share, Overestimation, and Participation in Post- Election Polls by Pre-Election Vote Preference, ANES 1952-2008.

Year	*Winner*	*Party*	*Vote%*	*Victory Margin*	*ANES%*	*Overest.*	*%Winners Post-Elec.*	*%Losers Post-Elec.*	*Diff. Win-Lose*
1952	Eisenhower	R	55.2%	10.9%	57.9%	2.74%	90.9	91.0%	-0.11%
1956	Eisenhower	R	57.4%	15.4%	59.5%	2.08%	100.0	100.0%	0.00%
1960	Kennedy	D	49.7%	0.2%	49.3%	-0.38%	92.9	96.1%	-3.25%
1964	Johnson	D	61.1%	22.6%	67.4%	6.34%	94.4	92.6%	1.74%
1968	Nixon	R	43.4%	0.7%	47.6%	4.20%	88.7	86.7%	2.02%
1972	Nixon	R	60.7%	23.2%	63.6%	2.94%	87.0	84.6%	2.46%
1976	Carter	D	50.1%	2.1%	49.7%	-0.38%	84.5	88.2%	-3.76%
1980	Reagan	R	50.8%	9.7%	50.8%	0.07%	90.1	87.0%	3.11%
1984	Reagan	R	58.8%	18.2%	57.7%	-1.10%	90.9	89.4%	1.54%
1988	Bush	R	53.4%	7.7%	52.3%	-1.10%	90.1	88.2%	1.84%
1992	Clinton	D	43.0%	5.6%	47.6%	4.62%	90.8	90.9%	-0.03%
1996	Clinton	D	49.2%	8.5%	52.9%	3.68%	90.7	90.1%	0.55%
2000	Bush	R	47.9%	-0.5%	45.5%	-2.38%	88.7	85.8%	2.99%
2004	Bush	R	50.7%	2.5%	50.1%	-0.67%	91.2	86.7%	4.50%
2008	Obama	D	52.9%	7.3%	53.7%	0.87%	91.8	91.9%	-0.13%

Table 5: Logistic Regression of ANES Cumulative File Post-Election Survey Response Probability on Pre-Election Vote Preference and Other Covariates

	Model 1	Model 2 w/year effects	Model 3 w/year effects
	β	β	β
	(s.e.)	**(s.e.)**	**(s.e.)**
Pre-election vote preference	0.25***	0.26**	0.23***
(1=Winner, 0=Other)	(0.05)	(0.05)	(0.06)
African American			-0.15
(1=Yes, 0=No)			(0.09)
Other non-white			-0.43***
(1=Yes, 0=No)			(0.29)
Age			0.00
			(0.00)
High school or less			0.10
(1=Yes, 0=No)			(0.09)
Some college			0.25*
(1=Yes, 0=No)			(0.11)
College or advanced degree			0.43**
(1=Yes, 0=No)			(0.12)
Income			-0.05
			(0.02)
Gender			-0.09
(1=Male, 0=Female)			(0.05)
Interest in politics			0.16
(1=Not much, 3=Very interested)			(0.04)***
Party identification			0.03*
(1=Strong Dem., 7=Strong Repub.)			(0.01)
Intercept	2.06***	2.83***	2.43***
	(0.03)	(0.89)	(0.17)
N	16,726	16,726	14,358
Model χ^2	23.94***	267.92***	310.22***
$Pseudo\,R^2$	0.002	0.025	0.034

Significance levels: *** $p < .001$ ** $p < .01$ * $p < 05$ † $p < .10$

Coefficients for year fixed effects in models 2 and 3 available from authors.

Figure 1: Histograms of Classification Error Rates for 2009 California Special Election Pre- and Post-Election Bootstrapped Samples.

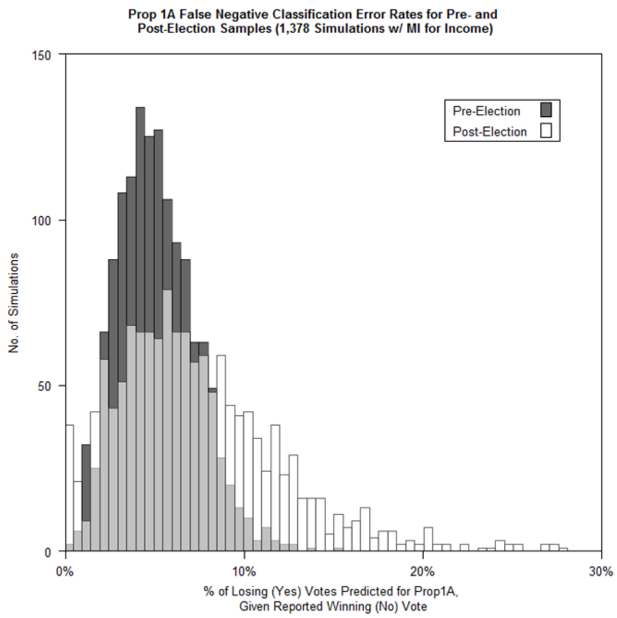

a. Prop. 1A

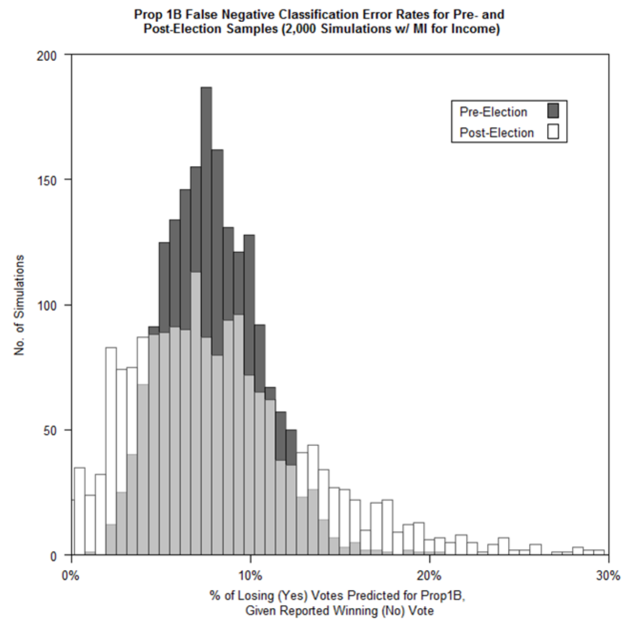

b. Prop. 1B

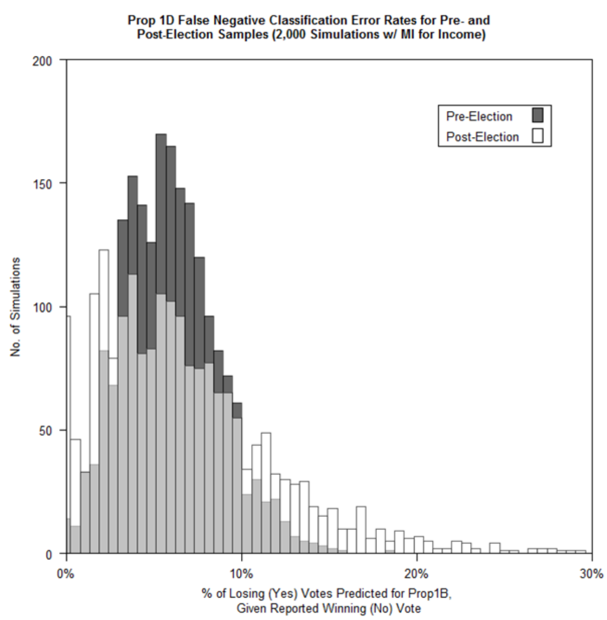

c. Prop. 1D

Figure 2: Time Trend in Probability of Respondent Support for 2009 California Special Election Propositions.

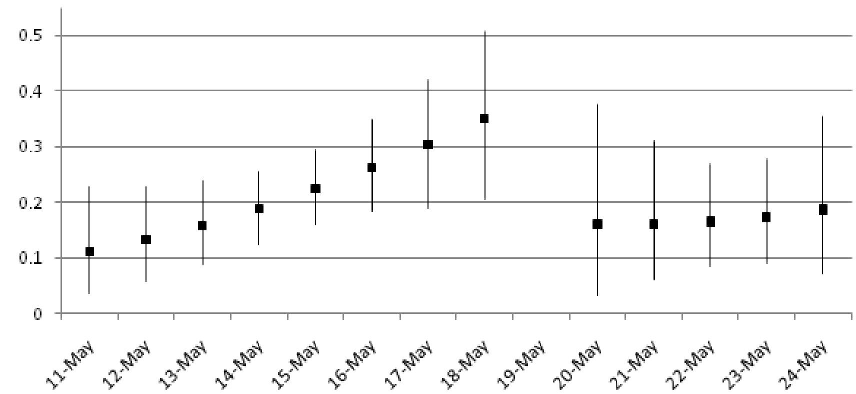

a. Prop. 1A

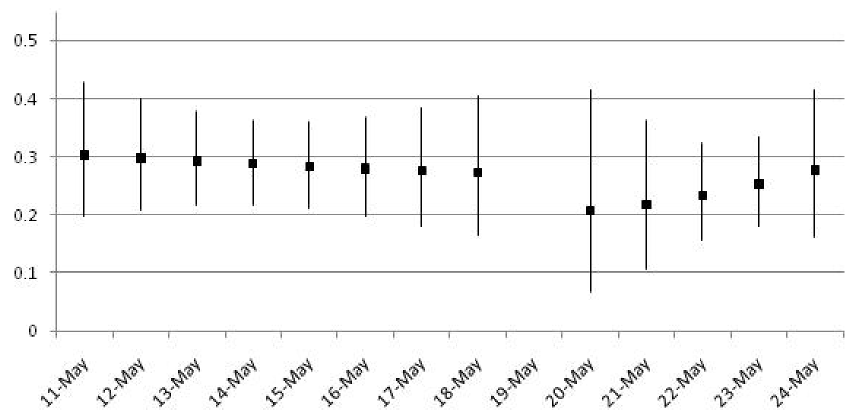

b. Prop. 1B

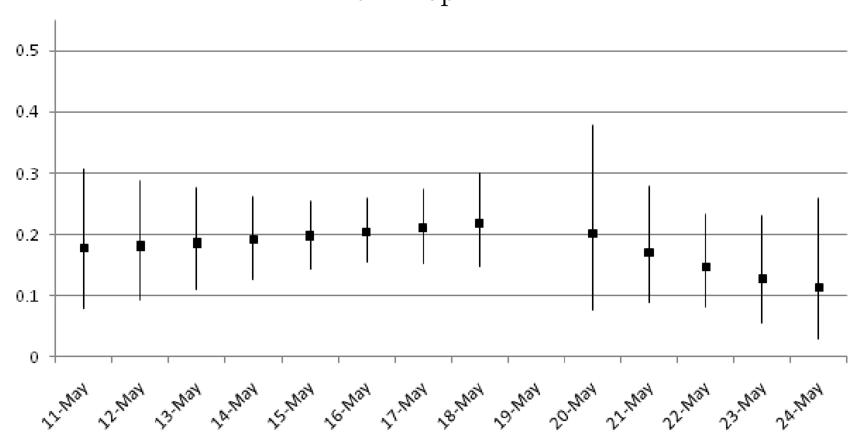

c. Prop. 1D

Figure 3: Histogram of Classification Error Rates for ANES 1952-2008 Simulated Pre- and Post-Election Cross-sections.