

GOOD EXCUSES: UNDERSTANDING WHO VOTES WITH AN IMPROVED TURNOUT QUESTION

BRIAN DUFF
MICHAEL J. HANMER
WON-HO PARK
ISMAIL K. WHITE

Abstract Utilizing studies which validate voter turnout, previous researchers have been able to identify a strong tendency for individuals to report voting when they in fact did not. In this article, we assess the effectiveness of a new turnout question on reducing voter over-reporting in the National Election Study. Providing respondents with socially acceptable excuses for not voting, we found that this alternate question significantly reduces the over-reporting of turnout in the 2002 National Election Study by about 8 percentage points. Moreover, our analysis reveals that with the new question wording, estimates of the turnout rate for those usually thought to be the least likely to vote are considerably lower than estimates using the traditional question. Thus, not only did the experiment work to significantly reduce over-reporting, the new question provides deeper insights into the voting behavior of the American electorate that has implications for both scholars and reformers.

For scholars interested in American elections, there is probably no more important question asked of American National Election Study (NES) respondents than the one that inquires if they voted in the previous election. Unfortunately, a portion of the people who did not vote in the previous

BRIAN DUFF is an assistant professor of political science at the University of New England. MICHAEL J. HANMER is an assistant professor of government at Georgetown University. WON-HO PARK is an assistant professor of political science at the University of Florida. ISMAIL K. WHITE is an assistant professor of government at the University of Texas-Austin. We greatly appreciate the encouragement and guidance of Nancy Burns and Don Kinder. Helpful comments on earlier versions were also provided by the members of the National Election Studies Board of Overseers, Allyson Holbrook, the editor, and the anonymous referees. We would also like to thank Mike Traugott, Matt Beckmann, and Dick Niemi for engaging in helpful discussions, and Diana Watral for research assistance. The authors gratefully acknowledge the support of the National Election Studies Research Fellowship program. All errors are our own. Author names are listed alphabetically. Address Correspondence to Ismail K. White; e-mail: whiteik@mail.la.utexas.edu.

election will report to the NES that they did. This over-report can be inferred by comparing the reported turnout of the NES sample to the turnout of the population at large. It has also been confirmed in various studies that have used voter rolls to verify the turnout of survey respondents.

The two factors most often identified as driving the over-reporting of voting are problems with the respondent's memory and the social desirability of saying one voted (particularly in the midst of answering questions about politics). The traditional NES turnout question attempts to mitigate these two factors with its introductory script¹ that acknowledges socially acceptable reasons for not voting. Respondents might also be reminded that they were especially busy this year.

In 2000, the NES introduced a new version of the turnout question, which all respondents received. This version had an identical introductory script, but offered the respondents a chance to seize onto one of the socially acceptable excuses in giving their response. After the introductory script, respondents were asked "Which of the following statements best describes you: 1. *I did not vote (in the election this November)*; 2. *I thought about voting this time but didn't*; 3. *I usually vote but didn't this time*; 4. *I am sure I voted.*" This new version seemed successful in reducing the size of the over-report of turnout in 2000, if only by a small percentage (table 1). But absent an experimental design using both the new and old questions, we cannot be certain if the new question accounted for the lower over-report in 2000. It was also difficult to determine whether the new version affected different social groups in different ways, and therefore might be introducing new bias into the study of voting. The 2002 NES offers an opportunity to investigate these questions because it used an experimental design to randomly assign half its respondents to the new version of the turnout question and half to the traditional version. We take advantage of the experimental design to assess the effectiveness of this new version of the turnout question on reducing the rate of over-reporting. The experimental design allows us to say more definitively that the new question does indeed mitigate the problem of overestimating the turnout rate, by about 8% ($p = 0.002$).²

We offer further analysis to determine if this new turnout question affects different groups of respondents in different ways. In particular, we investigate the extent to which the new version might systematically filter out certain groups that tend to over-report, and not others. When examined against the analysis of validated vote results presented by Belli, Traugott, and Beckmann (2001) our study provides insight into possible biases created by the new turnout question. Belli, Traugott and Beckmann found that over-reporters

1. The question reads: *In talking to people about elections, we often find that a lot of people were not able to vote because they weren't registered, they were sick, or they didn't have time. How about you—did you vote in the elections this November? 1. Yes, voted. 5. No, didn't vote.*

2. Tests confirmed that the random assignment procedure was implemented successfully.

Table 1. Official and NES Turnout Rates 1994–2002

Year	Official, based on VEP	NES Standard Question	NES New Question	NES Over-report, Standard Question	NES Over-report, New Question
1994	41.1%	55.7%	NA	14.6	NA
1996	51.8%	71.8%	NA	20.0	NA
1998	38.4%	52.1%	NA	13.7	NA
2000	54.3%	NA	72.1%	NA	17.7
2002	40.2%	64.9%	56.9%	24.7	16.6

NOTE.—The Official turnout rate uses total votes cast for the highest office as reported by the FEC.

Voting Eligible Population (VEP) data are from Michael McDonald's web site: <http://elections.gmu.edu/Turnout%201980-2004.xls>, visited 3/12/06.

All results from the NES are weighted.

NA indicates that the result is not applicable due to the absence of the question version for given years.

The 1996 NES suffers from problems with the sample (Burden 2000, 2003; McDonald 2003; Martinez 2003; and The National Election Studies 1996). We utilized the weight provided by the NES, V960005B, to adjust for sampling error.

tend to be similar to actual voters both demographically and attitudinally. While the NES' experimental design does not allow as straightforward an examination of the question of whether the decline in over-reporting occurred among those who are most similar to actual voters, we examine the data for indications that this is the case and find none.

The most striking finding that emerges from an analysis of the turnout wording experiment included in the 2002 NES is that the traditional turnout question masks the very low levels at which the poorest, least knowledgeable, and least self-regarding in terms of political efficacy turn out to vote in a midterm election. An analysis of the new turnout question indicates that these groups of citizens are left out of the electoral process to an extent that could not be uncovered using the traditional turnout question.

Vote Over-reporting

The over-reporting of voting in surveys is a well-recognized problem among survey researchers and those interested in understanding voting behavior. Depending upon how one measures both actual turnout and misreporting (Silver, Anderson, and Abramson 1986; Martinez 2003; McDonald 2003), survey over-reporting tends to range from 8 to 14 percentage points (Belli, Traugott, and Beckmann 2001) and by some measures has exceeded 20 points in recent elections (Burden 2000; Martinez 2003; McDonald 2003; and table 1). Such high levels of over-reporting can have important consequences

for research attempting to explain voter turnout, especially if there is something systematic about who is likely to over-report.

Attempts to understand better the source of this problem have resulted in two distinct approaches to assessing vote over-reporting: voter validation studies and survey experiments. While each approach addresses the questions of which respondents tend to over-report and why they do so, validation studies have met with much more success than the survey experiments in identifying who over-reports. Yet, survey question experiments hold promise for building the over-report solution into the survey instrument itself, and avoiding the expensive and sometimes inaccurate validation process (Presser, Traugott, and Traugott 1990).

VALIDATION STUDIES

By comparing respondents' self-reported turnout in surveys to official voting records, researchers have sought to assess the accuracy of responses to turnout questions (Clausen 1968; Traugott and Katosh 1979). Although these validation studies have proven to be a useful tool in identifying both the extent of vote over-reporting and its demographic and attitudinal correlates, there is a fair amount of variability in the results across studies, deriving from differences in approaches to measuring vote over-reporting.³ While these studies have produced a long and varied list of attitudinal and social determinants of over-reporting, the race of the respondent stands out as the most consistent predictor of vote over-reporting. Validation studies indicate that African Americans are anywhere from two to three times more likely to over-report voting than white Americans (Abramson and Claggett 1984). The most common explanation for this finding is that the historical struggle for the right to vote is so salient to black respondents that they experience a higher than usual internal pressure to report having voted. Race aside, over-reporters tend to be more similar to validated voters than to non-voters in their attitudes and beliefs about politics (Belli, Traugott, and Beckmann 2001). These results provide some evidence for the idea that over-reporters are those individuals who face the strongest internal and societal pressures to produce socially desirable responses. That is, blacks, and those who care more about politics are more likely to be aware of the norm of voting and see voting as tool for social change.

Of course, validation studies fall short of solving the over-report problem within the survey context. The expense of validating votes—in both time and money—provides incentive to seek alternative means of addressing survey respondents' tendency to inform researchers that they voted when they did not. Moreover, there is evidence of complications with validation studies

3. For more on this point see Appendix table A1 in the on-line version of the article. For an overall review, see Belli, Traugott and Beckmann (2001).

stemming from reliance on official state records that may be unreliable in a troubling, systematic way (Presser, Traugott, and Traugott 1990).

EXPERIMENTAL STUDIES

As an alternative to validation studies, researchers have sought ways to reduce over-reporting by manipulating the survey itself. These studies have explored variations in the survey conditions and the wording and answer choices of the turnout question meant to reduce over-reporting by targeting the respondents' likely incentive to provide socially desirable responses or ability to recall accurately their behavior.

Such experiments, however, have not been entirely successful at reducing over-reporting. Presser (1990) found that manipulating whether or not the respondent was asked to report earlier instances of political participation did not reduce over-reporting, nor did preceding the turnout question with an item asking information that only a person who voted would know. Abelson, Loftus, and Greenwald (1992) report the results of three survey-based experiments, each of which fails to find significant results. Perhaps this failure was due to their decision to limit their analysis to citizens whose registration could be validated. That is, by overcoming the hurdle of registration (Wolfinger and Rosenstone 1980) the citizens in their studies have performed at least one political act, thus suggesting some attachment to the political system and an awareness of its norms. As we discuss later, this explanation is supported by our results. Belli, Traugott, and Rosenstone (1994) found that encouraging respondents to scrutinize the sources of their memories also had little effect on vote over-reporting. Operating under the idea that over-reporting is the result of confusion among respondents about the election in which they actually voted, these scholars found that exposing respondents to either the traditional turnout question or a version of the question which encouraged them to think harder about having voted, along with additional answer choices (the same as used in the 2000 and 2002 NES) resulted in no statistically significant differences in vote over-reporting. Again, we believe Belli, Traugott, and Rosenstone's (1994) reliance on validated registrants may have limited their ability to reduce over-reporting.⁴ Holbrook and Krosnick (2003) utilized experimental techniques designed to reduce social desirability pressures by assuring respondents of the anonymity of their responses, but also failed to find evidence of reduced vote over-reporting.⁵

4. In fact, Belli, Traugott and Rosenstone (1994) note that their sample selection procedures might have influenced their results.

5. Though the approach used by Holbrook and Krosnick (2003) was aimed at reducing social desirability pressures, the question wording differs substantially from the wording used by the NES.

However, there is an example of success. Recently, Belli et al. (1999) tackled the problem of over-reporting by treating social desirability and source confusion as forces that might act together. Adding both an introductory text designed to reduce memory failures and “excuse” answer choices (the same as used in the 2000 and 2002 NES) that are more socially acceptable than a simple “no,” they found that over-reporting was reduced by about 9 percentage points. What remains unclear is what the separate effects of the enhanced introductory text and the additional answer choices are. Our results speak to this question since the NES changed their response options, but left their introductory text as it was.

Hypotheses

Our hypotheses are guided by recent work that demonstrates the difficulty of isolating the separate effects of social desirability and source confusion (Belli, Traugott, and Rosenstone 1994; Belli et al. 1999) and recognition of the failures of previous survey-based approaches to the reduction of over-reporting. That is, we agree that social desirability pressures and source confusion may exist simultaneously and acknowledge that survey-based experiments that could fully disentangle these effects have yet to be designed. Though our ability to explain the degree to which each of these explanations is responsible for a given result is limited, we sketch our general expectations in the following text and later discuss the totality of the results in light of these expectations.

According to the social desirability hypothesis, those who have not voted but who are strongly aware of the social desirability of being a voter should be more likely to fall into one of the new answer choice categories. This awareness, while widespread, should be especially concentrated among those who are most interested, engaged, and have the most to gain from the system (Silver, Anderson, and Abramson 1986). Of course, if the pressure is sufficiently strong, much more than the addition of two socially acceptable excuses will be necessary to overcome the urge to falsely report having voted. If this is the case, we would expect that those with little to moderate stake in the system and those with low to moderate levels of engagement would be more likely to choose one of the new answer categories, but those with higher levels will be unaffected.

Similarly, those who have voted in the past, but cannot accurately separate recent behavior from previous behavior should find solace in the new answer choice which allows them to count themselves as a “usual” voter. Because people in this category are likely to have high levels of resources and be interested in politics, if the new answer choices help reduce source confusion, we again would expect to see that those most likely to vote are attracted to the new answer choices. However, without the expanded introductory text

employed to jog respondents' memories (Belli, Traugott, and Rosenstone 1994; Belli et al. 1999) the answer choices fit more closely with a social desirability explanation. Evidence to suggest the answer choices led to a reduction in source confusion would be most apparent if over-reporting is lowest among those who were interviewed long after the election, when memories are more likely to fade.

Results

RECENT RATES OF OVER-REPORTING

The relevance of the 8-point drop in turnout that resulted from the new NES question can be seen by placing it in the context of several recent elections. Consider first the turnout rate in the three most recent midterm elections (table 1).⁶ While turnout across these elections was quite similar, differing by <3 percentage points, turnout over-reporting using the traditional NES questions varied from 14 to 25 points, a swing four times larger than the fluctuation in turnout.⁷ The largest rate of over-reporting (25%) was found in 2002 among those who were asked the traditional NES question. Even with the additional answer choices, the rate of over-reporting in 2002 (17%) was higher than in both 1994 and 1998. Explaining the increase in over-reporting that results from the traditional NES question in 2002 is beyond the scope of this article.⁸

6. We use official returns from the Federal Election Commission with the Voting Eligible Population (VEP) as the base (McDonald and Popkin 2001). Calculations using the Voting Age Population (VAP) show turnout to be 2 to 3 points lower than the VEP estimates.

7. Our decisions regarding measurement of over-reporting are made out of necessity. That is, without the virtue of a validation study we must treat differences in self-reported turnout and official turnout rates as the result of over-reporting. Other factors such as response rates, contact rates, testing, interview mode, and panel mortality could all boost the NES turnout rate (Burden 2000, 2003; Martinez 2003; McDonald 2003), but reliable methods to adjust fully for all of these factors have not been developed and doing so is beyond the scope of the present endeavor. We do, however, employ the NES provided weights designed to account for some of these features of the data. Moreover, given the experimental design, to the extent that these forces make a difference they are expected to be equally present in each half-sample, and thus should under or over estimate turnout equally across conditions. Though the new question might lead to more false-negatives, unfortunately, without a validation study this possibility cannot be tested. In sum, we recognize that other forces might be at work, some of which we are simply unable to address, but with the experimental design and appropriate weighting we can mitigate several of these concerns and thus treat lower rates of turnout from the new question as more accurate.

8. The 2002 NES was conducted entirely by telephone, included pre- and post-election interviews as well as panel and non-panel cases. Thus, these factors might influence differences in the rate of over-reporting when 2002 is compared with other midterm elections. Due to the advantage of the experimental design, which forms the basis for the remainder of our results, these issues are

OVER-REPORTING BY INDIVIDUAL CHARACTERISTICS

That the overall rate of over-reporting was reduced is important on its own, especially in light of the numerous attempts that have failed. But for those interested in understanding political behavior, the degree to which over-reporting varies across individuals is a more central concern. As previously noted, the literature indicates that over-reporting is not constant across groups of people with different demographic and attitudinal characteristics. By virtue of the 2002 NES' experimental design we can study turnout across versions of the question and individual characteristics. Using a series of simple probit models, we begin with an investigation of the degree to which the strength of the relationship between turnout and various demographic and attitudinal variables differs across the two turnout questions. Next, using difference of means tests, we turn to a deeper analysis that compares turnout rates across treatments to determine for whom the new question reduces over-reporting and by how much. That is, we isolate individual characteristics and compare reported turnout rates for those who received the new question with those who received the traditional question. An examination of the characteristics of those who select each answer choice, with a focus on who selects the "excuses," follows. We conclude with an analysis of the extent to which our understanding of turnout, using probit models that control for demographics and attitudes, changes when the probability of voting is modeled using the new turnout question. See Appendix A for variable coding.

Table 2 presents the results from a series of bivariate analyses of the overall effect of a standard set of demographic and attitudinal variables on turnout, as measured with the traditional and new question wording.⁹ We report the slope coefficient for the respective variables when using the traditional NES question (control), the slope coefficient when using the new NES question (treatment), the difference between the two slopes, and *p*-value for the test that this difference is equal to zero.

common to both those who receive the traditional and new turnout question and thus do not impact the comparisons across these conditions. For a lively debate on the factors leading to the rise in over-reporting over time see Burden (2000, 2003), McDonald (2003), and Martinez (2003).

9. The results were generated by pooling respondents from each treatment and running a probit estimation on a binary dependent variable set to 1 for those who reported voting and set to 0 for those providing a valid response that indicated that they did not vote, with the independent variable of interest, a dummy variable for the treatment, and an interaction of the independent variable and the treatment dummy. The *p*-value reported in table 2 is the significance level of this interaction term. Of course, equivalent results are obtained when the probit model is estimated separately for the two groups with the traditional and new (coded as vote or did not vote) questions as the respective dependent variables and the independent variable of interest as the only right hand side variable.

Table 2. Bivariate Relationships Between Demographic and Attitudinal Variables and Voting by Treatment (Source: 2002 NES (weighted))

	Control Coeff.	Treatment Coeff.	Difference	<i>p</i> -value
Demographics				
Age	0.02	0.03	-0.01	0.66
Education	0.21	0.27	-0.06	0.26
Income	0.09	0.18	-0.09	0.03
Race (White)	0.19	0.30	-0.11	0.52
Sex (Male)	0.37	0.24	0.13	0.37
Homeownership	0.53	0.74	-0.21	0.19
Years in the Community	0.14	0.18	-0.04	0.63
Interview Week	-0.24	-0.18	-0.06	0.55
Attitudes				
Partisan Strength	0.70	0.46	0.24	0.07
Campaign Interest	1.05	0.95	0.10	0.44
Care Who Wins House Election	0.43	0.60	-0.18	0.05
Information Level (Interviewer assessment)	0.27	0.43	-0.16	0.03
External Efficacy	0.56	0.78	-0.21	0.25
Internal Efficacy	2.12	1.91	0.21	0.49

NOTE.—Control refers to the traditional NES turnout question.

Treatment refers to the new NES turnout question with additional answer choices.

The results were generated by pooling respondents from the control and treatment conditions and running a probit estimation on a binary dependent variable set to 1 for those who reported voting and set to 0 for those providing a valid response that indicated that they did not vote, with the independent variable of interest, a dummy variable for the treatment, and an interaction of the independent variable and the treatment dummy. The *p*-value reported is the significance level of this interaction term.

For the most part, the results in table 2 reveal few significant differences in the effect of the independent variables across the two NES turnout questions. Though the coefficients on the demographic variables in the treatment condition tend to be larger, the only variable for which a statistically significant difference appears is income. In the experimental condition the coefficient for income is twice as large as in the control condition. Thus, in the treatment condition, as income increases, the probability of reporting turnout increases at a faster rate.

The results with respect to attitudes, in terms of the strength of the relationship under the new question, as opposed to the traditional question are mixed. Based on the size of the difference between the coefficients, stronger relationships are found with the new question for caring who wins the

House election, information level,¹⁰ and external efficacy, but only caring who wins the election and information level reach statistical significance. This suggests that turnout rates increase at a faster rate with the new question as the amount one cares and the amount one appears to know about politics increase, respectively. For partisan strength, a weaker relationship is obtained with the new question. This result is driven by those who describe themselves as strong Democrats, for whom the turnout gap across conditions was 7 times larger than the turnout gap across conditions among strong Republicans.

The results in table 2 provide an important first step in our attempt to understand who votes across the experimental conditions but these simple models potentially conceal vital information for assessing the effectiveness of the new question. Further evaluation of the new question requires an examination of which types of individuals are less likely to over-report with the new question and what, if any, new insights into turnout rates across subpopulations can be gained. Table 3 presents the self-reported turnout rate across the two versions of the question, the difference in turnout between the control (traditional wording) and treatment (new wording) conditions, and the *p*-value (based on a two-tailed *t*-test) for this difference by individual characteristics.¹¹

The most striking result found in table 3 is that some of the largest differences are found among the least likely to vote. That is, some of the largest differences and those which are statistically significant appear among respondents who are young, have low levels of education, do not own homes, are new to their community, care little about the outcome of the House election, have low levels of political knowledge, and do not score highly on the efficacy scales. The pattern is somewhat noisy and we uncover several exceptions, but the new question's tendency to affect disproportionately the least educated, the least knowledgeable, and the least interested in politics is distinct.

The new question's effect among the least likely to vote has important normative implications. The new question provides much lower estimates of turnout among those least likely to vote than are obtained with the traditional question, with turnout rates dropping into the low double and even single digits. These very low estimates of turnout for some groups, when they are asked the new question, reveals that for scholars and reformers who treat low turnout as a problem, the degree of under-representation for these groups is far worse than previously thought.

10. Preferably, the respondents' knowledge would be assessed through questions that tap factual information about politics. These sorts of questions were not included in the 2002 NES. Therefore, we use the interviewer's assessment of the respondent's knowledge. Although factual questions were included on the 2000 NES, using 2000 knowledge would mean losing all cases from the 2002 fresh cross-section.

11. That is, we calculated the mean turnout rate across each condition for each category of our independent variables and performed a difference of means test.

Table 3. Turnout Across Conditions by Demographic and Attitudinal Factors, 2002 NES (weighted)

Demographics	Control	Treatment	Difference	<i>p</i> -value
Age				
18–24	0.34	0.16	0.18	0.141
25–30	0.57	0.30	0.27	0.012
31–39	0.57	0.61	–0.04	0.505
40–49	0.70	0.65	0.04	0.414
50–59	0.80	0.71	0.09	0.062
60–64	0.84	0.54	0.30	0.001
65 and up	0.82	0.75	0.07	0.134
Education				
0–8 Years ^a	0.60	0.54	0.06	0.732
9–11 Years	0.49	0.27	0.21	0.091
High School Grad	0.59	0.45	0.14	0.007
1–3 College No Degree	0.58	0.57	0.01	0.900
Junior College Grad	0.63	0.60	0.03	0.696
College Grad	0.87	0.82	0.05	0.220
Advanced Degree	0.87	0.83	0.05	0.399
Income				
\$0–\$14,999	0.52	0.33	0.19	0.038
\$15,000–\$34,999	0.61	0.47	0.14	0.017
\$35,000–\$49,999	0.64	0.54	0.10	0.149
\$50,000–\$64,999	0.73	0.53	0.20	0.006
\$65,000–\$84,999	0.70	0.75	–0.04	0.495
More than \$84,999	0.70	0.70	0.00	0.971
Race				
Non-White	0.59	0.48	0.11	0.063
White	0.66	0.60	0.07	0.021
Sex				
Female	0.59	0.52	0.07	0.055
Male	0.73	0.62	0.11	0.005
Homeownership				
Non-Homeowners	0.50	0.36	0.15	0.014
Homeowners	0.70	0.64	0.06	0.040
Years in the Community				
Less than 1 Year	0.61	0.39	0.22	0.108
1–2 Years	0.55	0.48	0.08	0.446
3–5 Years	0.52	0.48	0.04	0.653
6 Years or Longer	0.68	0.60	0.08	0.005
Interview Week				
First Week	0.73	0.61	0.11	0.000
Second Week	0.45	0.51	–0.06	0.273
Three or Later Weeks	0.66	0.49	0.17	0.029

Continued

Table 3. Continued

Attitudes	Control	Treatment	Difference	<i>p</i> -value
Partisan Strength				
Independent	0.42	0.30	0.12	0.261
Weak	0.60	0.55	0.05	0.166
Strong	0.86	0.70	0.16	0.000
Campaign Interest				
Not Much Interested	0.25	0.16	0.09	0.137
Somewhat	0.65	0.59	0.07	0.075
Very Much Interested	0.92	0.83	0.09	0.003
Care Who Wins House Election				
Not at all	0.42	0.11	0.31	0.013
Not very much	0.53	0.33	0.20	0.000
Pretty Much	0.63	0.67	-0.04	0.374
Very Much	0.86	0.75	0.11	0.004
Information Level (Interviewer assessment)				
Low ^a	0.39	0.05	0.34	0.048
2	0.44	0.22	0.22	0.033
3	0.57	0.48	0.09	0.080
4	0.73	0.68	0.05	0.241
High	0.73	0.71	0.02	0.718
External Efficacy				
Least Efficacious	0.57	0.38	0.19	0.005
0.25	0.45	0.51	-0.06	0.554
0.5	0.60	0.52	0.08	0.154
0.75	0.67	0.66	0.01	0.906
Most Efficacious	0.74	0.68	0.06	0.120
Internal Efficacy				
Least Efficacious ^b	0.42	0.16	0.26	0.045
0.17 ^a	0.24	0.18	0.05	0.724
0.33	0.49	0.46	0.04	0.530
0.50	0.54	0.45	0.09	0.177
0.67	0.72	0.71	0.01	0.886
0.83	0.89	0.74	0.15	0.009
Most Efficacious	0.94	0.84	0.10	0.016

NOTE.—Control refers to the traditional NES turnout question.

Treatment refers to the new NES turnout question with additional answer choices.

Significance levels are based on *t*-tests (two-tailed).

^aIndicates that the *N* is <20 in the control and treatment conditions respectively.

^bIndicates that the *N* is <20 in the control condition.

Furthermore, the results presented in table 3, when examined against the analysis of validated vote results presented by Belli, Traugott, and Beckmann (2001), suggest that new biases may be introduced by the new turnout question. In essence, Belli, Traugott and Beckmann's findings comparing validated voters to self-reported voters indicate that over-reporters tend to be similar to actual voters in terms of their demographic and attitudinal characteristics. To the extent that a pattern emerges in the data presented in table 3, it is the opposite of the pattern Belli, Traugott and Beckmann's findings would lead us to expect. This means that even as the new turnout question helps uncover the extent to which the least well-off fail to vote, it might exaggerate the extent to which the most highly educated and politically knowledgeable vote. With these two general points in mind, a more detailed discussion of the results in table 3 follows.

The results for age¹² demonstrate both the overall pattern and one of the exceptions to it. NES respondents under 30, who are least likely to be voters, seem to be much more affected by the new question than their counterparts between the ages of 30 and 60. Reported turnout for 18–24-year-olds is more than cut in half, to 16% by the new question, a substantively interesting drop even though it is not statistically significant, and reported turnout for respondents 25–30 is cut approximately in half as well, to 30%. Thus, the new question paints a very different picture of voting among young people in 2002. Comparing again across conditions, for none of the age groups between 30 and 60 is the proportional drop in reported turnout nearly so large as it was for those under 30. An odd exception occurs among the 60–64 group, among whom those asked the new question reported voting at only 54%, as compared with 84% for those answering the traditional question.¹³

The results for education also follow the general pattern, while for income a slight deviation is apparent. With respect to education, the most movement occurs among those who least resemble expected voters, with turnout dipping as low as 27% among individuals who attended but did not graduate from high school when receiving the new question as opposed to 49% with the traditional question. With regards to income, the new question's effectiveness in reducing over-reporting sneaks well up into the middle class, but evaporates among those making \$65,000+ annually.¹⁴

12. Age is specified as a continuous variable for table 2, but here is split into meaningful categories for ease of exposition.

13. This might have something to do with people in this age group relocating for retirement. Although we cannot be certain of this explanation, this age category contains a large number of respondents who "usually" vote, and might have missed voting in the 2002 election because of the burden of re-registering, and were happy to seize upon that excuse. Just over 28% of 60–64-year-olds in the experimental condition reported that they "usually" vote as opposed to 10% who "thought" about voting.

14. The results for homeownership, and years in the community also suggest effects among the least and most likely to vote.

One of the most regular findings of previous research is that non-whites are more likely to over-report. Seemingly, the pressure among non-whites to conform to the idea that voting is a mechanism to improve their standing (Belli, Traugott, and Beckmann 2001), or prove they are not outsiders keeps many from falling into the new answer choices.

The second half of table 3 presents the results by attitudinal characteristics. Looking at these attitudinal measures some of the largest differences occur among the least likely to vote, resulting in some of the lowest overall turnout rates. However, the pattern is not particularly distinct. The results for partisan strength and interest in the campaign, for example, do not appear to follow this, or any, pattern; and the results for feelings of external and internal efficacy are mixed. But when we turn to “care who wins” and the interviewer’s assessment of the respondent’s knowledge about politics, the pattern re-emerges.

Those who care little about the House election are dramatically less likely to report voting when asked the new “excuse” question, (53% with the traditional question versus 33%) and those who don’t care at all drop even more (42–11%). For those who care very much, the drop (86–75%) is a good deal smaller. Even more striking differences are found when political knowledge, as assessed by the interviewer, is considered. The effect of the new turnout question is extraordinary among the least informed, whose reported turnout drops from 39% with the old question to just 5% with the new question, while the new question makes virtually no difference among respondents who the interviewers assessed as well-informed.¹⁵ For reformers interested in higher and more equal rates of participation, these results will be rather disheartening. However, for people who would prefer an informed and engaged electorate, the results suggest that very few of the least informed and engaged cast ballots in 2002.

We have highlighted a particular pattern of results that emerges from the data presented in table 3: in large part, the new turnout question has the most impressive effects on the reported turnout of those whose demographic and attitudinal characteristics make them most unlike the typical habitual voter. What makes this pattern of results especially notable, and worth highlighting despite the exceptions that emerge, is that Belli, Traugott, and Beckmann’s (2001) validation study would lead us to expect precisely the opposite pattern. If the new question had an equal likelihood of influencing any given respondent to give an accurate response to the turnout question, we would expect the most dramatic effects to be among the most well-off, the most educated, the most interested, the most knowledgeable and the

15. Although the *N* is somewhat small in lowest information category, when the two lowest information levels are combined, the reported rate of turnout remains low in the treatment condition (18%) and is 25 percentage points lower than reported turnout in the control condition; the 25 point difference is statistically different from 0 at $p = 0.005$.

most efficacious, since these are the groups that Belli, Traugott, and Beckmann found most frequently answered the turnout question inaccurately. Given their results, if we had found no pattern at all in terms of which groups were most affected by the turnout question that would have pointed to the possibility that the new question is introducing bias. The fact that to the extent there is an identifiable pattern, it is the opposite direction of the pattern that Belli, Traugott, and Beckmann's research leads one to predict seems to us to introduce a point of caution for those who study turnout using the new NES question. The new question seems to affect disproportionately those over-reporters who are least similar to the typical voter, and leave those over-reporters most similar to typical voters relatively unaffected.

To test the effect of the new turnout question on over-reporting resulting from "source confusion," we examine the impact that the length of time between the election and the interview had on respondents' reported turnout. Respondents who were interviewed more than 2 weeks after the election had a turnout rate 17 points lower in the treatment condition; while those interviewed within the first week had a turnout rate 11 points lower in the treatment condition. However, this difference is not statistically significant. Thus, the new answer choices may help reduce over-reporting among respondents who might have trouble remembering whether or not they voted in the most recent election, but it is difficult to say this definitively.

Given that the answer choices seem to tap social desirability pressures more directly than source confusion, and since unlike the findings of Belli et al. (1999) (who incorporated both a longer introductory text and the "excuse" answer choices), we failed to find strong evidence that over-reporting decreased among those interviewed long after the election, we are persuaded to believe that the reduction in social desirability pressures rather than source confusion is the driving force behind decreases in over-reporting that result from the new question wording. As noted above, this conclusion should be viewed as preliminary until additional theoretical and empirical work can be done to tease out the separate effects of these two mechanisms. Since Belli et al.'s study (1999) found greater effects as time from the interview increases, there may be distinct effects of a longer introductory text versus new answer choices.

In sum, the additional answer choices do not seem to be enough to reduce turnout reports among those who feel the strongest pressure to produce a socially acceptable response, or those who are likely to have voted in the past but who have murky memories. Given that these types of individuals are not widely affected by the new response options, it is not surprising that survey-based experiments on validated registrants (Abelson, Loftus, and Greenwald 1992; Belli, Traugott, and Rosenstone 1994), a group with a mix of resources, attitudes, and previous behavior that would predict a high likelihood of voting, did not reduce over-reporting. For those who are least likely to care about the election and whose demographic characteristics

suggest they probably vote infrequently, if at all, appearances seem to be enough; i.e., these voters do not feel sufficient pressure to report the most desirable response and more willingly latch onto the new answer choices. Not only do the results suggest that the new answer choices cannot fully overcome the social desirability pressure or constraints on memory, but these factors appear more prevalent among the least engaged than previously thought. Perhaps the most consequential result is that the levels of turnout found among citizens with low levels of resources and little attachment to politics point to the possibility of a more severe degree of inequality in regards to who votes and who does not.

COMPARISON OF TRAITS ACROSS ANSWER CHOICES

Having investigated the types of citizens who are less likely to falsely report when they can report non-voting while still appearing to be civic minded, we now turn to a comparison of individuals across answer choice categories. That is, we examine who says “*I thought about voting this time but didn’t*” and who says “*I usually vote but didn’t this time*” in comparison with reported voters and those who chose the traditional did not vote response (“*No, I didn’t vote*” or “*I did not vote*” in the traditional and new NES questions, respectively). The profile of the respondents who opted for the excuses is substantively interesting: how similar or distinctive are they from those who claimed to have voted and those who readily admitted to abstaining?

The results are presented in table 4, which shows the mean values for a number of independent variables traditionally associated with turnout across the response categories in the control and experimental conditions.

What makes it difficult to interpret the results in table 4 is that the respondents who fall into the two “excuse” categories are a mix of respondents who would have claimed to vote if asked the traditional question, and those who would have admitted to not voting even if they were not offered an excuse. While the experimental design reveals that the new version of the turnout question reduced over-reporting by 8%, more than 31% of respondents who were asked the new question seized upon one of the “excuse” categories. We might suppose that the “usually” category, which was chosen by 17% of respondents, consists of an equal mix of respondents who would have reported voting and those who would have reported not voting if asked the traditional question. Since previous research (Belli, Traugott, and Beckmann 2001) indicates that over-reporters are similar to voters in terms of their demographics and attitudes, and dissimilar to non-voters, we could expect the “usually” category to fall about half way between those who say they definitively voted, and those who say they did not vote without the luxury of an excuse. This is not what we found.

Table 4. Mean Values of Predictors by Experimental Condition and Voting Outcome

	Control Group		Experimental Group				Total
	Did Not Vote	Voted	Did Not Vote	Thought About Voting But Did Not	Usually Vote But Did Not	Sure Voted	
Number of Cases ^a	172	500	61	53	96	462	1344
Weighted Number of Cases	236	436	79	97	114	382	1344
Weighted Percentages	35.07%	64.93%	11.72%	14.37%	17.03%	56.87%	
Demographics							
Age	37.59	47.99	45.05	31.87	40.06	48.55	44.32
Education	3.44	4.19	3.16	3.37	3.78	4.48	3.99
Income	3.30	3.68	2.85	3.24	3.31	3.96	3.58
White (%)	0.74	0.79	0.77	0.59	0.76	0.79	0.77
South (%)	0.40	0.29	0.36	0.34	0.37	0.28	0.32
Male (%)	0.32	0.46	0.43	0.43	0.38	0.51	0.44
Own Home (%)	0.64	0.80	0.75	0.56	0.54	0.83	0.74
Political Characteristics							
Party ID	2.90	2.99	2.61	2.46	2.81	3.06	2.92
Partisan Strength	0.98	1.34	1.16	1.14	1.07	1.34	1.23
Interest	1.65	2.31	1.61	1.64	1.87	2.33	2.08
Care Who Wins	2.55	3.03	2.41	2.41	2.69	3.16	2.87
Int. Efficacy Scale	0.44	0.65	0.36	0.41	0.54	0.66	0.57
Ext. Efficacy Scale	0.53	0.65	0.36	0.48	0.59	0.68	0.60
Knowledge (IWER Assess)	3.39	3.82	3.00	3.26	3.57	3.96	3.67
Contextual Variables							
Interview Day	8.25	7.01	7.12	9.23	8.62	6.97	7.52
Interview Week	1.64	1.44	1.53	1.65	1.60	1.45	1.51

^aThe numbers represent the cases of the dependent variable before conditioning on the independent variables: due to item non-response, the N is smaller in some cells.

Rather, table 4 indicates that even respondents in the “usually” category are more similar to non-voters than voters on most measures.¹⁶

When examining a number of variables associated with turnout, we found a pattern that does not parallel the findings of Belli, Traugott, and Beckmann (2001). Those researchers found that in the categories of education, caring about the election, efficacy, interest in the election, partisan strength, and knowledge, over-reporters “are significantly closer to validated voters than to admitted non-voters” (Belli, Traugott, and Beckmann 2001, 489). Turning to the 2002 NES and comparing the respondents who seized upon the new excuse categories to those who report voting and those who chose the simple “did not vote” category, the “excuse” respondents more closely resemble those who report they simply did not vote, in both the control and experimental conditions.¹⁷

For example, non-voters who “thought about voting,” when compared with those who report voting in the control condition, have significantly lower levels of education, less interest in the election, care less about the outcome, have lower levels of efficacy, and less political knowledge. When the same “thought about voting” respondents are compared with the group who answered the traditional turnout question by saying they did not vote, the differences in the means are smaller, and none of the differences are statistically significant.

A similar pattern emerges even when comparing those non-voters who said they “usually vote” to respondents asked the traditional turnout question. The “usually” group is significantly different from reported voters in their partisan strength, interest in the election, caring who wins, and internal efficacy. When compared with non-voters in the control group, however, the “usually” group is only statistically distinguishable on internal efficacy.

This general pattern repeats itself when one looks at only those respondents who were asked the new version of the NES question. Comparing respondents in the excuse categories to those who said they are sure they voted produces a host of statistically significant differences. That is, self-reported voters are older, have higher levels of education, have larger incomes, are more likely to own homes, have greater interest in the campaign, care more who wins, and are more knowledgeable about politics than

16. While this result, when viewed along with the results presented in table 3, might raise concerns about systematic bias in who is affected by the new version of the turnout question, the ambiguity introduced by the mix of those who would have admitted not voting and over-reporters in the excuse categories makes it difficult to demonstrate definitively the degree to which bias is introduced. Our estimates serve as a baseline, leaving this question to be addressed more fully in future research that incorporates a validation study into the design.

17. Table A2, available in the appendix to the on-line version of the article, reports the mean differences for each comparison and results from tests of statistical significance.

those who fall into the excuse categories. Comparing the excuse respondents with those who simply say they did not vote, however, reveals very few differences that rise to conventional levels of statistical significance. The one notable exception to this rule is that those non-voters who say they “usually vote” are more similar to voters than they are to simple (no excuse) non-voters in their levels of internal efficacy.

When comparing respondents across the excuse options we see little to distinguish those who select the “thought about voting” option from those who select the “usually vote” option. We found only 2 statistically significant results: the average age among respondents who “usually vote” is 8 years older ($p=0.037$) than those who “thought about voting;” and “usual” voters score slightly higher on the internal efficacy score ($p=0.052$).

UNDERSTANDING TURNOUT

The above analysis exploited the experimental design to cull what can be discovered by focusing on bivariate relationships. We now investigate the extent to which the new question alters our understanding of who votes in a multiple regression framework. The number of model specifications that have been used in previous research to predict turnout is too large to summarize here; for the sake of simplicity, our turnout model is based on the specification devised by Rosenstone and Hansen (1993).

Table 5 presents the results for a probit model, run separately for those asked the traditional NES question and those asked the new version, which includes demographic and attitudinal variables along with a variable to indicate whether or not the respondent had been contacted by a political party about voting.¹⁸ The two columns on the left represent the coefficients and p -values for the model on respondents with the traditional question, and the last two are corresponding results from the experimental group.

Across both conditions, education, age, caring who wins the election, partisan strength, being contacted by a party, and residency in the south are significant predictors of turnout. For education, age, and caring who wins, the size of the coefficient was larger for the experimental question; in the case of caring who wins the election, the coefficient more than doubled.

Several differences in terms of sign and statistical significance are also apparent. In the treatment condition, the sign on the income coefficient is positive, but the effect cannot be distinguished from zero. Surprisingly,

18. We obtain similar results when all respondents are pooled, and we include interactions between the predictor variables and a dummy variable indicating whether the respondent received the new turnout question or the traditional question.

Table 5. Modeling Turnout by Condition (Demographics and Attitudes) (Source: 2002 NES (weighted))

Independent Variables	Control coeff.	control <i>p</i> -value	Treatment coeff.	Treatment <i>p</i> -value
Income	-0.096	0.042	0.053	0.246
Education	0.188	0.000	0.248	0.000
Unemployed	-0.304	0.324	-0.265	0.483
Age	0.048	0.029	0.080	0.000
Age Squared	0.000	0.103	-0.001	0.003
Years in the Community	0.021	0.778	0.022	0.783
Church Attendance	0.032	0.529	0.002	0.972
Home Ownership	0.153	0.337	0.389	0.015
Border South	-0.286	0.335	0.131	0.613
Southern State	-0.259	0.085	-0.325	0.030
Race - Black	0.032	0.883	-0.007	0.973
Registration Date	-0.003	0.699	0.000	0.968
External Efficacy	0.257	0.140	0.288	0.080
Partisan Strength	0.688	0.000	0.370	0.002
Care Who Wins	0.136	0.101	0.421	0.000
Contacted by Party	0.754	0.000	0.550	0.000
Constant	-2.956	0.000	-5.399	0.000
Number of Cases	589		595	
Log Likelihood	-287.66		-289.45	

the effect of income is negative and statistically significant in the control condition.¹⁹

Home ownership and external efficacy are not significant predictors for the traditional question but are significant when modeling the experimental question. For external efficacy, the magnitude of the change in the coefficient is actually very small, but the variable nonetheless is significant at the $p < 0.10$ level in the experimental question. The change in the coefficient for home ownership across the two models is larger—it more than doubles.

Although testing the differences in the probit estimates from the two groups is not straightforward, we compared the marginal effects of the

19. The bivariate correlation between turnout and income in the control group is 0.10 and is weaker than that in the treatment group (0.22). This is consistent with and can be explained by our earlier findings that there is a disproportionately large number of control group false reporters at the lower end of the income scale, compared with the treatment group. This over-report problem in the control group is strong enough to make the coefficient on income turn significantly negative when controlling for other demographic variables, most importantly, education. Although we want to stress that the new NES question does not in any way fully capture the true turnout rate, we think this provides a good example of one of the problems over-reporting using the traditional NES question may create.

independent variables across conditions.²⁰ Our tests revealed that only two of the independent variables have different effects across conditions: respondent income and the degree to which the respondent cares who wins. Though there is a sizeable difference for care who wins, on the whole, the new NES question does not do much to revise our knowledge of the effect of various demographic and attitudinal variables on turnout when looked at in a multivariate context.

Conclusion

Ultimately, we believe these excuses are quite good, and that it would be best to include the new response categories on future versions of the NES survey and other post-election surveys. The new question basically works in that it significantly reduces over-reporting, and furthermore does so without the inclusion of a lengthy and cumbersome introductory script. For most scholars who make use of NES respondents' answers to the turnout question, this is a significant improvement that is unlikely to introduce further complications. The new question does not get perfect information, but what it does seem to do is reduce precisely the false reports of voting in the data that are likely to be the most disharmonious and the most likely to obscure the overall theme. Over-reporters who are very dissimilar to habitual voters tend to seize on the excuses.

What over-reporting the new turnout question does remove will allow researchers to highlight better and explore what is undoubtedly among the most significant findings regarding who votes and who does not vote in the U.S.—that those Americans who lack resources, and the confidence that they can understand politics and participate meaningfully, vote at astonishingly low rates. As a practical consequence of these findings, advocates of higher participation rates should allocate additional resources toward educating the citizenry about politics and the role of the citizen in a democratic society (Hanmer 2004). Efforts that fail to do so are unlikely to overcome the abysmal turnout rates found among those who have little concern for the outcome or little knowledge of politics.

However, for researchers who want to understand not merely who typically votes, but to dig deeper into the question of which occasional, and nearly habitual voters mobilize or de-mobilize in particular elections, in response to particular sorts of campaigns, or despite particular obstacles etc., the new

20. We computed the marginal effect of each independent variable, holding the other variables at their means (cardinal variables) or their medians (dummy variables). The “delta method” provides standard errors of fitted probabilities, and we checked whether the estimated effects in one condition fell within the confidence interval of their counterparts in the other condition. See Herron (1999) for more details.

question does not offer much new hope. The same can be said for those scholars interested in the group of citizens that resembles habitual voters in attitudes and demographics and that nonetheless does not vote. NES data for respondents in both of these groups will continue to be problematic, even with the new question.

To the extent that the over-reporting of voting is a result of social desirability, our analysis seems to indicate that a sense of the social desirability of being a voter is pervasive. But our analysis shows that this widespread social desirability effect is very deep in some spots and very shallow in others. The new NES response categories seem to be very good at mopping up the shallow spots—for example, among those who do not care who wins and have very little political knowledge, reported turnout drops dramatically (more than 75%) with the new question. Since the factors we expected to be associated with social desirability were not good predictors of whether respondents were affected by the new version of the turnout question it does not seem that the new question does consistently well where the sense of social desirability runs deep.

Inclusion of the new response categories are probably a cost effective substitute for validation studies, which represent a time consuming and often problematic approach to finding out which respondents *really* voted. However, we believe it would be worthwhile to combine the new question with a validation study at least once, especially as states implement electronic voter files that will make validation less costly. Our analysis leads us to predict that the likely result would be similar to the findings of Belli, Traugott, and Beckmann (2001). When using the new turnout question, there will be fewer over-reporters, but those who remain will resemble confirmed voters much more than admitted non-voters. With the new question this pattern should be even more pronounced. Our analysis indicates that this would be the case, but using the new question in combination with a validation study would be the best test.

Variable Coding

Age: in years 18–99.

Black: 1 = black only, 0 = all else.

Border South: 1 = lives in a border state (KY, MO, MD, OK, WV), 0 = all else.

Care Who Wins: 1 = Not at all, 2 = Not very much, 3 = Pretty Much, 4 = Very Much.

Church Attendance: 1 = attends church, 0 = does not.

Contacted by Party: 1 = was contacted by a political party, 0 = was not.

Education: 1 = 0–8 Years, 2 = 9–11 Years, 3 = High School Grad, 4 = 1–3 College No Degree, 5 = Junior College Grad, 6 = College Grad, 7 = Advanced Degree.

External Efficacy Scale: ranges from 0 to 1 (average of answers to 2 questions).

Home Ownership: 1 = owns a home, 0 = does not.

Income: 1 = \$0–\$14,999, 2 = \$15,000–\$34,999, 3 = \$35,000–\$49,999, 4 = \$50,000–\$64,999, 5 = \$65,000–\$84,999, 6 = more than \$84,999.

Interest: 1 = Not much interested in the campaign, 2 = Somewhat, 3 = Very Much Interested.

Internal Efficacy Scale: ranges from 0 to 1 (average of answers to 3 questions).

Interview Day: ranges from 1–31 (number of days after the election that the interview was conducted).

Interview Week: 1 = within the first week, 2 = within the second week, 3 = within the third or fourth weeks (number of weeks after the election that the interview was conducted).

Knowledge (Interviewer Assessment): ranges from 1 to 5 (1 = low, 5 = high).

Male: 1 = male, 0 = female.

Party ID: 0 = strong democrat, 1 = weak democrat, 2 = independent-democrat, 3 = independent-independent, 4 = independent-republican, 5 = weak republican, 6 = strong republican.

Partisan Strength: 0 = independent, 1 = weak partisan, 2 = strong partisan.

Registration Date: number of days prior to the election that registration ends (0–30 days).

Southern State: 1 = lives in the South (states of the old Confederacy), 0 = all else.

White: 1 = white only, 0 = all else.

Years in the Community: 0 = <1 year, 1 = 1–2 years, 2 = 3–5 years, 3 = 6 or more.

Supplementary Data

Supplementary data are available online at <http://pubopq.oxfordjournals.org/>.

References

- Abelson, Robert P., Elizabeth F. Loftus, and Anthony G. Greenwald. 1992. "Attempts to Improve the Accuracy of Self-Reports of Voting." In *Questions about Questions*, ed. Judith M. Tanur, pp. 138–153. New York: Russell Sage Foundation.
- Abramson, Paul R., and William Claggett. 1984. "Race-related Differences in Self-reported and Validated Turnout." *Journal of Politics* 46:719–38.
- Anderson, Barbara A., and Brian D. Silver. 1986. "Measurement and Mismeasurement of the Validity of the Self-Reported Vote." *American Journal of Political Science* 30: 771–85.

- Belli, Robert F., Santa Traugott, and Steven J. Rosenstone. 1994. "Reducing Over-Reporting of Voter Turnout: An Experiment Using a Source Monitoring Framework." NES Technical Reports Number 35 (see available technical reports at <http://www.umich.edu/~nes/>).
- Belli, Robert F., Michael W. Traugott, Margaret Young, and Katherine A. McGonagle. 1999. "Reducing Vote Over-Reporting in Surveys: Social Desirability, Memory Failure, and Source Monitoring." *Public Opinion Quarterly* 63:90–108.
- Belli, Robert F., Michael W. Traugott, and Matthew N. Beckmann. 2001. "What Leads to Voting Overreports? Contrasts of Overreporters to Validated Voters and Admitted Nonvoters in the American National Election Studies." *Journal of Official Statistics* 17(4): 479–98.
- Bernstein, Robert., Anita Chadha, and Robert Montjoy. 2001. "Overreporting Voting: Why it Happens and Why it Matters." *Public Opinion Quarterly* 65:22–44.
- Burden, Barry C. 2000. "Voter Turnout and the National Elections Studies." *Political Analysis* 8:389–98.
- Burden, Barry C. 2003. "Internal and External Effects on the Accuracy of NES Turnout: Reply." *Political Analysis* 11:193–5.
- Clausen, Aage. 1968. "Response Validity: Vote Report." *Public Opinion Quarterly* 32: 588–606.
- Hanmer, Michael J. 2004. *From Selection to Election and Beyond: Understanding the Causes and Consequences of Electoral Reform in America*. Ph.D. dissertation, University of Michigan.
- Herron, Michael C. 1999. "Post-Estimation Uncertainty in Limited Dependent Variable Models." *Political Analysis* 8:83–98.
- Hill, Kim Quaile, and Patricia A. Hurley. 1984. "Nonvoters in Voters' Clothing: The Impact of Voting Behavior Misreporting on Voting Behavior Research." *Social Science Quarterly* 45:519–35.
- Holbrook, Allyson L., and Jon A. Krosnick. 2003. "Vote Over-Reporting: A Test of the Social Desirability Hypothesis" Unpublished Manuscript.
- Katosh, John P., and Michael W. Traugott. 1981. "The Consequences of Validated and Self-Reported Voting Measures." *Public Opinion Quarterly* 45:519–35.
- Martinez, Michael D. 2003. "Comment on "Voter Turnout and the National Election Studies." *Political Analysis* 11:187–92.
- McDonald, Michael P. 2003. "On the Over-Report Bias of the National Election Survey." *Political Analysis* 11:180–6.
- MacDonald, Michael P., and Samuel L. Popkin. 2001. "The Myth of the Vanishing Voter." *American Political Science Review* 95:963–74.
- The National Election Studies. *The 1996 National Election Study* [ICPSR 6896]. Ann Arbor, MI: University of Michigan, Center for Political Studies.
- Presser, Stanley. 1990. "Can Context Changes Reduce Vote Over-Reporting?" *Public Opinion Quarterly* 54:586–93.
- Presser, Stanley, Michael W. Traugott, and Santa Traugott. 1990. "Vote 'Over' Reporting in Surveys: The Records or the Respondents?" NES Technical Reports No. nes010157 (see available technical reports at <http://www.umich.edu/~nes/>).
- Rosenstone, Steven J., and Mark J. Hansen. 1993. *Mobilization, Participation, and Democracy in America*. New York: Macmillan.
- Sigelman, Lee. 1982. "The Nonvoting Voter in Voting Research." *American Journal of Political Science* 26:47–56.
- Silver, Brian D., Barbara A. Anderson, and Paul R. Abramson. 1986. "Who Overreports Voting?" *American Political Science Review* 80:613–24.
- Traugott, Michael W., and John P. Katosh. 1979. "Response Validity in Surveys of Voting Behavior." *Public Opinion Quarterly* 43:359–77.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who Votes?* New Haven, CT: Yale University Press.